

Distributed Stochastic Variance Reduced Gradient Methods and A Lower Bound for Communication Complexity

Jason D. Lee · Qihang Lin · Tengyu Ma ·
Tianbao Yang

Received: date / Accepted: date

Abstract We study distributed optimization algorithms for minimizing the average of convex functions. The applications include empirical risk minimization problems in statistical machine learning where the datasets are large and have to be stored on different machines. We design a distributed stochastic variance reduced gradient algorithm that, under certain conditions on the condition number, simultaneously achieves the optimal parallel runtime, amount of communication and rounds of communication among all distributed first-order methods up to constant factors. Our method and its accelerated extension also outperform existing distributed algorithms in terms of the rounds of communication as long as the condition number is not too large compared to the size of data in each machine. We also prove a lower bound for the number of rounds of communication for a broad class of distributed first-order methods including the proposed algorithms in this paper. We show that our accelerated distributed stochastic variance reduced gradient algorithm achieves this lower bound so that it uses the fewest rounds of communication among all distributed first-order algorithms.

Keywords Distributed Optimization · Communication Complexity · Machine Learning · First-Order Method

1 Introduction

In this paper, we consider the *distributed optimization* problem of minimizing the average of N convex functions in \mathbb{R}^d , i.e.,

$$\min_{x \in \mathbb{R}^d} \left\{ f(x) := \frac{1}{N} \sum_{i=1}^N f_i(x) \right\} \quad (1)$$

using m machines. For simplicity, we assume $N = mn$ for an integer n with $m \ll n$ but all of our results can be easily generalized for a general N . Here, $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ for $i = 1, \dots, N$ is convex and L -smooth, meaning that f_i is differentiable and its gradient ∇f_i is L -Lipschitz continuous¹, i.e., $\|\nabla f_i(x) - \nabla f_i(y)\| \leq L\|x - y\|$, $\forall x, y \in \mathbb{R}^d$, and their average f is μ -strongly convex, i.e., $\|\nabla f(x) - \nabla f(y)\| \geq \mu\|x - y\|$, $\forall x, y \in \mathbb{R}^d$. We call $\kappa = \frac{L}{\mu}$ the *condition number* of function f . Note that the function f itself can be L_f -smooth,

Jason D. Lee
Department of Electrical Engineering and Computer Science, UC Berkeley
E-mail: jasondlee@berkeley.edu

Qihang Lin
Tippie College of Business, The University of Iowa
E-mail: qihang-lin@uiowa.edu

Tengyu Ma
Department of Computer Science, Princeton University
E-mail: tengyu@cs.princeton.edu

Tianbao Yang
Department of Computer Science, The University of Iowa
E-mail: tianbao-yang@uiowa.edu

¹ In this paper, the norm $\|\cdot\|$ represents Euclidean norm.

namely, $\|\nabla f(x) - \nabla f(y)\| \leq L_f \|x - y\|$, $\forall x, y \in \mathbb{R}^d$, for a constant $L_f \leq L$. Let x^* be the unique optimal solution of (1) and a solution \hat{x} is called an ϵ -optimal solution² for (1) if $f(\hat{x}) - f(x^*) \leq \epsilon$.

One of the most important applications of problem (1) is *empirical risk minimization* (ERM) in statistics and machine learning. Suppose there exists a set of i.i.d. samples $\{\xi_1, \xi_2, \dots, \xi_N\}$ from an unknown distribution D of a random vector ξ . An ERM problem can be formulated as

$$\min_{x \in \mathbb{R}^d} \frac{1}{N} \sum_{i=1}^N \phi(x, \xi_i) \quad (2)$$

where x represents a group of parameters of a predictive model, ξ_i is the i th data point, and $\phi(x, \xi)$ is a loss function. Note that (2) has the form of (1) with each function $f_i(x)$ being $\phi(x, \xi_i)$. Typically, the data point ξ is given as a pair (a, b) where $a \in \mathbb{R}^d$ is a feature vector and $b \in \mathbb{R}$ is either a continuous (in regression problems) or a discrete response (for classification problems). The examples of loss function $\phi(x, \xi)$ with $\xi = (a, b)$ include: square loss in linear regression where $a \in \mathbb{R}^d$, $b \in \mathbb{R}$, and $\phi(x, \xi) = (a^T x - b)^2$; logistic loss in logistic regression where $a \in \mathbb{R}^d$, $b \in \{1, -1\}$, and $\phi(x, \xi) = \log(1 + \exp(-b(a^T x)))$; smooth hinge loss where $a \in \mathbb{R}^d$, $b \in \{1, -1\}$, and

$$\phi(x, \xi) = \begin{cases} 0 & \text{if } ba^T x \geq 1 \\ \frac{1}{2} - ba^T x & \text{if } ba^T x \leq 0 \\ \frac{1}{2}(1 - ba^T x)^2 & \text{otherwise.} \end{cases}$$

To improve the statistical generalization properties of the model learned from (2), a regularization term $\frac{\lambda}{2} \|x\|^2$ is often added to (2) and the problem becomes a regularized ERM problem

$$\min_{x \in \mathbb{R}^d} \frac{1}{N} \sum_{i=1}^N \phi(x, \xi_i) + \frac{\lambda}{2} \|x\|^2 \quad (3)$$

which still takes the form of (1) with $f_i(x) = \phi(x, \xi_i) + \frac{\lambda}{2} \|x\|^2$. The parameter λ is called an regularization parameter. As argued by [24, 25, 23, 28], for ERM problem, the value of λ is typically in the order of $\Theta(1/\sqrt{N}) = \Theta(1/\sqrt{mn})$.

We consider a situation where all N functions are initially stored in the same large storage space that has limited computation power. We assume that each of the m machines we use to solve (1) has a limited memory space of C so that it can load at most C of the N functions in (1). In the case of ERM, this means each machine can load at most C data points among $\{\xi_1, \xi_2, \dots, \xi_N\}$ in its memory. Since the data point ξ_i uniquely defines f_i in ERM, in the rest of the paper, we will call f_i a data point i or a function i interchangeably.

Throughout the whole paper, we assume that

Assumption 1 *The memory space C of each machine satisfies $n < C < N$ and the quantity $\tilde{n} \equiv C - n$ satisfies $\tilde{n} \geq cn$ for a universal constant $c > 0$.*

The inequality $C < N$ forces us to use more than one, if not all, the machines for solving (1). The quantity \tilde{n} represents the remaining space in each machine after we evenly allocate N data points onto m machines. The inequality $\tilde{n} \geq cn$ means each machine still has $\Omega(n)$ memory space after such an allocation of data. This can happen when either the machine capacity C or the number of machines m is large enough.

We also assume that we can load the same function to multiple machines so that different machines may share some functions, so the sets of functions in all machines do not necessarily form a partition of $\{f_i\}_{i \in [N]}$. Since no machine can access all N functions, we have to solve (1) by distributed algorithms that alternate between a local computation procedure at each machine, and a round of communication to synchronize and share information among the machines.

1.1 Communication efficiency and runtime

To facilitate the theoretical study, we use the following simplified message passing model from the distributed computation literature [9, 5]: We assume the communication occurs in rounds – in each

² If \hat{x} is a random variable generated by a stochastic algorithm, we call it an ϵ -optimal solution if $\mathbb{E}[f(\hat{x}) - f(x^*)] \leq \epsilon$.

round, (a subset of) machines exchanges messages and, between two rounds, the machines only compute based on their local information (local data points and messages received before).

Given this state of affairs, we study the distributed optimization problem with three performance metrics in mind.

- **Local parallel runtime:** The longest running time of m machines spent in local computation, measured in the number of gradient computations, i.e., computing $\nabla f_i(x)$ for any i . We also refer it as “runtime” for simplicity.
- **The amount of communication:** The total amount of communication among m machines and the center, measured by the number of vectors³ of size d transmitted.
- **Rounds of communication:** How many times all machines have to pause their local computation and exchange messages. We also refer it as “rounds” for simplicity.

We will study these performance metrics for the algorithms we propose and compare with other existing techniques. However, the main focus of this paper is the rounds of communication.

1.2 Summary of contributions

In this paper, we first propose a *distributed stochastic variance reduced gradient* (DSVRG) method, which is simple and easy to implement – it is essentially a distributed implementation of a well-known single-machine stochastic variance reduced gradient (SVRG) method [11, 26, 12]. We show that the proposed DSVRG algorithm requires $O((1 + \frac{\kappa}{n}) \log(1/\epsilon))$ rounds of communication to find an ϵ -optimal solution for (1) under Assumption 1. The corresponding parallel runtime is $O((n + \kappa) \log(1/\epsilon))$ and the associated amount of communication is $O((m + \frac{\kappa}{n}) \log(1/\epsilon))$.

Given these performance metrics of DSVRG, we further ask a key question:

How can we achieve the optimal parallel runtime, the optimal amount of communication, and the optimal number of rounds of communication simultaneously for solving (1)?

This paper answers this seemingly ambitious question affirmatively in a reasonable situation: When $\kappa = \Theta(n^{1-2\delta})$ with a constant $0 < \delta < \frac{1}{2}$, with an appropriate choices for the parameters in DSVRG (shown in Corollary 1), DSVRG finds an ϵ -optimal solution for (1) with a parallel runtime of $O(n)$, an $O(m)$ amount of communication and $O(1)$ rounds of communication for any $\epsilon = \frac{1}{n^s}$ where s is any positive constant. Here, the notation O hides a logarithmic term of the optimality gap of an initial solution for DSVRG, which is considered as a constant in the whole paper.

We want to point out that $\kappa = \Theta(n^{1-2\delta})$ is a typical setting for machine learning applications. For example, as argued by [24, 25, 23, 28], for ERM, the condition number κ is typically in the order of $\Theta(\sqrt{N}) = \Theta(\sqrt{mn})$. Therefore, when the number of machines m is not too large, e.g., when $m \leq n^{0.84}$, we have that $\kappa = \Theta(\sqrt{mn}) \leq n^{0.9}$ (so that $\delta = 0.05$). Moreover, $\epsilon = n^{-10}$ (so that $s = 10$) is certainly a high enough accuracy for most machine learning applications since it exceeds the machine precision of real numbers, and typically people choose $\epsilon = \Theta(\frac{1}{N})$ in empirical risk minimization.

These performance guarantees of DSVRG, under the specific setting where $\kappa = O(n^{1-2\delta})$ and $\epsilon = O(\frac{1}{n^s})$, are optimal up to constant factors among all distributed first-order methods. First, to solve (1), all m machines together need to compute at least $\Omega(N)$ gradients [1] in total so that each function in $\{f_i\}_{i=1, \dots, N}$ can be accessed at least once. Therefore, at least one machine needs to compute at least $\Omega(n)$ gradients in parallel given any possible allocation of functions. Second, the amount of communication is at least $\Omega(m)$ for even simple Gaussian mean estimation problems [3], which is a special case of (1). Third, at least $O(1)$ rounds of communication is needed to integrate the computation results from machines into a final output.

Furthermore, using the generic acceleration techniques developed in [8] and [14], we propose a *distributed accelerated stochastic variance reduced gradient* (DASVRG) method that further improves the theoretical performance of DSVRG. Under Assumption 1, we show that DASVRG requires only $\tilde{O}((1 + \sqrt{\frac{\kappa}{n}}) \log(1/\epsilon))$ rounds of communication to find an ϵ -optimal solution, leading to better theoretical performance than DSVRG. Also, we show that the runtime and the amount of communication for DASVRG are $\tilde{O}((n + \sqrt{n\kappa}) \log(1/\epsilon))$ and $\tilde{O}(m + m\sqrt{\frac{\kappa}{n}} \log(1/\epsilon))$, respectively. We also prove a

³ We will only consider communicating data points, or iterates x . For simplicity, we assume that the data point and iterates are of the same dimension, but this can be easily generalized.

⁴ If $n = 10^5$, then $n^{0.8} = 10^4$ which is already much more than the number of machines in most clusters.

Algorithm	Rounds	Parallel Runtime	Assumptions
DSVRG	$(1 + \frac{\kappa}{n}) \log \frac{1}{\epsilon}$	$G(n + \kappa) \log \frac{1}{\epsilon}$	Assumption 1
DASVRG	$(1 + \sqrt{\frac{\kappa}{n}}) \log(1 + \frac{\kappa}{n}) \log \frac{1}{\epsilon}$	$G(n + \sqrt{n\kappa}) \log(1 + \frac{\kappa}{n}) \log \frac{1}{\epsilon}$	Assumption 1
DISCO (quad)	$(1 + \frac{\sqrt{\kappa}}{n^{.25}}) \log \frac{1}{\epsilon}$	$Q(1 + \frac{\sqrt{\kappa}}{n^{.25}}) \log \frac{1}{\epsilon}$	(3), $\xi_i \stackrel{iid}{\sim} D$
DISCO (non-quad)	$d^{.25} \left((1 + \frac{\sqrt{\kappa}}{n^{.25}}) \log \frac{1}{\epsilon} + \frac{\kappa^{1.5}}{n^{.75}} \right)$	$Qd^{.25} \left((1 + \frac{\sqrt{\kappa}}{n^{.25}}) \log \frac{1}{\epsilon} + \frac{\kappa^{1.5}}{n^{.75}} \right)$	(3), $\xi_i \stackrel{iid}{\sim} D$
DANE (quad)	$(1 + \frac{\kappa^2}{n}) \log \frac{1}{\epsilon}$	$Q(1 + \frac{\kappa^2}{n}) \log \frac{1}{\epsilon}$	(3), $\xi_i \stackrel{iid}{\sim} D$
CoCoA+	$\kappa \log \frac{1}{\epsilon}$	$G(n + \sqrt{\kappa n}) \kappa \log \frac{1}{\epsilon}$	(3)
Accel Grad	$\sqrt{\kappa f} \log \frac{1}{\epsilon}$	$Gn\sqrt{\kappa f} \log \frac{1}{\epsilon}$	

Table 1 Rounds and runtime of different distributed optimization algorithms in a general setting. Let G be the computation cost of the gradient of f_i . For typical problems such as logistic regression, we have $G = O(d)$. Let Q be the cost of solving a linear system. We have $Q = O(d^3)$ if exact matrix inversion is used and have $Q = O(dn\sqrt{\kappa} \log \frac{1}{\epsilon})$ if an ϵ -approximate inverse is found by an accelerated gradient method [17].

lower bound on the rounds of communication that shows any first-order distributed algorithm needs $\tilde{\Omega}(\sqrt{\frac{\kappa}{n}} \log(1/\epsilon))$ rounds of communication. It means DASVRG is optimal in that it uses the least number of rounds of communication. Since our lower bound indeed can be applied to a broad class of distributed first-order algorithms, it is interesting by itself. Here, and in the rest of the paper, \tilde{O} and $\tilde{\Omega}$ hide some logarithmic terms of κ , N , m and n .

The rest of this paper is organized as follows. In Section 2, we compared the theoretical performance of our methods with some existing work in distributed optimization. In Section 3 and Section 4, we propose our DSVRG and DASVRG algorithms, respectively, and discuss their theoretical guarantee. In Section 5, we prove a lower bound on the number of rounds of communication that a distributed algorithm needs, which demonstrates that DASVRG is optimal. Finally, we present the numerical experiments in Section 6, and conclude the paper in Section 7.

2 Related Work

Recently, there have been several distributed optimization algorithms proposed for problem (1). We list several of them, including a distributed implementation of the accelerated gradient method (Accel Grad) by Nesterov [17]⁵, in Table 1 and present their rounds and runtime for a clear comparison. The algorithms proposed in this paper are DSVRG and DASVRG.

The distributed dual coordinate ascent method, including DisDCA [27], CoCoA [10] and CoCoA+ [15], is a class of distributed coordinate optimization algorithms which can be applied to the conjugate dual formulation of (3). In these methods, each machine only updates n dual variables contained in a local problem defined on the n local data points. Any optimization algorithm can be used as a subroutine in each machine as long as it reduces the optimality gap of the local problem by a constant factor. According to [15,10], CoCoA+ requires $O(\kappa \log(1/\epsilon))$ rounds of communication to find an ϵ -optimal solution⁶. If the accelerated SDCA method [21,22] is used as the subroutine in each machine, the total runtime for CoCoA+ is $O((n + \sqrt{\kappa n}) \kappa \log(1/\epsilon))$. Therefore, both DSVRG and DASVRG have lower runtime and communication than CoCoA+, and the other distributed dual coordinate ascent variants.

Assuming the problem (1) has the form of (3) with ξ_i 's i.i.d. sampled from a distribution D (denoted by $\xi_i \stackrel{iid}{\sim} D$), the DANE [25] and DISCO [28] algorithms require $O((1 + \frac{\sqrt{\kappa}}{n^{.25}}) \log(1/\epsilon))$ and $O((1 + \frac{\kappa^2}{n}) \log(1/\epsilon))$ rounds of communication, respectively. Hence, DSVRG uses fewer rounds of communication than DANE and fewer than DISCO when $\kappa \leq n^{1.5}$. DASVRG always uses fewer rounds of communication than DISCO and DANE. Note that, for these four algorithms, the rounds can be very small in the ‘‘big data’’ case of large n . Indeed, as n increases, all four methods require only $O(\log \frac{1}{\epsilon})$ rounds which is independent of the condition number κ .

Moreover, DISCO and DANE have large running times due to solving a linear system each round, which is not practical for problems of large dimensionality. As an alternative, Zhang and Xiao [28] suggest solving the linear system with an inexact solution using another optimization algorithm, but

⁵ This is the accelerated gradient method by Nesterov [17] except that the data points are distributed in m machines to parallelize the computation of $\nabla f(x)$.

⁶ CoCoA+ has a better theoretical performance than CoCoA. According to [15], CoCoA+ is equivalent to DisDCA with ‘‘practical updates’’ [27] under certain choices of parameters.

Algorithm	Rounds	Runtime	Assumptions
DSVRG	$(1 + \sqrt{\frac{m}{n}}) \log \frac{1}{\epsilon}$	$G(n + \sqrt{mn}) \log \frac{1}{\epsilon}$	Assumption 1
DASVRG	$(1 + (\frac{m}{n})^{.25}) \log(1 + \sqrt{\frac{m}{n}}) \log \frac{1}{\epsilon}$	$G(n + n^{.75} m^{.25}) \log(1 + \sqrt{\frac{m}{n}}) \log \frac{1}{\epsilon}$	Assumption 1
DISCO (quad)	$m^{.25} \log \frac{1}{\epsilon}$	$Qm^{.25} \log \frac{1}{\epsilon}$	(3), $\xi_i \stackrel{iid}{\sim} D$

Table 2 Rounds and runtime of DSVRG, DASVRG and DISCO when $\kappa = \Theta(\sqrt{N})$. The coefficients Q and G are defined as in Table 1.

this still has large runtime for ill-conditioned problems. The runtimes of DISCO and DANE are shown in Table 1 with $Q = O(d^3)$ to represent the time for taking matrix inverse and $Q = O(dn\sqrt{\kappa} \log \frac{1}{\epsilon})$ to represent the time when an accelerated first-order method is used for solving the linear system. In both case, the runtimes of DISCO and DANE can be higher than those of DSVRG or DASVRG when d is large. Furthermore, DANE only has the theoretical guarantee mentioned above when it is applied to quadratic problems, for example, regularized linear regression. Also, DISCO only applies to self-concordant functions with easily computed Hessian⁷, and makes strong statistical assumptions on the data points. On the contrary, DSVRG and DASVRG works for a more general problem (1) and do assume $\xi_i \stackrel{iid}{\sim} D$ for (3).

We also make the connection to the recent lower bounds [2] for the rounds of communication needed by distributed optimization. Arjevani and Shamir [2] prove that, for a class of δ -related functions (see [2] for the definition) and, for a class of algorithms, the rounds of communication achieved by DISCO is optimal. However, as mentioned above, DASVRG needs fewer rounds than DISCO. This is not a contradiction since DASVRG does not fall into the class of algorithms subject to the lower bound in [2]. In particular, the algorithms concerned by [2] can only use the n local data points from the initial partition to update the local solutions while DASVRG samples and utilizes a second set of data points in each machine in addition to those n data points.

Building on the work of [2], we prove a new lower bound showing that any distributed first-order algorithm requires $\tilde{O}(\sqrt{\frac{\kappa}{n}} \log \frac{1}{\epsilon})$ rounds. This lower bound combined with the convergence analysis of DASVRG shows that DASVRG is optimal in the number of rounds of communication.

In Table 2, we compare the rounds and runtime of DSVRG, DASVRG and DISCO in the case where $\kappa = \Theta(\sqrt{N}) = \Theta(\sqrt{mn})$, which is a typical setting for ERM problem as justified in [28, 25, 24, 23]. We only compare our methods against DISCO, since it uses the fewest rounds of communication among other related algorithms. Let us consider the case where $n > m$, which is true in almost any reasonable distributed computing scenario. We can see from Table 2 that the rounds needed by DSVRG is lower than that of DISCO. In fact, both DSVRG and DASVRG use $O(\log \frac{1}{\epsilon})$ rounds of communication, which is almost a constant for many practical machine learning applications⁸.

3 Distributed SVRG

In this section, we consider a *distributed stochastic variance reduced gradient* (DSVRG) method that is based on a parallelization of SVRG [11, 26, 12]. SVRG works in multiple stages and, in each stage, one batch gradient is computed using all N data points and $O(\kappa)$ iterative updates are performed with only one data point processed in each. Our distributed algorithm randomly partitions the N data points onto m machines with n local data points on each to parallelize the computation of the batch gradient in SVRG. Then, we let the m machines conduct the iterative update of SVRG in serial in a “round-robin” scheme, namely, let all machine stay idle except one machine that performs a certain steps of iterative updates of SVRG using its local data and pass the solution to the next machine. However, the only caveat in this idea is that the iterative update of SVRG requires an unbiased estimator of $\nabla f(x)$ which can be constructed by sampling over the whole data set. However, the unbiasedness will be lost if each machine can only sample over its local data. To address this issue, we use the remaining $\tilde{n} = C - n$ memory space of each machine to store a second set of data which is uniformly sampled from the whole

⁷ The examples in [28] all take the form of $f_i(x) = g(a_i^T x)$ for some function g on \mathbb{R}^1 , which is more specific than $\phi(x, \xi_i)$. Under this form, it is relatively easy to compute the Hessian of f_i .

⁸ Typically $\epsilon \in (10^{-6}, 10^{-2})$, so $\log \frac{1}{\epsilon}$ is always less than 20.

data set before the algorithm starts. If each machine samples over this dataset, the unbiased estimator will be still available so that the convergence property can be inherited from the single-machine SVRG.

To load the second dataset mentioned above onto each machine, we design an efficient data allocation scheme which reuses the randomness of the first partitioned dataset to construct this second one. We show that this method helps to increase the overlap between the first and second dataset so that it requires smaller amount of communication than the direct implementation.

3.1 An efficient data allocation procedure

To facilitate the presentation, we define a *multi-set* as a collection of items where some items can be repeated. We allow taking the union of a regular set S and a multi-set R , which is defined as a regular set $S \cup R$ consisting of the item in either S or R without repetition.

We assume that a random partition S_1, \dots, S_m of $[N]$ can be constructed efficiently. A straightforward data allocation procedure is to prepare a partition S_1, \dots, S_m of $[N]$ and then sample a sequence of Q i.i.d. indices r_1, \dots, r_Q uniformly with replacement from $[N]$. After partitioning r_1, \dots, r_Q into m multi-sets $R_1, \dots, R_m \subset [N]$, we allocate data $\{f_i \mid i \in S_j \cup R_j\}$ to machine j . Since S_1, \dots, S_m has occupied n of the memory in each machine, Q can be at most $\tilde{n}m$. Note that the amount of communication in distributing S_1, \dots, S_m is exactly N which is necessary for almost all distributed algorithms. However, this straightforward procedure requires an extra $O(Q)$ amount of communication for distributing $R_j \setminus S_j$.

To improve the efficiency of data allocation, we propose a procedure which reuses the randomness of S_1, \dots, S_m to generate the indices r_1, \dots, r_Q so that the the overlap between S_j and R_j can be increased which helps reduce the additional amount of communication for distributing R_1, \dots, R_m . The key observation is that the concatenation of S_1, \dots, S_m is a random permutation of $[N]$ which has already provided enough randomness needed by R_1, \dots, R_m . Hence, it will be easy to build the i.i.d. samples r_1, \dots, r_Q by adding a little additional randomness on top of S_1, \dots, S_m . With this observation in mind, we propose our data allocation procedure in Algorithm 1.

Algorithm 1 Data Allocation : $DA(N, m, Q)$

Input: Index set $[N]$, the number of machines m , and the length of target sequence Q .

Output: A random partition S_1, \dots, S_m of $[N]$, indices $r_1, \dots, r_Q \in [N]$, multi-sets $R_1, \dots, R_m \subset [N]$, and data $\{f_i \mid i \in S_j \cup R_j\}$ stored on machine j for each $j \in [m]$.

Center samples r_1, \dots, r_Q and R_1, \dots, R_m as follows:

- 1: Randomly partition $[N]$ into m disjoint sets S_1, \dots, S_m of the same size $n = \frac{N}{m}$.
- 2: Concatenate the subsets S_1, \dots, S_m into a random permutation i_1, \dots, i_N of $[N]$ so that $S_j = \{i_{(j-1)n+1}, \dots, i_{jn}\}$.
- 3: **for** $\ell = 1$ to Q **do**
- 4: Let

$$r_\ell = \begin{cases} i_\ell & \text{with probability } 1 - \frac{\ell-1}{N} \\ i_{\ell'} & \text{with probability } \frac{1}{N} \text{ for } \ell' = 1, 2, \dots, \ell-1. \end{cases}$$

- 5: **end for**
- 6: Let

$$R_j = \begin{cases} \{r_{(j-1)\tilde{n}+1}, \dots, r_{j\tilde{n}}\} & \text{if } j = 1, \dots, \lceil Q/\tilde{n} \rceil - 1 \\ \{r_{\lceil Q/\tilde{n} \rceil - 1 \tilde{n} + 1}, \dots, r_Q\} & \text{if } j = \lceil Q/\tilde{n} \rceil \\ \emptyset & \text{if } j = \lceil Q/\tilde{n} \rceil + 1, \dots, m. \end{cases} \quad (4)$$

Distribute data points to machines:

- 7: Machine j acquires data points in $\{f_i \mid i \in S_j \cup R_j\}$ from the storage center.
-

The correctness and the expected amount of communication of Algorithm 1 are characterized as follows.

Lemma 1 *The sequence r_1, \dots, r_Q generated in Algorithm 1 has the same joint distribution as a sequence of i.i.d. indices uniformly sampled with replacement from $[N]$. Moreover, the expected amount of communication for distributing $\cup_{i=1}^m \{f_i \mid i \in R_j \setminus S_j\}$ is at most $\frac{Q^2}{N}$.*

Proof Conditioned on $i_1, \dots, i_{\ell-1}$ and $r_1, \dots, r_{\ell-1}$, the random index i_ℓ has uniform distribution over $[N] \setminus \{i_1, \dots, i_{\ell-1}\}$. Therefore, by Line 3 in Algorithm 1, the conditional distribution of the random index r_ℓ , conditioning on $i_1, \dots, i_{\ell-1}$ and $r_1, \dots, r_{\ell-1}$, is a uniform distribution over $[N]$. Hence, we complete the proof of the first claim of the lemma.

To analyze the amount of communication, we note that $i_\ell \neq r_\ell$ with probability $\frac{\ell-1}{N}$. Suppose $i_\ell \in S_j$ for some j . We know that the data point f_{r_ℓ} needs to be transmitted to machine j separately from S_j only if $i_\ell \neq r_\ell$. Therefore, the expected amount of communication for distributing $\cup_{i=1}^m \{f_i | i \in R_j \setminus S_j\}$ is upper bounded by $\sum_{\ell=1}^Q \frac{\ell-1}{N} \leq \frac{Q^2}{N}$. \square

According to Lemma 1, besides the (necessary) N amount of communication to distribute S_1, \dots, S_m , Algorithm 1 needs only $\frac{Q^2}{N}$ additional amount of communication to distribute R_1, \dots, R_m thanks to the overlaps between S_j and R_j for each j . This additional amount is less than the $O(Q)$ amount required by the straightforward method when $Q \leq N$.

Note that, in the DSVRG algorithm we will introduce later, we need $Q = O(\kappa \log(1/\epsilon))$. When $\kappa = \Theta(\sqrt{N})$ in a typical ERM problem, we have $\frac{Q^2}{N} = \frac{\kappa^2 (\log(1/\epsilon))^2}{N} = O((\log(1/\epsilon))^2)$ which is typically much less than the N amount of communication in distributing S_1, \dots, S_m . In other words, although DSVRG does require additional amount of communication to allocate the data than other algorithms, this additional amount is nearly negligible.

3.2 DSVRG algorithm and its theoretical guarantees

With the data $\{f_i | i \in S_j \cup R_j\}$ stored on machine j for each $j \in [m]$ after running Algorithm 1, we are ready to present the distributed SVRG algorithm in Algorithm 2 and Algorithm 3.

We start SVRG in machine k with $k = 1$ initially at an initial solution $\tilde{x}^0 \in \mathbb{R}^d$. At the beginning of stage ℓ of SVRG, all m machines participate in computing a batch gradient h^ℓ in parallel using the data indexed by S_1, \dots, S_m . Within stage ℓ , in each iteration, machine k samples one data f_i from its local data indexed by R_k to construct a stochastic gradient $\nabla f_i(x_t) - \nabla f_i(\tilde{x}^\ell) + h^\ell$ and performs the iterative update. Since R_k is a multi-set that consists of indices sampled with replacement from $[N]$, the unbiasedness of $\nabla f_i(x_t) - \nabla f_i(\tilde{x}^\ell) + h^\ell$, i.e., the property

$$\mathbb{E}_{i \sim R_k} [f_i(x_t) - \nabla f_i(\tilde{x}^\ell) + h^\ell] = \mathbb{E}_{i \sim N} [f_i(x_t) - \nabla f_i(\tilde{x}^\ell) + h^\ell] = \nabla f(x_t) \quad (5)$$

is guaranteed. After this iteration, i is removed from R_k .

The m machines do the iterative updates in the order from machine 1 to machine m . Once the current active machine, says machine k , has removed all of its samples in R_k (so that $R_k = \emptyset$), then it must pass the current solution and the running average of all solutions generated in the current stage to machine $k + 1$. At any time during the algorithm, there is only one machine updating the solution x_t and the other $m - 1$ machines only contribute in computing the batch gradient h^ℓ . We want to emphasize that it is important that machines should never use any samples in R_j 's more than once since, otherwise, the stochastic gradient $\nabla f_i(x_t) - \nabla f_i(\tilde{x}^\ell) + h^\ell$ will lose its unbiasedness. We describe formally each stage of this algorithm in Algorithm 2 and the iterative update in Algorithm 3.

Note that an implicit requirement of Algorithm 2 is $TK = Q \leq \tilde{n}m$. Because one element in R_k is removed in each *iterative update* (Line 3) of Algorithm 3, this update cannot be performed any longer once each R_k becomes empty. Hence, the condition $TK = Q$ ensures that the number of iterative updates matches the total number of indices contained in R_1, R_2, \dots, R_m . Note that, in the worse case, S_j and R_j may not overlap so that the size of data $\{f_i | i \in S_j \cup R_j\}$ stored on machine j can be $|S_j| + |R_j|$. That is why the size of R_j is only $\tilde{n} = C - n$. The condition $TK = Q \leq \tilde{n}m$ is needed here so the total amount of data stored in all machines do not exceed the total capacity Cm . The convergence of Algorithm 2 is established by the following theorem.

Theorem 1 *Suppose $0 < \eta < \frac{1}{4L}$ and $TK = Q \leq \tilde{n}m$. Algorithm 2 guarantees*

$$\mathbb{E} [f(\tilde{x}^K) - f(x^*)] \leq \left(\frac{1}{\mu\eta(1-4L\eta)T} + \frac{4L\eta(T+1)}{(1-4L\eta)T} \right)^K [f(\tilde{x}^0) - f(x^*)]. \quad (6)$$

In particular, when $\eta = \frac{1}{16L}$, $T = 96\kappa$ and $TK = Q \leq \tilde{n}m$, Algorithm 2 needs $K = \frac{1}{\log(9/8)} \log \left(\frac{f(\tilde{x}^0) - f(x^)}{\epsilon} \right)$ stages to find an ϵ -optimal solution.*

Algorithm 2 Distributed SVRG (DSVRG)

Input: An initial solution $\tilde{x}^0 \in \mathbb{R}^d$, data $\{f_i\}_{i=1,\dots,N}$, the number of machine m , a step length $\eta < \frac{1}{4L}$, the number of iterations T in each stage, the number of stages K , and a sample size $Q = TK$.

Output: \tilde{x}^K

- 1: Use DA to generate (a) a random partition S_1, \dots, S_m of $[N]$, (b) Q i.i.d. indices r_1, \dots, r_Q uniformly sampled with replacement from $[N]$, (c) m multi-sets $\mathcal{R} = \{R_1, \dots, R_m\}$ defined as (4), and (d) data $\{f_i \mid i \in S_j \cup R_j\}$ stored on machine j .
 - 2: $k \leftarrow 1$
 - 3: **for** $\ell = 0, 1, 2, \dots, K - 1$ **do**
 - 4: Center sends \tilde{x}^ℓ to each machine
 - 5: **for** machine $j = 1, 2, \dots, m$ **in parallel do**
 - 6: Compute $h_j^\ell = \sum_{i \in S_j} \nabla f_i(\tilde{x}^\ell)$ and send it to center
 - 7: **end for**
 - 8: Center computes $h^\ell = \frac{1}{N} \sum_{j=1}^m h_j^\ell$ and send it to machine k
 - 9: $(\tilde{x}^{\ell+1}, \mathcal{R}, k) \leftarrow \text{SS-SVRG}(\tilde{x}^\ell, h, \mathcal{R}, k, \eta, T)$
 - 10: **end for**
-

Algorithm 3 Single-Stage SVRG: SS-SVRG($\tilde{x}, h, \mathcal{R}, k, \eta, T, \{f_i\}_{i=1,\dots,N}$)

Input: A solution $\tilde{x} \in \mathbb{R}^d$, data $\{f_i\}_{i=1,\dots,N}$, a batch gradient h , m multi-sets $\mathcal{R} = \{R_1, R_2, \dots, R_m\}$, the index of the active machine k , a step length $\eta < \frac{1}{4L}$, and the number of iterations T .

Output: The average solution \bar{x}_T , the updated multi-sets $\mathcal{R} = \{R_1, R_2, \dots, R_m\}$, and the updated index of active machine k .

- 1: $x_0 = \tilde{x}$ and $\bar{x}_0 = \mathbf{0}$
- 2: **for** $t = 0, 1, 2, \dots, T - 1$ **do**
- 3: Machine k samples an instance i from R_k and computes

$$x_{t+1} = x_t - \eta(\nabla f_i(x_t) - \nabla f_i(\tilde{x}) + h), \quad \bar{x}_{t+1} = \frac{x_{t+1} + t\bar{x}_t}{t+1}, \quad R_k \leftarrow R_k \setminus \{i\}$$

- 4: **if** $R_k = \emptyset$ **then**
 - 5: x_{t+1} and \bar{x}_{t+1} are sent to machine $k+1$
 - 6: $k \leftarrow k+1$
 - 7: **end if**
 - 8: **end for**
-

Proof In the iterative update given in Line 3 of Algorithm 3, a stochastic gradient $\nabla f_i(x_t) - \nabla f_i(\tilde{x}) + h$ is constructed with h being the batch gradient $\nabla f(\tilde{x})$ and i sampled from R_k in the active machine k . Since i is one of the indices r_1, \dots, r_Q , each of which is sampled uniformly from $[N]$, this stochastic gradient is unbiased estimator of $\nabla f(x_t)$. Therefore, the path of solutions $\tilde{x}^0, \tilde{x}^1, \tilde{x}^2, \dots$ generated by Algorithm 2 has the same distribution as the ones generated by single-machine SVRG so that the convergence result for the single-machine SVRG can be directly applied to Algorithm 2. The inequality (6) has been shown in Theorem 1 in [26] for single-machine SVRG, which now also holds for Algorithm 2.

When $\eta = \frac{1}{16L}$ and $T = 96\kappa$, it is easy to show that

$$\frac{1}{\mu\eta(1-4L\eta)T} + \frac{4L\eta(T+1)}{(1-4L\eta)T} \leq \frac{1}{\mu\eta(1-4L\eta)T} + \frac{8L\eta}{(1-4L\eta)} = \frac{2}{9} + \frac{2}{3} = \frac{8}{9} \quad (7)$$

so that Algorithm 2 needs $K = \frac{1}{\log(9/8)} \log\left(\frac{f(\tilde{x}^0) - f(x^*)}{\epsilon}\right)$ stages to find an ϵ -optimal solution. \square

By Theorem 1, DSVRG can find an ϵ -optimal solution for (1) after $K = O(\log(1/\epsilon))$ stages with $T = O(\kappa)$ iterative updates (Line 3 of Algorithm 3) in each stage. Therefore, there are $O(\kappa \log(1/\epsilon))$ iterative updates in total so that Q must be $TK = O(\kappa \log(1/\epsilon))$. Since the available memory space requires $Q \leq \tilde{n}m = (C - n)m$. We will need at least $C = \Omega(n + \frac{\kappa}{m} \log(1/\epsilon))$ in order to implement DSVRG.

Under the assumptions of Theorem 1, including Assumption 1 (so $\tilde{n} = \Omega(n)$), we discuss the theoretical performance of DSVRG as follows.

- **Local parallel runtime:** Since one gradient is computed in each iterative update and n gradients are computed in parallel to construct the batch gradient, the total local parallel runtime for DSVRG to find an ϵ -optimal solution is $O((n+T)K) = O((n+\kappa) \log(1/\epsilon))$
- **Communication:** There is a fixed amount of $O(N)$ communication needed to distribute the partitioned data S_1, \dots, S_m to m machines. When Algorithm 1 is used to generate R_1, \dots, R_m , the

additional amount of communication to complete the data allocation step of DSVRG is $O(Q^2/N) = O(\kappa^2 \log^2(1/\epsilon)/N)$ in expectation according to Lemma 1. During the algorithm, the batch gradient computations and the iterative updates together require $O((m + \frac{T}{\tilde{n}})K) = O((m + \frac{\kappa}{\tilde{n}}) \log(1/\epsilon))$ amount of communication.

- **Rounds of communication:** Since DSVRG needs one round of communication to compute a batch gradient at each stage (one call of SS-SVRG) and one round after every \tilde{n} iterative update (when $R_k = \emptyset$), it needs $O(K + \frac{TK}{\tilde{n}}) = O((1 + \frac{\kappa}{\tilde{n}}) \log(1/\epsilon))$ rounds of communication in total to find an ϵ -optimal solution. We note that, if the memory space C in each machine is large enough, the value of \tilde{n} can larger than κ so that the rounds of communication needed will be only $O(\log(1/\epsilon))$.

3.3 Regimes where DSVRG is Optimal

In this subsection, we consider a scenario where $\kappa = \Theta(n^{1-2\delta})$ with a constant $0 < \delta < \frac{1}{2}$ and $\epsilon = \frac{1}{n^s}$ with positive constant s . We show that with a different choices of η , T and K , DSVRG can find an ϵ -optimal solution for (1) with the optimal parallel runtime, the optimal amount of communication, and the optimal number of rounds of communication simultaneously.

Proposition 1 *Suppose $\kappa \leq \frac{n^{1-2\delta}}{32}$ with a constant $0 < \delta < \frac{1}{2}$ and we choose $\eta = \frac{1}{16n^\delta L}$, $T = n$ and $K = \frac{1}{\log(n^\delta/2)} \log\left(\frac{f(\tilde{x}^0) - f(x^*)}{\epsilon}\right)$ in Algorithm 2. Also, suppose Assumption 1 holds and $TK = Q \leq \tilde{n}m$. Algorithm 2 finds ϵ -optimal solution for (1) with $O(\frac{\log(1/\epsilon)}{\delta \log n})$ rounds of communications, $O(\frac{n \log(1/\epsilon)}{\delta \log n})$ total parallel runtime, and $O(\frac{m \log(1/\epsilon)}{\delta \log n})$ amount of communication.*

In particular, when $\epsilon = \frac{1}{n^s}$ with a positive constant s , Algorithm 2 finds ϵ -optimal solution for (1) with $O(1)$ rounds of communications, $O(n)$ total parallel runtime, and $O(m)$ amount of communication.

Proof Since $\eta = \frac{1}{16n^\delta L} < \frac{1}{4L}$, Algorithm 2 guarantees (6) according to Theorem 1. With $T = n$ and $\eta = \frac{1}{16n^\delta L}$, we have

$$\begin{aligned} \frac{1}{\mu\eta(1-4L\eta)T} + \frac{4L\eta(T+1)}{(1-4L\eta)T} &\leq \frac{1}{\mu\eta(1-4L\eta)T} + \frac{8L\eta}{(1-4L\eta)} = \frac{16n^\delta \kappa}{(1-1/(4n^\delta))n} + \frac{1}{2n^\delta(1-1/(4n^\delta))} \\ &\leq \frac{1}{2n^\delta(1-1/(4n^\delta))} + \frac{1}{2n^\delta(1-1/(4n^\delta))} \leq \frac{2}{n^\delta} \end{aligned}$$

Hence, $\mathbb{E}[f(\tilde{x}^K) - f(x^*)] \leq \epsilon$ can be implied by the inequality (6) as $K = \frac{1}{\log(n^\delta/2)} \log\left(\frac{f(\tilde{x}^0) - f(x^*)}{\epsilon}\right)$.

Under Assumption 1, we have $\tilde{n} = \Omega(n)$ so that the number of rounds of communication Algorithm 2 needs is $O(K + \frac{TK}{\tilde{n}}) \leq O(K + \frac{TK}{n}) = O(\frac{\log(1/\epsilon)}{\delta \log n})$. Moreover, the total parallel needed is $O((n+T)K) = O(\frac{n \log(1/\epsilon)}{\delta \log n})$ and the amount of communication needed is $O((m + \frac{T}{\tilde{n}})K) = O(\frac{m \log(1/\epsilon)}{\delta \log n})$. The second conclusion can be easily derived by replacing ϵ with $O(\frac{1}{n^s})$. \square

The justification for the scenario where $\kappa = \Theta(n^{1-2\delta})$ and $\epsilon = \frac{1}{n^s}$ and why these performance guarantees are optimal have been discussed in Section 1.2.

4 Accelerated Distributed SVRG

In this section, we use the generic acceleration techniques in [8] and [14] to further improve the theoretical performance of DSVRG and obtain a *distributed accelerated stochastic variance reduced gradient* (DASVRG) method.

4.1 DASVRG algorithm and its theoretical guarantees

Following [8] and [14], we define a *proximal function* for $f(x)$ as

$$f_\sigma(x; y) \equiv f(x) + \frac{\sigma}{2} \|x - y\|^2 = \frac{1}{N} \sum_{i=1}^N \tilde{f}_i(x; y), \quad (8)$$

where $\tilde{f}_i(x; y) = f_i(x) + \frac{\sigma}{2} \|x - y\|^2$, $\sigma \geq 0$ is a constant to be determined later and $y \in \mathbb{R}^d$ is a *proximal point*. The condition number of this proximal function is $\kappa(f_\sigma) \equiv \frac{L+\sigma}{\mu+\sigma}$ which can be smaller than κ when σ is large enough.

Given an algorithm, denoted by \mathcal{A} , that can be applied to (1), the acceleration scheme developed in [8] and [14] is an iterative method that involves inner and outer loops and uses \mathcal{A} as a sub-routine in its outer loops. In particular, in p -th outer iteration of this acceleration scheme, the algorithm \mathcal{A} is applied to find a solution for the p -th proximal problem defined on a proximal point y_{p-1} , namely,

$$f_p^* \equiv \min_{x \in \mathbb{R}^d} f_\sigma(x; y_{p-1}) \text{ for } p = 1, 2, \dots, P. \quad (9)$$

The algorithm \mathcal{A} does not need to solve (9) to optimality but only needs to generate an approximate solution \hat{x}_p with an accuracy ϵ_p in the sense that

$$f_\sigma(\hat{x}_p; y_{p-1}) - f_p^* \leq \epsilon_p. \quad (10)$$

When $\kappa(f_\sigma)$ is smaller than κ , finding such an \hat{x}_p is easier than finding an ϵ -optimal solution for (1).

Then, the acceleration scheme uses \hat{x}_p to construct a new proximal point y_p using an extrapolation update as $y_p = \hat{x}_p + \beta_p(\hat{x}_p - \hat{x}_{p-1})$, where $\beta_p \geq 0$ is an extrapolation step length. After that, the $p+1$ -th proximal problem is constructed based on y_p which will be solved in the next outer iteration. With an appropriately chosen value for σ , it is shown by [8] and [14] that, for many existing \mathcal{A} including SAG [20, 19], SAGA [6], SDCA [21], SVRG [11] and Finito/MISO [7, 16], this acceleration scheme needs a smaller runtime for finding an ϵ -optimal solution than applying algorithm \mathcal{A} directly to (1).

Given the success of this acceleration scheme in the single-machine setting, it will be promising to also apply this scheme to the DSVRG to further improve its theoretical performance. Indeed, this can be done by choosing \mathcal{A} in this aforementioned acceleration scheme to be DSVRG. Then, we can obtain the DASVRG algorithm. In particular, in the p -th outer iteration of the aforementioned acceleration scheme, we use DSVRG to solve the proximal problem (9) in a distributed way up to an accuracy ϵ_p . We present DASVRG in Algorithm 4 where $\tilde{f}_i(x; y) = f_i(x) + \frac{\sigma}{2} \|x - y\|^2$.

Algorithm 4 Distributed Accelerated SVRG (DASVRG)

Input: An initial solution $\hat{x}_0 \in \mathbb{R}^d$, data $\{f_i\}_{i=1, \dots, N}$, the number of machine m , a step length $\eta < \frac{1}{4L}$, the number of iterations T in each stage of DSVRG, the number of stages K of DSVRG, the number of outer iterations P in the acceleration scheme, a sample size $TKP = Q \leq \tilde{n}m$, and a parameter $\sigma \geq 0$.

Output: \hat{x}_P

- 1: Use DA to generate (a) a random partition S_1, \dots, S_m of $[N]$, (b) Q of i.i.d. indices r_1, \dots, r_Q uniformly sampled with replacement from $[N]$, (c) m multi-sets $\mathcal{R} = \{R_1, \dots, R_m\}$ defined as (4), and (d) data $\{f_i \mid i \in S_j \cup R_j\}$ stored on machine j .
 - 2: $k \leftarrow 1$
 - 3: Initialize $q = \frac{\mu}{\mu+\sigma}$, $y_0 = \hat{x}_0$ and $\alpha_0 = \sqrt{q}$
 - 4: **for** $p = 1, 2, \dots, P$ **do**
 - 5: Center computes $\tilde{x}^0 = \hat{x}_{p-1}$ and sends y_{p-1} to each machine
 - 6: **for** $\ell = 0, 1, 2, \dots, K-1$ **do**
 - 7: Center sends \tilde{x}^ℓ to each machine
 - 8: **for** machine $j = 1, 2, \dots, m$ **in parallel do**
 - 9: Compute $h_j^\ell = \sum_{i \in S_j} \nabla \tilde{f}_i(\tilde{x}^\ell; y_{p-1})$ and send it to center
 - 10: **end for**
 - 11: Center computes $h^\ell = \frac{1}{N} \sum_{j=1}^m h_j^\ell$ and sends it to machine k
 - 12: $(\tilde{x}^{\ell+1}, \mathcal{R}, k) \leftarrow \text{SS-SVRG}(\tilde{x}^\ell, h^\ell, \mathcal{R}, k, \eta, T, \{\tilde{f}_i(x; y_{p-1})\}_{i=1, \dots, N})$
 - 13: **end for**
 - 14: Machine k computes $\hat{x}_p = \tilde{x}^K$ and sends \hat{x}_p to center
 - 15: Center computes $\alpha_p \in (0, 1)$ from the equation $\alpha_p^2 = (1 - \alpha_p)\alpha_{p-1}^2 + q\alpha_p$.
 - 16: Center computes $y_p = \hat{x}_p + \beta_p(\hat{x}_p - \hat{x}_{p-1})$, where $\beta_p = \frac{\alpha_{p-1}(1 - \alpha_{p-1})}{\alpha_{p-1}^2 + \alpha_p}$.
 - 17: **end for**
-

Proposition 2 Suppose $\eta = \frac{1}{16L}$, $T = 96\kappa(f_\sigma)$,

$$K = \frac{1}{\log(9/8)} \log \left(\frac{4}{2 - \sqrt{q}} + \frac{10368\sigma}{\mu q (1 - \frac{\sqrt{q}}{2})^2} \right) = O \left(\log \left(1 + \frac{\sigma}{\mu} \right) \right). \quad (11)$$

and $TKP = Q \leq \tilde{n}m$. The solution \hat{x}_p generated in Algorithm 4 satisfies (10) with

$$\epsilon_p = \frac{2}{9} [f(\hat{x}_0) - f(x^*)] \left(1 - \frac{\sqrt{q}}{2}\right)^p, \text{ for } p = 1, 2, \dots, P, \quad (12)$$

Moreover, Algorithm 4 finds an ϵ -optimal solution for (1) after $P = O\left(\sqrt{\frac{\mu+\sigma}{\mu}} \log(1/\epsilon)\right)$ outer iterations.

Proof Due to the unbiasedness of $\nabla f_i(x_t) - \nabla f_i(\hat{x}^\ell) + h^\ell$ (5), conditioning on \hat{x}_{p-1} , the solution path $\tilde{x}^0, \tilde{x}^1, \tilde{x}^2, \dots$ generated within the p -th outer loop of Algorithm 4 has same distribution as the solution path generated by applying single-machine SVRG to (9) with an initial solution of \hat{x}_{p-1} . Hence, all the convergence results of SVRG can be applied.

According to (7) in the proof of Theorem 1, the choices of η and T ensure $\mathbb{E}[f_\sigma(\hat{x}_p; y_{p-1}) - f_p^*] \leq \left(\frac{8}{9}\right)^K [f_\sigma(\hat{x}_{p-1}; y_{p-1}) - f_p^*]$. Using this result and following the analysis in Section B.2 in [14], we can show that \hat{x}_p satisfies (10) with ϵ_p given by (12) if K is set to (11).

Therefore, according to Theorem 3.1 in [14] with $\rho = \frac{\sqrt{q}}{2}$, Algorithm 4 guarantees that $\mathbb{E}[f(\hat{x}_P) - f(x^*)] \leq \frac{32}{q} \left(1 - \frac{\sqrt{q}}{2}\right)^{P+1} [f(\hat{x}_0) - f(x^*)]$ by choosing K as 11. This means Algorithm 4 finds an ϵ -optimal solution for (1) after $P = O\left(\frac{1}{\sqrt{q}} \log(1/\epsilon)\right) = O\left(\sqrt{\frac{\mu+\sigma}{\mu}} \log(1/\epsilon)\right)$ outer loops.

The condition $TKP = Q \leq \tilde{n}m$ here is only to guarantee that we have enough samples in R_1, R_2, \dots, R_m to finish a total of TKP iterative updates. \square

We will choose $\sigma = \frac{L}{n}$ in DASVRG and obtain the following theorem.

Theorem 2 Suppose $\eta = \frac{1}{16L}$, $T = 96\kappa(f_\sigma)$, K is chosen as (11), $TKP = Q \leq \tilde{n}m$ and $\sigma = \frac{L}{n}$. Algorithm 4 finds an ϵ -optimal solution for (1) after $O\left((1 + \sqrt{\frac{\kappa}{n}}) \log\left(1 + \frac{\kappa}{n}\right) \log(1/\epsilon)\right)$ calls of SS-SVRG.

Proof When $\sigma = \frac{L}{n}$, according to Proposition 2, DASVRG needs $P = O\left(\sqrt{\frac{\mu+\sigma}{\mu}} \log(1/\epsilon)\right) = O\left((1 + \sqrt{\frac{\kappa}{n}}) \log(1/\epsilon)\right)$ outer iterations with $K = O\left(\log\left(1 + \frac{\sigma}{\mu}\right)\right) = O\left(\log\left(1 + \frac{\kappa}{n}\right)\right)$ calls of SS-SVRG in each according to (11). Hence, a total of $KP = O\left((1 + \sqrt{\frac{\kappa}{n}}) \log\left(1 + \frac{\kappa}{n}\right) \log(1/\epsilon)\right)$ calls of SS-SVRG are needed. \square

When $\sigma = \frac{L}{n}$, we have $\kappa(f_\sigma) = \frac{L+\sigma}{\mu+\sigma} = \frac{nL+L}{n\mu+L} \leq n+1$. According to Theorem 2, DASVRG can find an ϵ -optimal solution for (1) after $\tilde{O}\left((1 + \sqrt{\frac{\kappa}{n}}) \log(1/\epsilon)\right)$ calls of SS-SVRG. Note that each call of SS-SVRG involves $T = O(\kappa(f_\sigma)) = O(n)$ iterative updates. Therefore, there are $Q = TKP = \tilde{O}\left((n + \sqrt{n\kappa}) \log(1/\epsilon)\right)$ iterative updates in total. Since the available memory space requires $Q \leq \tilde{n}m = (C-n)m$, we can derive from the inequality $\tilde{O}\left((n + \sqrt{n\kappa}) \log(1/\epsilon)\right) \leq (C-n)m$ that we need at least $C = \tilde{O}\left(n + \left(\frac{n}{m} + \frac{\sqrt{n\kappa}}{m}\right) \log(1/\epsilon)\right)$ in order to implement DASVRG.

Under the assumptions of Theorem 2, including Assumption 1 (so $\tilde{n} = \Omega(n)$), we summarize the theoretical performance of DASVRG as follows.

- **Local parallel runtime:** Since each call of SS-SVRG involves a batch gradient computation and $T = O(n)$ iterative update, the total runtime of DASVRG is $O((n+T)KP) = \tilde{O}(n + \sqrt{n\kappa}) \log(1/\epsilon)$.
- **The amount of communication:** Similar to DSVRG, we need a fixed amount $O(N)$ communication to distribute S_1, \dots, S_m and an additional amount of $O(\kappa^2 \log^2(1/\epsilon)/N)$ communication to distribute R_1, \dots, R_m to m machines. During the algorithm, the batch gradient computations and the iterative updates together require $O\left((m + \frac{T}{n})KP\right) = \tilde{O}\left((m + \frac{T}{n})(1 + \sqrt{\frac{\kappa}{n}}) \log(1/\epsilon)\right) = \tilde{O}\left((m + m\sqrt{\frac{\kappa}{n}}) \log(1/\epsilon)\right)$ amounts of communication.
- **The number of rounds of communication:** Since each call of SS-SVRG needs $1 + \frac{T}{n}$ rounds communication, the rounds of DASVRG is $O\left((1 + \frac{T}{n})KP\right) = \tilde{O}\left((1 + \sqrt{\frac{\kappa}{n}}) \log(1/\epsilon)\right)$.

Recall that DSVRG needs $O\left((1 + \frac{\kappa}{n}) \log(1/\epsilon)\right)$ rounds of communication which is more than DASVRG. The crucial observation here is that, although in the single-machine setting, the acceleration scheme of [8, 14] only helps when $\kappa \geq N$, in the distributed setting, it helps to reduce the rounds as long as κ is larger than the number of local samples n .

5 Lower Bounds on Rounds of Communication

In this section, we prove that, under Assumption 1, any distributed first-order method will require at least $\tilde{\Omega}(\sqrt{\frac{\kappa}{n}} \log(1/\epsilon))$ rounds to find an ϵ -optimal solution for (1) with both partitioned data and i.i.d. sampled data in each machine. This lower bound is matched by the upper bound of the rounds needed by DASVRG in Section 4 up to some logarithmic terms. We note that we are working under different scenarios than [2]: In [2], the authors assumed that the only property of the data that an algorithm can exploit is that the local sums are δ -related for a $\delta \approx 1/\sqrt{n}$, and proved that the number of rounds is at least $\Omega(\sqrt{\kappa}/n^{1/4})$. The DASVRG algorithm exploits the fact that data is randomly partitioned and outperforms their lower bound. This suggests that δ -relatedness shouldn't be the only property that an algorithm exploits, and motivates us to prove a new (matching) lower bound by assuming the data is randomly partitioned.

5.1 A lower bound for rounds of communication

We first consider a family of algorithms which consist of a data distribution stage where the functions $\{f_i\}_{i \in [N]}$ are distributed onto m machines, and a distributed computation stage where, in each round, machines can not only use first-order (gradient) information of the functions stored locally but also apply preconditioning using local second-order information (Hessian matrix).

Definition 1 (Distributed (extended) first-order algorithms \mathcal{F}_α) We say an algorithm \mathcal{A} for solving (1) with m machines belongs to the family \mathcal{F}_α ($\mathcal{A} \in \mathcal{F}_\alpha$) of distributed first-order algorithms if it distributes $\{f_i\}_{i \in [N]}$ to m machines only once at the beginning such that:

1. The index set $[N]$ is randomly and evenly partitioned into S_1, S_2, \dots, S_m with $|S_j| = n$ for $j = 1, \dots, m$.
2. A multi-set R_j of size αn is created by sampling with replacement from $[N]$ for $j = 1, \dots, m$, where $\alpha \geq 0$ is a constant.
3. Let $S'_j = S_j \cup R_j$. Machine j acquires functions in $\{f_i | i \in S'_j\}$ for $j = 1, \dots, m$.

and let the machines do the following operations in rounds:

1. Machine j maintains a local set of vectors W_j initialized to be $W_j = \{\mathbf{0}\}$.
2. In each round, for arbitrarily many times, machine j can add any w to W_j if w satisfies (c.f. [2])

$$\begin{aligned} \gamma w + \nu \nabla F_j(w) \in \text{span} \{w', \nabla F_j(w'), (\nabla^2 F_j(w') + D)w'', (\nabla^2 F_j(w') + D)^{-1}w'' \mid \\ w', w'' \in W_j, D \text{ diagonal, } \nabla^2 F_j(w') \text{ and } (\nabla^2 F_j(w') + D)^{-1} \text{ exists} \} \end{aligned} \quad (13)$$

for some γ, μ such that $\gamma \nu \neq 0$, where $F_j = \sum_{i \in U_j \subset S'_j} f_i$ with an arbitrary subset U_j of S'_j .

3. At the end of the round, all machines can simultaneously send any vectors in W_j to any other machines, and machines can add the vectors received from other machines to its local working set.
4. The final output is a vector in the linear span of one W_j .

We define $\mathcal{A}(\{f_i\}_{i \in [N]}, H)$ as the output vector of \mathcal{A} when it is applied to (1) for H rounds with the inputs $\{f_i\}_{i \in [N]}$.

Besides the randomness due to the data distribution stage, the algorithm \mathcal{A} itself can be a randomized algorithm. Hence, the output $\mathcal{A}(\{f_i\}_{i \in [N]}, H)$ can be a random variable.

We would like to point out that, although the algorithms in \mathcal{F}_α can use the local second-order information like $\nabla^2 f_i(x)$ in each machine, Newton's method is still not contained in \mathcal{F}_α since Newton's method requires the access to the global second-order information such as $\nabla^2 f(x)$ (machines are not allowed to share matrices with each other). That being said, one can still use a distributed iteration method which can multiply $\nabla^2 f_i(x)$ to a local vector in order to solve the inversion of $\nabla^2 f(x)$ approximately. This method will lead to a distributed inexact Newton method such as DISCO [28]. In fact, both DANE [25] and DISCO [28] belong to \mathcal{F}_α with $\alpha = 0$. Suppose, in Assumption 1, the capacity of each machine C is given such that $\tilde{n} = cn$. The DSVRG and DASVRG algorithms proposed in this paper belong to \mathcal{F}_c with $\alpha = c$.

We are ready to present the lower bounds for the rounds of communications.

Theorem 3 Suppose $\kappa \geq n$ and there exists an algorithm $\mathcal{A} \in \mathcal{F}_\alpha$ with the following property:

“For any $\epsilon > 0$ and any N convex functions $\{f_i\}_{i \in [N]}$ where each f_i is L -smooth and f defined in (1) is μ -strongly convex, there exists H_ϵ such that the output $\hat{x} = \mathcal{A}(\{f_i\}_{i \in [N]}, H_\epsilon)$ satisfies $\mathbb{E}[f(\hat{x}) - f(x^*)] \leq \epsilon$.”

Then, when $m \geq \max\{\exp(\frac{\alpha}{\max\{1, \alpha\}} e^{\frac{2}{\max\{1, \alpha\}} + 1}), (e + \max\{1, \alpha\})^2\}$, we must have

$$H_\epsilon \geq \left(\frac{\sqrt{\kappa/n - 1}}{4\sqrt{2}((e + \max\{1, \alpha\}) \log m)^{3/2}} \right) \log \left(\left(1 - \frac{1}{(e + e^\alpha)(e + \max\{1, \alpha\})^2} \right) \frac{\mu n \|w^*\|^2}{4\epsilon} \right) \geq \Omega \left(\sqrt{\frac{\kappa}{n(\log m)^3}} \log \left(\frac{\mu n \|w^*\|^2}{\epsilon} \right) \right).$$

We want to emphasize that Theorem 3 holds without assuming Assumption 1.⁹ In the definition of \mathcal{F}_α , we allow the algorithm to access both randomly partitioned data and independently sampled data, and allow the algorithm to use local Hessian for preconditioning. This makes our lower bounds in Theorem 3 **stronger**: Even with an algorithm more powerful than first-order methods (in terms of the class of operations it can take) and with more options in distributing data, the number of rounds needed to find an ϵ -optimal solution still cannot be reduced.

We note that the condition $\kappa \geq n$ in Theorem 3 is necessary. Recall that, when $\kappa \leq \frac{n^{1-2\delta}}{32} < n$ for a constant $0 < \delta < \frac{1}{2}$, we showed in Proposition 1 in Subsection 3.3 that $O(\frac{\log(1/\epsilon)}{\delta \log n})$ rounds is enough for DSVRG. Therefore, the lower bound $H \geq \tilde{\Omega}(\sqrt{\frac{\kappa}{n}} \log(\frac{1}{\epsilon})) = \tilde{\Omega}(\frac{1}{n^\delta} \log(\frac{1}{\epsilon}))$ won't be true for the case when $\kappa \leq \frac{n^{1-2\delta}}{32} < n$.

5.2 Proof for the lower bound

Given a vector $x \in \mathbb{R}^d$ and a set of indices $D \subset [d]$, we use x_D to represent the sub-vector of x that consists of the coordinates of x indexed by D .

Definition 2 A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is *decomposable* with respect to a partition D_1, \dots, D_r of coordinates $[d]$ if f can be written as $f(x) = g^l(x_{D_1}) + \dots + g^l(x_{D_r})$, where $g^l : \mathbb{R}^{|D_l|} \rightarrow \mathbb{R}$ for $l = 1, \dots, r$. A set of functions $\{f_i\}_{i \in [N]}$ is *simultaneously decomposable* w.r.t a partition D_1, \dots, D_r if each f_i is decomposable w.r.t D_1, \dots, D_r .

It follows the Definition 1 and Definition 2 straightforwardly that:

Proposition 3 Suppose the functions $\{f_i\}_{i \in [N]}$ in (1) are simultaneously decomposable with respect to a partition D_1, \dots, D_r so that $f_i(x) = \sum_{l=1}^r g_i^l(x_{D_l})$ with functions $g_i^l : \mathbb{R}^{|D_l|} \rightarrow \mathbb{R}$ for $i = 1, \dots, N$ and $l = 1, \dots, r$. We have

$$x_{D_l}^* = \arg \min_{w \in \mathbb{R}^{|D_l|}} \left\{ \bar{g}^l(w) \equiv \frac{1}{N} \sum_{i=1}^N g_i^l(w) \right\}, \text{ for } l = 1, 2, \dots, r. \quad (14)$$

where x^* is the optimal solution of (1).

Moreover, any algorithm $\mathcal{A} \in \mathcal{F}_\alpha$, when applied to $\{f_i\}_{i \in [N]}$, becomes decomposable with respect to the same partition D_1, \dots, D_r in the following sense: For $l = 1, \dots, r$, there exists an algorithm¹⁰ $\mathcal{A}_l \in \mathcal{F}_\alpha$ such that, after any number of rounds H ,

$$\mathbb{E}[f(\hat{x}) - f(x^*)] = \sum_{l=1}^r \mathbb{E}[\bar{g}^l(\hat{w}^l) - \bar{g}^l(x_{D_l}^*)],$$

where $\hat{x} = \mathcal{A}(\{f_i\}_{i \in [N]}, H) \in \mathbb{R}^d$ and $\hat{w}^l = \mathcal{A}_l(\{g_i^l\}_{i \in [N]}, H) \in \mathbb{R}^{|D_l|}$ for $l = 1, 2, \dots, r$.

Proof The proof of this proposition is straightforward. Since $\{f_i\}_{i \in [N]}$ in (1) are simultaneously decomposable with respect to a partition D_1, \dots, D_r , we have

$$f(x) = \frac{1}{N} \sum_{i=1}^N \sum_{l=1}^r g_i^l(x_{D_l}) = \sum_{l=1}^r \bar{g}^l(x_{D_l})$$

⁹ If the capacity C is such that $\tilde{n} = 0$ and $c = 0$, the lower bound given by Theorem 3 still applies to the algorithms in \mathcal{F}_α with $\alpha = 0$.

¹⁰ Here and throughout the paper, an algorithm is essentially considered as a sequence of operations allowed by Definition 1.

so that the problem (1) can be solved by solving (14) for each l separately and $x_{D_l}^*$ must be the solution of the l -th problem in (14).

In addition, the function F_j in Definition 1 is also decomposable with respect to the same partition D_1, \dots, D_r . As a result, its gradient $\nabla F_j(x)$ also has a decomposed structure in the sense that $[\nabla F_j(x)]_{D_l}$ only depends on x_{D_l} for $l = 1, 2, \dots, r$. Similarly, its Hessian matrix $\nabla^2 F_j(x)$ is a block diagonal matrix with r blocks and the l -th block only depends on x_{D_l} . These properties ensure that each operation as (13) conducted by \mathcal{A} can be decomposed into r independent operations as (13) and applied on $x_{D_1}, x_{D_2}, \dots, x_{D_r}$ separately. The data distribution and the sequence of operations conducted by \mathcal{A} on x_{D_l} can be viewed as an algorithm $\mathcal{A}_l \in \mathcal{F}_\alpha$ applied to $\{g_i^l\}_{i \in [N]}$ so that $\hat{w}^l = \mathcal{A}_l(\{g_i^l\}_{i \in [N]}, H)$ is indeed the subvector \hat{x}_{D_l} of the vector $\hat{x} = \mathcal{A}(\{f_i\}_{i \in [N]}, H)$ for $l = 1, 2, \dots, r$ and any H . \square

Now we are ready to give the proof for Theorem 3.

Proof (Theorem 3)

Let $\mu' = \mu n$, $\kappa' = \frac{L}{\mu'} = \frac{\kappa}{n}$ and $k = (e + \max\{1, \alpha\}) \log m$. Since $m \geq \max\{\exp(\frac{\alpha}{\max\{1, \alpha\}} e^{\max\{1, \alpha\} + 1}), (e + \max\{1, \alpha\})^2\}$, we can show that $v \equiv \frac{N}{nk} = \frac{m}{(e + \max\{1, \alpha\}) \log m} \geq \frac{e + \max\{1, \alpha\}}{2 \log(e + \max\{1, \alpha\})} \geq \frac{e}{2} > 1$ for any $\alpha \geq 0$.¹¹ For the simplicity of notation, we will only prove Theorem 3 when k and v are both integers. The general case can be proved by a very similar argument only with more sophisticated notations.

We first use the machinery developed by [2, 17, 13] to construct k functions on \mathbb{R}^b where $b = uk$ for any integer $u \geq 1$. In particular, for $i, j = 1, \dots, b$, let $\delta_{i,j}$ be an $b \times b$ matrix with its (i, j) entry being one and others being zeros. Let M_0, M_1, \dots, M_{b-1} be $b \times b$ matrices defined as

$$M_i = \begin{cases} \delta_{1,1} & \text{for } i = 0 \\ \delta_{i,i} - \delta_{i,i+1} - \delta_{i+1,i} + \delta_{i+1,i+1} & \text{for } 1 \leq i \leq b-2 \\ \delta_{b-1,b-1} - \delta_{b-1,b} - \delta_{b,b-1} + \frac{\sqrt{\kappa'+k-1}+3\sqrt{k}}{\sqrt{\kappa'+k-1}+\sqrt{k}} \delta_{b,b} & \text{for } i = b-1. \end{cases}$$

For $s \in [k]$, let $\Sigma_s = \sum_{i=0}^{u-1} M_{ik+s-1}$. For example, when $u = 2$ and $k = 3$ (so $b = 6$), the matrices Σ_s 's are given as follows.

$$\Sigma_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \Sigma_2 = \begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \Sigma_3 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & -1 & \frac{\sqrt{\kappa'+k-1}+3\sqrt{k}}{\sqrt{\kappa'+k-1}+\sqrt{k}} & -1 \end{bmatrix}.$$

We define k functions $p_1, \dots, p_k : \mathbb{R}^b \rightarrow \mathbb{R}$ as follows

$$p_s(w) = \begin{cases} \frac{L}{4} \left[\left(1 - \frac{\mu'}{L}\right) \frac{w^T \Sigma_1 w}{2} - \left(1 - \frac{\mu'}{L}\right) e_1^T w \right] + \frac{\mu'}{2} \|w\|^2 & \text{for } s = 1 \\ \frac{L}{4} \left[\left(1 - \frac{\mu'}{L}\right) \frac{w^T \Sigma_s w}{2} \right] + \frac{\mu'}{2} \|w\|^2 & \text{for } s = 2, \dots, k, \end{cases} \quad (15)$$

where $e_1 = (1, 0, \dots, 0)^T \in \mathbb{R}^b$, and denote their average by $\bar{p} = \frac{1}{k} \sum_{s=1}^k p_s$. According to the condition $\kappa \geq n$, we have $1 - \frac{\mu'}{L} \geq 0$ so that p_s for any $s \in [k]$ and \bar{p} are all μ' -strongly convex functions. It is also easy to show that $\lambda_{\max}(\Sigma_s) \leq 4$ so that ∇p_s has a Lipschitz continuity constant of $L \left(1 - \frac{\mu'}{L}\right) + \mu' = L$ and p_s is L -smooth.

Next, we characterize the optimal solution of $\min_{w \in \mathbb{R}^b} \bar{p}(w)$.

Lemma 2 *Let $h \in \mathbb{R}$ be the smaller root of the equation*

$$h^2 - 2 \left(\frac{\kappa' - 1 + 2k}{\kappa' - 1} \right) h + 1 = 0,$$

namely,

$$h = \frac{\sqrt{\kappa' + k - 1} - \sqrt{k}}{\sqrt{\kappa' + k - 1} + \sqrt{k}}.$$

Then, $w^ = (w_1^*, w_2^*, \dots, w_b^*)^T \in \mathbb{R}^b$ with*

$$w_j^* = h^j, \text{ for } j = 1, 2, \dots, b \quad (16)$$

is the optimal solutions of $\min_{w \in \mathbb{R}^b} \bar{p}(w)$.

¹¹ Here, we use the factor that $\frac{x}{\log x}$ is monotonically increasing on $[e, +\infty)$.

Proof By definition, we provide the following explicit formulation of $\bar{p}(w)$

$$\bar{p}(w) = \frac{L' - \mu'}{4k} \left[\frac{1}{2} w^T \left(\sum_{s=1}^k \Sigma_s \right) w - e_1^T w \right] + \frac{\mu'}{2} \|w\|^2.$$

Observing that

$$\sum_{s=1}^k \Sigma_s = \begin{bmatrix} 2 & -2 & 0 & 0 & 0 & 0 \\ -1 & 2 & -2 & 0 & 0 & 0 \\ 0 & -1 & 2 & -1 & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & -1 & 2 & -1 \\ 0 & 0 & 0 & 0 & -1 & \frac{\sqrt{\kappa'+k-1}+3\sqrt{k}}{\sqrt{\kappa'+k-1}+\sqrt{k}} \end{bmatrix},$$

and following [13], we can show that w^* must satisfy the following optimality conditions

$$\begin{aligned} w_2^* - 2 \left(\frac{\kappa' - 1 + 2k}{\kappa' - 1} \right) w_1^* + 1 &= 0 \\ w_{j+1}^* - 2 \left(\frac{\kappa' - 1 + 2k}{\kappa' - 1} \right) w_j^* + w_{j-1}^* &= 0, \text{ for } j = 2, 3, \dots, b-1 \\ - \left(\frac{\sqrt{\kappa'+k-1}+3\sqrt{k}}{\sqrt{\kappa'+k-1}+\sqrt{k}} + \frac{4k}{\kappa' - 1} \right) w_b^* + w_{b-1}^* &= 0 \end{aligned} \quad (17)$$

We can easily verify that $w_j^* = h^j$ for $j = 1, 2, \dots, b$ satisfy all equations (17) and is the optimal solution of $\min_{w \in \mathbb{R}^b} \bar{p}(w)$. \square

We claim that $\{p_s\}_{s \in [k]}$ has the following property which directly follows our construction.

Lemma 3 *Suppose U is a strict subset of $\{p_s\}_{s \in [k]}$, and q is an arbitrary linear combination of p_s in U . The Hessian of q is a block diagonal matrix where each block has a size of at most k .*

Proof Suppose $p_{s'}$ is not in U for some $s' \in [k]$. Since q is a linear combination of p_s 's in U , according to the construction in (15), the Hessian of q is a linear combination of one diagonal matrix and all Σ_s 's except $\Sigma_{s'}$, which is a tridiagonal matrix. We note that $\Sigma_{s'}$ is the only matrix among all Σ_s 's that has non-zero entries in the positions $(s' - 1 + ik, s' + ik)$ and $(s' + ik, s' - 1 + ik)$ for $i = 0, 1, \dots, u-1$ and these positions are periodically repeated with a period of k . Therefore, without $\Sigma_{s'}$ involved in the linear combination, the tridiagonal Hessian becomes block diagonal with each block of a size at most k . \square

To complete the proof of Theorem 3, the following lemma is critical. This lemma tells us that the property given by Lemma 3 forces the machines to perform a large number of rounds of communication in order to minimize \bar{p} whenever $\{p_s\}_{s \in [k]}$ do not appear together in any machine.

Lemma 4 *Suppose b (or u) is large enough. Let $\{g_i\}_{i \in [N]}$ be functions on \mathbb{R}^b that consists of v copies of $\{p_s\}_{s \in [k]}$ defined as (15) and $(n-1)vk$ zero functions, that is,*

$$g_i(w) = \begin{cases} p_s(w) & \text{if } i = s, s+k, s+2k, \dots, s+(v-1)k \\ 0 & \text{if } i \geq vk+1. \end{cases} \quad (18)$$

Let $\bar{g} = \frac{1}{N} \sum_{i=1}^N g_i$. We have $w^* = \arg \min_{w \in \mathbb{R}^b} \bar{p}(w) = \arg \min_{w \in \mathbb{R}^b} \bar{g}(w)$ where w^* is defined as (16).

Suppose an algorithm $\mathcal{A} \in \mathcal{F}_\alpha$ is applied to $\{g_i\}_{i \in [N]}$. Let \mathcal{E} be the random event that none of the machines has all functions in $\{p_s\}_{s \in [k]}$ (in either S_j or R_j) after the data distribution stage of \mathcal{A} and let $\hat{w} = \mathcal{A}(\{g_i\}_{i \in [N]}, H)$. Then, to ensure $\mathbb{E}[g(\hat{w}) - g(w^*) | \mathcal{E}] \leq \epsilon$, we need $H = \left(\frac{\sqrt{\kappa'+k-1}-\sqrt{k}}{4k\sqrt{k}} \right) \log \left(\frac{\mu \|w^*\|^2}{4\epsilon} \right)$.

Proof Since $\bar{g} = \frac{1}{N} \sum_{i=1}^N g_i = \frac{1}{nk} \sum_{s=1}^k p_s = \frac{\bar{p}}{n}$, we have $w^* = \arg \min_{x \in \mathbb{R}^b} \bar{g}(x) = \arg \min_{x \in \mathbb{R}^b} \bar{p}(x)$ by Lemma 2, where w^* is defined as in (16).

Let $E_0 = \{\mathbf{0}\}$ and E_t be the linear space spanned by the unit vectors e_1, \dots, e_t for $t = 1, \dots, b$. Suppose event \mathcal{E} happens. Every machine will only have a strict subset U of $\{p_s\}_{s \in [k]}$. Lemma 3 guarantees that, under algorithm \mathcal{A} , if machine j starts one round with a set of working vectors $W_j \subset E_t$, then W_j is always contained by the space E_{t+k} after this round. Therefore, we can show that, at the beginning of round ℓ in algorithm \mathcal{A} , if $\cup_{j=1}^m W_j \subset E_t$, then at the end of round ℓ (and at the beginning of round

$\ell + 1$), we have $\cup_j W_j \subset E_{t+kk}$. Using this finding and the fact that $\cup_j W_j = \{\mathbf{0}\} = E_0$ initially, we conclude that, after H rounds in \mathcal{A} , $\cup_j W_j \subset E_{Hk}$. Let $t = Hk$. Since $\hat{w} = \mathcal{A}(\{g_i\}_{i \in [N]}, H)$, we must have $\hat{w} \in E_t$.

By (16), we can show that

$$\|w^*\|^2 = \sum_{j=1}^b (w_j^*)^2 = \sum_{j=1}^b h^{2j} = \frac{h^2(1-h^{2b})}{1-h^2}. \quad (19)$$

Following the analysis in [2, 17, 13] and using the μ' -strong convexity of \bar{p} , we have

$$\mathbb{E}[\bar{p}(\hat{w}) - \bar{p}(w^*) | \mathcal{E}] \geq \frac{\mu'}{2} \mathbb{E}[\|\hat{w} - w^*\|^2 | \mathcal{E}] \geq \frac{\mu'}{2} \sum_{j=t+1}^b \mathbb{E}[(w_j^*)^2 | \mathcal{E}] \geq \frac{\mu'}{2} \frac{h^{2t+2}(1-h^{2b-2t})}{1-h^2} \geq \frac{\mu' \|w^*\|^2 (h^{2t} - h^{2b})}{2(1-h^{2b})}$$

where the second inequality is because $\hat{w} \in E_t$ and the third inequality is due to (19). When b (or u) is large enough, the inequality above implies

$$\mathbb{E}[\bar{p}(\hat{w}) - \bar{p}(w^*) | \mathcal{E}] \geq \frac{\mu' \|w^*\|^2 h^{2t}}{4} = \frac{\mu' \|w^*\|^2}{4} \left(\frac{\sqrt{\kappa' + k - 1} - \sqrt{k}}{\sqrt{\kappa' + k - 1} + \sqrt{k}} \right)^{2t}.$$

Based on this inequality, when $\mathbb{E}[\bar{g}(\hat{w}) - \bar{g}(w^*) | \mathcal{E}] \leq \epsilon$, or equivalently, when $\mathbb{E}[\bar{p}(\hat{w}) - \bar{p}(w^*) | \mathcal{E}] \leq n\epsilon$, we must have

$$\log \left(\frac{\mu' \|w^*\|^2}{4n\epsilon} \right) \leq 2t \log \left(\frac{\sqrt{\kappa' + k - 1} + \sqrt{k}}{\sqrt{\kappa' + k - 1} - \sqrt{k}} \right) \leq 2t \log \left(1 + \frac{2\sqrt{k}}{\sqrt{\kappa' + k - 1} - \sqrt{k}} \right) \leq \frac{4t\sqrt{k}}{\sqrt{\kappa' + k - 1} - \sqrt{k}}$$

which further implies

$$H = \frac{t}{k} \geq \left(\frac{\sqrt{\kappa' + k - 1} - \sqrt{k}}{4k\sqrt{k}} \right) \log \left(\frac{\mu' \|w^*\|^2}{4\epsilon} \right)$$

□

We now complete the proof of Theorem 3 by constructing N special functions $\{f_i\}_{i \in [N]}$ on \mathbb{R}^d with $d = nb$ for a sufficiently large b (or u) based on $\{p_s\}_{s \in [k]}$, so that any algorithm $\mathcal{A} \in \mathcal{F}_\alpha$, when applied to $\{f_i\}_{i \in [N]}$, will need at least the targeted amount of rounds of communication.

We partition the set of indices $[d] = \{1, \dots, d\}$ into n disjoint subsets D_1, D_2, \dots, D_n with $|D_j| = b$ and $D_j = \{b(j-1) + 1, \dots, bj\}$. For any $j \in [n]$ and $s \in [k]$, let $q_{j,s}(x)$ be a function on \mathbb{R}^d such that $q_{j,s}(x) = p_s(x_{D_j})$, which means $q_{j,s}(x)$ only depends on the b coordinates of x indexed by D_j . Therefore, we obtain nk different functions $\{q_{j,s}\}_{j \in [n], s \in [k]}$. Finally, we define $\{f_i\}_{i \in [N]}$ to be a set that consists of v copies of $\{q_{j,s}\}_{j \in [n], s \in [k]}$ (recall that $N = vkn$ and $v \geq 1$ is an integer). Because

$$f(x) = \frac{1}{N} \sum_{i=1}^N f_i(x) = \frac{v}{vkn} \sum_{j=1}^n \sum_{s=1}^k q_{j,s}(x) = \frac{v}{vkn} \sum_{j=1}^n \sum_{s=1}^k p_s(x_{D_j}) = \frac{1}{nk} \sum_{j=1}^n \sum_{s=1}^k p_s(x_{D_j}) = \frac{1}{n} \sum_{j=1}^n \bar{p}(x_{D_j}) \quad (20)$$

and Lemma 2, the optimal solution x^* for (1) with $\{f_i\}_{i \in [N]}$ constructed as above is $x^* = (w^*, w^*, \dots, w^*)^T$ where $w^* \in \mathbb{R}^d$ is defined as (16) and is repeated for n times.

Now, we want to verify that functions $\{f_i\}_{i \in [N]}$ satisfy our assumptions. In fact, we have shown that p_s is L -smooth for each $s \in [k]$. Since f_i is either a zero function or equals $p_s(x_{D_j})$ for some $j \in [n]$ and $s \in [k]$, the function f_i is L -smooth for each $i \in [N]$ as well. Since \bar{p} is μ' -strongly convex (on \mathbb{R}^b) and $\mu' = n\mu$, the function f defined in (1) must be μ -strongly convex (on \mathbb{R}^d) according to the relationship (20).

According to its construction, $\{f_i\}_{i \in [N]}$ are simultaneously decomposable with respect to a partition D_1, \dots, D_n with $D_j = \{d(j-1) + 1, \dots, dj\}$ (see Definition 2). In particular, for any $i \in [N]$, $f_i(x) = \sum_{l=1}^n g_i^l(x_{D_l})$ where $g_i^l \in \{p_s\}_{s \in [k]}$ for exactly one $l \in [n]$ and $g_i^l = 0$ for other l 's. Moreover, for any $l \in [n]$, $\{g_i^l\}_{i \in [N]} = \{g_i\}_{i \in [N]}$ where $\{g_i\}_{i \in [N]}$ are defined as (18) such that $\bar{g}^l \equiv \frac{1}{N} \sum_{i=1}^N g_i^l = \frac{1}{N} \sum_{i=1}^N g_i = \bar{g}$.

By Proposition 3, \mathcal{A} can be decomposed with respect to the same partition D_1, D_2, \dots, D_n into $\mathcal{A}_1, \dots, \mathcal{A}_n \in \mathcal{F}_\alpha$ and \mathcal{A}_l is applied to $\{g_i\}_{i \in [N]}$. Following Definition 1, let S_1, \dots, S_m be the random partition of $[N]$ and R_1, \dots, R_m be set of i.i.d. indices uniformly drawn from $[N]$ with $|R_j| = \alpha n$. Let $S'_j = S_j \cup R_j$. Then, the algorithm \mathcal{A}_l will allocate $\{g_i | i \in S'_j\}$ to machine j and start the computation in rounds.

We now focus on the solution generated by \mathcal{A}_l for any l . For machine j in \mathcal{A}_l , let $Y_{1,j}$ be the number of functions in $\{p_s\}_{s \in [k]}$ (repetitions counted) that are contained in S_j and $Y_{2,j}$ be the number of functions in $\{p_s\}_{s \in [k]}$ (repetitions counted) that are contained in R_j .

Due to (18), the function g_i is not a zero function if and only if $1 \leq i \leq vk = m$. Hence, $Y_{1,j}$ has a hypergeometric distribution where $\text{Prob}(Y_{1,j} = r)$ equals the probability of r successes in n draws, without replacement, from a population of size N that contains exactly m successes. According to Chvatal [4], we have

$$\text{Prob}(Y_{1,j} \geq r) \leq \left(\frac{1}{r}\right)^r \left(\frac{n-1}{n-r}\right)^{n-r}$$

which, when $r = e \log m$, implies

$$\begin{aligned} \text{Prob}(Y_{1,j} \geq e \log m) &\leq \left(\frac{1}{e \log m}\right)^{e \log m} \left(\frac{n-1}{n-e \log m}\right)^{n-e \log m} \\ &= \left(\frac{1}{e \log m}\right)^{e \log m} \left(1 + \frac{e \log m - 1}{n - e \log m}\right)^{n-e \log m} \\ &< \left(\frac{1}{e \log m}\right)^{e \log m} e^{e \log m - 1} \\ &= \frac{1}{e} \left(\frac{1}{\log m}\right)^{e \log m} \\ &\leq \frac{1}{e} \left(\frac{1}{2 \log(e+1)}\right)^{e \log m} \\ &\leq \frac{1}{em^2} \end{aligned} \tag{21}$$

where the second inequality is because $(1 + \frac{1}{x})^x < e$ for any $x > 0$, the third inequality is due to the assumption that $m \geq \max\{\exp(\frac{\alpha}{\max\{1,\alpha\}} e^{\frac{2}{\max\{1,\alpha\}} + 1}), (e + \max\{1,\alpha\})^2\} \geq (e+1)^2$, and the last inequality is because $(2 \log(e+1))^e > e^2$.

On the other hand, we can represent $Y_{2,j} = \sum_{r \in R_j} \mathbf{1}_{r \leq vk}$ which is the sum of αn i.i.d. binary random variables $\mathbf{1}_{r \leq vk}$'s which equal one with a probability of $\frac{vk}{N} = \frac{1}{n}$ and zero with a probability of $1 - \frac{1}{n}$. By Chernoff inequality of multiplicative form, we have

$$\begin{aligned} \text{Prof}(Y_{2,j} \geq \max\{1,\alpha\} \log m) &\leq \left(\frac{e^{\frac{\max\{1,\alpha\} \log m - 1}{\alpha}}}{\left(\frac{\max\{1,\alpha\} \log m}{\alpha}\right)^{\frac{\max\{1,\alpha\} \log m}{\alpha}}}\right)^\alpha \\ &= \frac{e^{\max\{1,\alpha\} \log m - \alpha}}{\left(\frac{\max\{1,\alpha\} \log m}{\alpha}\right)^{\max\{1,\alpha\} \log m}} \\ &= \frac{1}{e^\alpha} \left(\frac{e\alpha}{\max\{1,\alpha\} \log m}\right)^{\max\{1,\alpha\} \log m} \\ &\leq \frac{1}{e^\alpha} \left(\frac{1}{e^{\frac{2}{\max\{1,\alpha\}}}}\right)^{\max\{1,\alpha\} \log m} \\ &= \frac{1}{e^\alpha m^2}, \end{aligned} \tag{22}$$

where the second inequality is because of the assumption that $m \geq \max\{\exp(\frac{\alpha}{\max\{1,\alpha\}} e^{\frac{2}{\max\{1,\alpha\}} + 1}), (e + \max\{1,\alpha\})^2\} \geq \exp(\frac{\alpha}{\max\{1,\alpha\}} e^{\frac{2}{\max\{1,\alpha\}} + 1})$.

Combining (21) and (22) for $j = 1, 2, \dots, m$ and using the union bound, we have

$$\text{Prob}(Y_{1,j} \geq e \log m \text{ for some } j \text{ or } Y_{2,j} \geq \max\{1,\alpha\} \log m \text{ for some } j) \leq \frac{1}{em} + \frac{1}{e^\alpha m} = \frac{1}{(e + e^\alpha)m},$$

which implies

$$\text{Prob}(Y_{1,j} + Y_{2,j} < (e + \max\{1,\alpha\}) \log m \text{ for } j = 1, 2, \dots, m) \geq 1 - \frac{1}{(e + e^\alpha)m}.$$

Therefore, we have shown that, with a probability of at least $1 - \frac{1}{(e + e^\alpha)m}$, all of the sets S'_1, \dots, S'_m contain fewer than $(e + \max\{1,\alpha\}) \log m = k$ functions from $\{p_s\}_{s \in [k]}$ (repetition counted). In other words, with a probability of at least $1 - \frac{1}{(e + e^\alpha)m}$, none of the sets S'_1, \dots, S'_m contains all of the functions

in $\{p_s\}_{s \in [k]}$. If the event that “none of the sets S'_1, \dots, S'_m contains all of functions of $\{p_s\}_{s \in [k]}$ ” (same as the event \mathcal{E} in Lemma 4) indeed happens in \mathcal{A}_l , we call \mathcal{A}_l *bad*. Then, we have actually proved

$$\text{Prob}(\mathcal{A}_l \text{ is bad}) \geq \left(1 - \frac{1}{(e + e^\alpha)m}\right) \geq \left(1 - \frac{1}{(e + e^\alpha)(e + \max\{1, \alpha\})^2}\right).$$

By Proposition 3 and $\{g_i^l\}_{i \in [N]} = \{g_i\}_{i \in [N]}$, after H rounds, the solutions $\hat{x} = \mathcal{A}(\{f_i\}_{i \in [N]}, H) \in \mathbb{R}^d$ and $\hat{w}^l = \mathcal{A}_l(\{g_i^l\}_{i \in [N]}, H) = \mathcal{A}_l(\{g_i\}_{i \in [N]}, H) \in \mathbb{R}^{|D_l|}$ for $l = 1, 2, \dots, n$ satisfy

$$\begin{aligned} \mathbb{E}[f(\hat{x}) - f(x^*)] &= \sum_{l=1}^n \mathbb{E}[\bar{g}^l(\hat{w}^l) - \bar{g}^l(x_{D_l}^*)] \\ &= \sum_{l=1}^n \mathbb{E}[\bar{g}(\hat{w}^l) - \bar{g}(x_{D_l}^*)] \\ &\geq \sum_{l=1}^n \mathbb{E}[\bar{g}(\hat{w}^l) - \bar{g}(x_{D_l}^*) | \mathcal{A}_l \text{ is bad}] \text{Prob}(\mathcal{A}_l \text{ is bad}) \\ &\geq \sum_{l=1}^n \mathbb{E}[\bar{g}(\hat{w}^l) - \bar{g}(x_{D_l}^*) | \mathcal{A}_l \text{ is bad}] \left(1 - \frac{1}{(e + e^\alpha)(e + \max\{1, \alpha\})^2}\right). \end{aligned}$$

Therefore, if $\mathbb{E}[f(\hat{x}) - f(x^*)] \leq \epsilon$, there must exist an $l \in [n]$ such that

$$\mathbb{E}[\bar{g}(\hat{w}^l) - \bar{g}(x_{D_l}^*) | \mathcal{A}_l \text{ is bad}] \left(1 - \frac{1}{(e + e^\alpha)(e + \max\{1, \alpha\})^2}\right) \leq \frac{\epsilon}{n}. \quad (23)$$

When \mathcal{A}_l is bad, after the data distribution stage, none of the m machines in \mathcal{A}_l has all functions in $\{p_s\}_{s \in [k]}$. According to Lemma 4, we know that to ensure (23), \mathcal{A}_l needs

$$\begin{aligned} H &\geq \left(\frac{\sqrt{\kappa' + k - 1} - \sqrt{k}}{4k\sqrt{k}}\right) \log \left(\left(1 - \frac{1}{(e + e^\alpha)(e + \max\{1, \alpha\})^2}\right) \frac{\mu n \|w^*\|^2}{4\epsilon} \right) \\ &\geq \left(\frac{\sqrt{\kappa' - 1}}{4\sqrt{2k}\sqrt{k}}\right) \log \left(\left(1 - \frac{1}{(e + e^\alpha)(e + \max\{1, \alpha\})^2}\right) \frac{\mu n \|w^*\|^2}{4\epsilon} \right) \end{aligned}$$

which is the desired lower bound after plugging in $k = (e + \max\{1, \alpha\}) \log m$. \square

6 Numerical Experiments

In this section, we conduct numerical experiments to compare our DSVRG and DASVRG algorithms with the DisDSCA [27] (with its practical updates) and a distributed implementation of the accelerated gradient method (Accel Grad) by Nesterov [17]. We apply these four algorithms to the ERM problem (3) with three datasets¹²: Covtype, Million Song and Epsilon. According to the types of data, the loss function $\phi(x, \xi)$ in (3) is chosen to be the square loss in ridge regression for Million Song and the logistic loss in logistic regression for the other two datasets. Following the previous work, we map the target variable of year from 1922 \sim 2011 into $[0, 1]$ for the Million Song data. We notice that the original Covtype and Million Song datasets are not very large (Covtype has 62M in the original size and Million Song has 450M in the original size), and both our algorithms and DisDCA can finish quickly on our server. Therefore, to make comparison among these algorithms in a more challenging setting, we conduct experiments using random Fourier features (RFF) [18] on Covtype and Million Song datasets. The RFF is a popular method for solving large-scale kernel methods by generating finite dimensional features, of which the inner product approximate the kernel similarity. We generate RFF corresponding to RBF kernel. Finally, Covtype data has $N = 522,911$ examples $d = 1,000$ features, Million Song data has $N = 463,715$ examples and $d = 2,000$ features. Since the original Epsilon data is large enough (12G), we use its original features.

The experiments are conducted on one server (Intel(R) Xeon(R) CPU E5-2667 v2 3.30GHz) with multiple processes with each process simulating one machine. We first choose the number of processes (machines) to be $m = 5$. To test the performances of algorithms for different condition numbers, we

¹² <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>

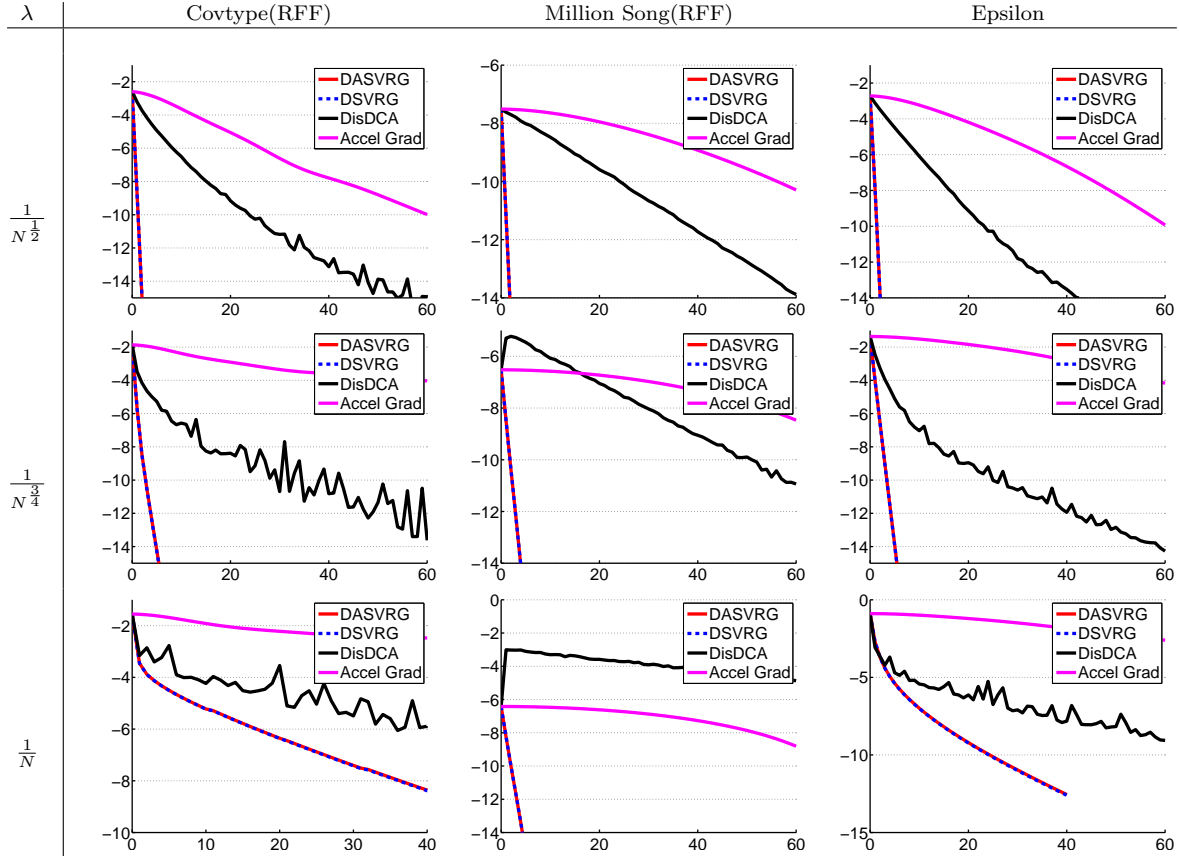


Fig. 1 Comparing the DSVRG and DASVRG methods with DisDCA and the accelerated gradient method (Accel Grad) in rounds.

choose the value of the regularization parameter λ in (3) to be $1/N^{0.5}$, $1/N^{0.75}$ and $1/N$. For each setting, L is computed as $\frac{\max_{i=1,\dots,N} \|a_i\|^2}{\gamma} + \lambda$ where $\frac{1}{\gamma}$ is the Lipschitz continuous constant of $\nabla_x \phi(x, \xi)$, and μ is equal to λ . We implement DSVRG by choosing $\eta = \frac{1}{L}$, $T = 10,000$ and $K = \frac{N}{T}$. For DASVRG, we choose $\eta = \frac{1}{L}$, $T = 10,000$, $K = 1$ and $P = \frac{N}{T}$. In both DSVRG and DASVRG, we directly choose $R_j = S_j$ (so that $|R_j| = \frac{N}{m}$ and $Q = N$) since it saves the time for data allocation and, in practice, gives performances very similar to the performances when R_j is sampled separately. For DisDCA, we use SDCA [21] as the local solver so that it is equivalent to the implementation of CoCoA+ with $\sigma' = m$ and $\gamma = 1$ as in the experiments in [15]. We run SDCA for $T = 10,000$ iterations in each round of DisDCA with $\frac{N}{T}$ rounds in total.

The numerical results are presented in Figure 1 and Figure 2. The horizontal axis presents the number of rounds of communication conducted by algorithms in Figure 1 and presents the parallel runtime (in seconds) used by the algorithms in Figure 2. In both figures, the vertical axis represents the logarithm of optimality gap. According to Figure 1 and Figure 2, the performances of all algorithms get worsen when λ decreases (so the condition number increases). We find that DSVRG and DASVRG have almost identical performances in rounds of communication and they both outperform the other two methods significantly. This shows the merit of our methods when applied to computer clusters with a high communication cost due to significant network delay. DSVRG and DASVRG have slightly different performances in runtime and they outperform the other two methods in Million Song data and obtain a comparable performance on Covtype data. DSVRG and DASVRG do not perform as good as DisDCA in runtime on Epsilon data.

To compare the performances of algorithms under different values of m . We choose the $m = 10$ and 15 and repeat the same experiments on Epsilon data. The numerical results based on the rounds and runtime are shown in Figure 3 and Figure 4 respectively. Similar to the case of $m = 5$, our DSVRG and DASVRG requires fewer rounds to reach the same ϵ -optimal solution but might require longer runtime on some dataset.

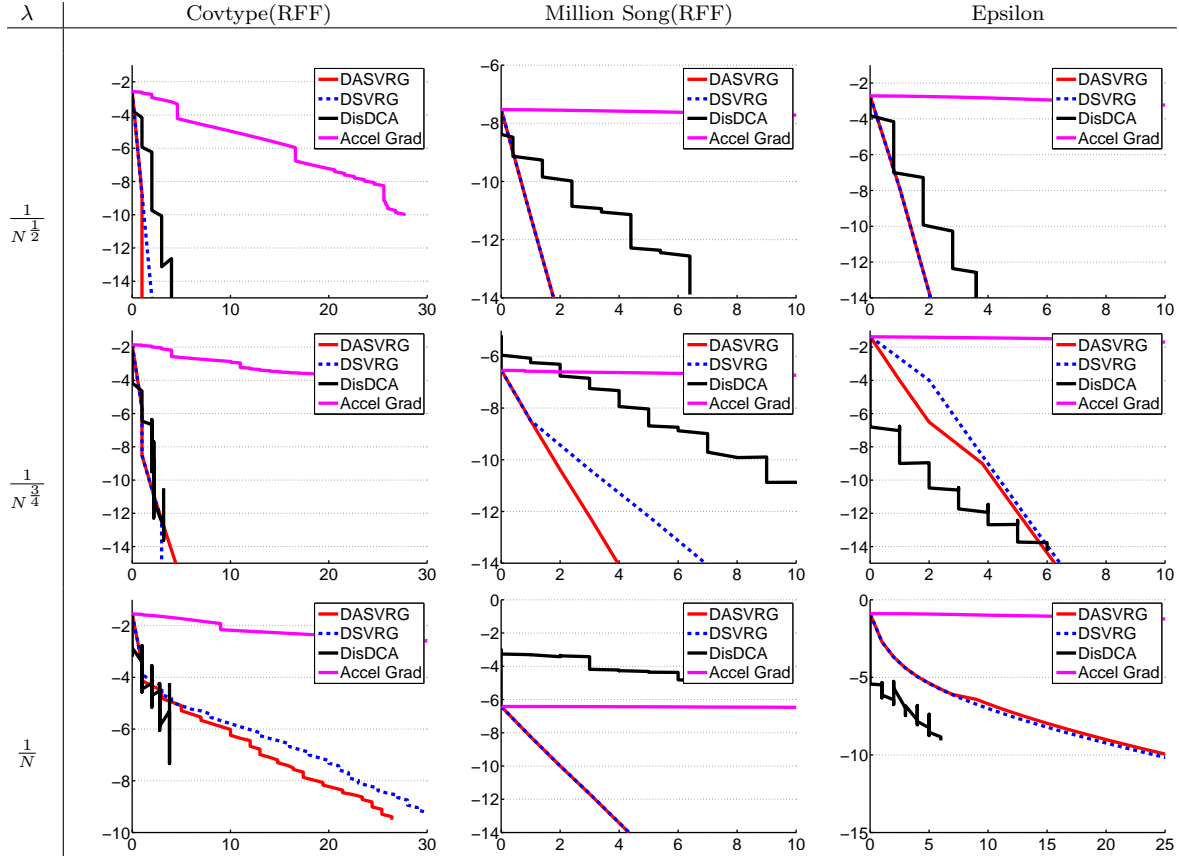


Fig. 2 Comparing the DSVRG and DASVRG methods with DisDCA and the accelerated gradient method (Accel Grad) in runtime.

7 Conclusion

We propose a DSVRG algorithm for minimizing the average of N convex functions which are stored in m machines. Our algorithm is a distributed extension of the existing SVRG algorithm, where we compute the batch gradients in parallel while let machines perform iterative updates in serial. Assuming sufficient memory in each machine, we develop an efficient data allocation scheme to store extra functions in each machine to construct the unbiased stochastic gradient in each iterative update. We provide theoretical analysis on the parallel runtime, the amount and the rounds of communication needed by DSVRG to find an ϵ -optimal solution, showing that it is optimal under all of these three metrics under some practical scenario. Moreover, we proposed a DASVRG algorithm that requires even fewer rounds of communication than DSVRG and almost all existing distributed algorithms using an acceleration strategy by [8] and [14].

Acknowledgements We would like to thank Roy Frostig, Hongzhou Lin, Lin Xiao, and Yuchen Zhang for numerous helpful discussions throughout various stages of this work.

References

1. A. Agarwal and L. Bottou. A Lower Bound for the Optimization of Finite Sums. Technical report, arXiv:1410.0723, October 2014.
2. Y. Arjevani and O. Shamir. Communication Complexity of Distributed Convex Learning and Optimization. Technical report, arXiv:1506.01900, June 2015.
3. M. Braverman, A. Garg, T. Ma, H. L. Nguyen, and D. P. Woodruff. Communication Lower Bounds for Statistical Estimation Problems via a Distributed Data Processing Inequality. *ArXiv e-prints*, June 2015.
4. Vaclav Chvatal. The tail of the hypergeometric distribution. *Discrete Mathematics*, 25:285–287, 1979.
5. Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.




Fig. 3 Comparing the DSVRG and DASVRG methods with DisDCA and the accelerated gradient method (Accel Grad) in rounds.

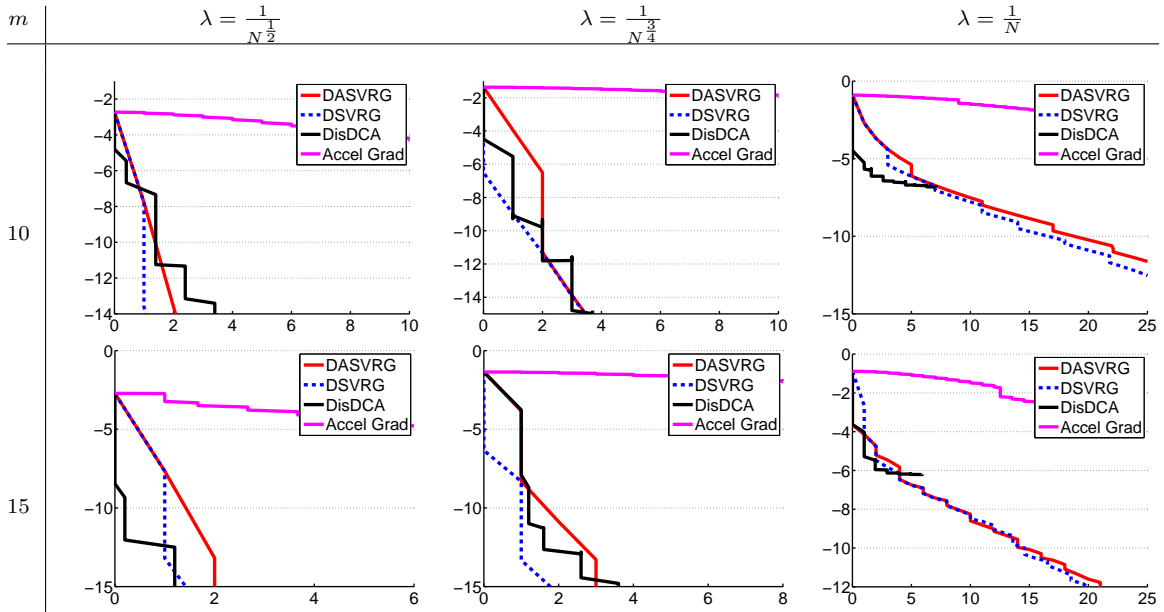


Fig. 4 Comparing the DSVRG and DASVRG methods with DisDCA and the accelerated gradient method (Accel Grad) in runtime.

6. Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in Neural Information Processing Systems*, pages 1646–1654, 2014.
7. Aaron Defazio, Justin Domke, and Tibrio S. Caetano. Finito: A faster, permutable incremental gradient method for big data problems. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pages 1125–1133, 2014.
8. R. Frostig, R. Ge, S. M. Kakade, and A. Sidford. Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization. *ArXiv e-prints*, June 2015.
9. William Gropp, Ewing Lusk, Nathan Doss, and Anthony Skjellum. A high-performance, portable implementation of the mpi message passing interface standard. *Parallel computing*, 22(6):789–828, 1996.
10. Martin Jaggi, Virginia Smith, Martin Takac, Jonathan Terhorst, Sanjay Krishnan, Thomas Hofmann, and Michael I. Jordan. Communication-efficient distributed dual coordinate ascent. In *Advances in Neural Information Processing Systems 27*, pages 3068–3076, 2014.

11. R. Johnson and T. Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems 26*, pages 315–323, 2013.
12. J. Konečný and P. Richtárik. Semi-stochastic gradient descent methods. Technical report, arXiv:1312.1666, 2015.
13. Guanghui Lan and Yi Zhou. An optimal randomized incremental gradient method. Technical report, arXiv:1507.02000, 2015.
14. Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. *arXiv preprint arXiv:1506.02186*, 2015.
15. Chenxin Ma, Virginia Smith, Martin Jaggi, Michael I Jordan, Peter Richtárik, and Martin Takáč. Adding vs. averaging in distributed primal-dual optimization. *arXiv preprint arXiv:1502.03508*, 2015.
16. Julien Mairal. Incremental majorization-minimization optimization with application to large-scale machine learning. *SIAM Journal on Optimization*, 25:829–855, 2015.
17. Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, Boston, 2004.
18. Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *NIPS*, pages 1177–1184, 2007.
19. N. Le Roux, M. Schmidt, and F. Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In *Advances in Neural Information Processing Systems 25*, pages 2672–2680, 2012.
20. Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. Technical report, arXiv:1309.2388, 2013.
21. S. Shalev-Shwartz and T. Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14:567–599, 2013.
22. S. Shalev-Shwartz and T. Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Mathematical Programming*, pages 1–41, 2014.
23. Shai Shalev-Shwartz, Ohad Shamir, Karthik Sridharan, and Nathan Srebro. Stochastic convex optimization. In *International Conference of Machine Learning*, 2009.
24. Ohad Shamir and Nathan Srebro. On distributed stochastic optimization and learning. In *Proceedings of the 52nd Annual Allerton Conference on Communication, Control, and Computing*, 2014.
25. Ohad Shamir, Nati Srebro, and Tong Zhang. Communication-efficient distributed optimization using an approximate newton-type method. In *Proceedings of the 31st International Conference on Machine Learning*, pages 1000–1008, 2014.
26. L. Xiao and T. Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
27. Tianbao Yang. Trading computation for communication: Distributed stochastic dual coordinate ascent. In *Advances in Neural Information Processing Systems*, Tianbao Yang.
28. Yuchen Zhang and Lin Xiao. Communication-efficient distributed optimization of self-concordant empirical loss. *arXiv preprint arXiv:1501.00263*, 2015.