

---

# Graph Oracle Models, Lower Bounds, and Gaps for Parallel Stochastic Optimization

---

Blake Woodworth  
Toyota Technological Institute at Chicago  
blake@ttic.edu

Jialei Wang  
University of Chicago  
jjialei@uchicago.edu

Brendan McMahan  
Google Toyota Technological Institute at Chicago<sup>1</sup>mcmahan@google.com nati@ttic.edu

Abstract

We suggest a general oracle-based framework that captures different parallel stochastic optimization settings described by a dependency graph, and derive generic lower bounds in terms of this graph. We then use the framework and derive lower bounds for several specific parallel optimization settings, including delayed updates and parallel processing with intermittent communication. We highlight gaps between lower and upper bounds on the oracle complexity, and cases where the “natural” algorithms are not known to be optimal.

## 1 Introduction

Recently, there has been great interest in stochastic optimization and learning algorithms that leverage parallelism, including e.g. delayed updates arising from pipelining and asynchronous concurrent processing; synchronous singleinstruction-multiple-data parallelism; and parallelism across distant devices. With the abundance of parallelization settings and associated algorithms, it is important to precisely formulate the problem, which allows us to ask questions such as “is there a better method for this problem than what we have?” and “what is the best we could possibly expect?”

Oracle models have long been a useful framework for formalizing stochastic optimization and learning problems. In an oracle model, we place limits on the algorithm’s access to the optimization objective, but not what it may do with the information it receives. This allows us to obtain sharp lower bounds, which can be used to argue that an algorithm is optimal and to identify gaps between current algorithms and what might be possible. Finding such gaps can be very useful—for example, the gap between the first order optimization lower bound of Nemirovski et al. [20] and the best known algorithms at the time inspired Nesterov’s accelerated gradient descent algorithm [21].

We propose an oracle framework for formalizing different parallel optimization problems. We specify the structure of parallel computation using a dependency graph which indicates how an algorithm accesses the oracle. Each node in the graph corresponds to a single stochastic oracle query, and that query (e.g. the point at which a gradient is calculated) must be computed using only oracle accesses in ancestors of the node. We generally think of each stochastic oracle access as being based on a single data sample, thus involving one or maybe a small number of vector operations.

---

<sup>1</sup> Part of this work was completed while visiting Google

In Section 3 we devise generic lower bounds for parallel optimization problems in terms of simple properties of the associated dependency graph, namely the length of the longest dependency chain and the total number of nodes. In Section 4 we study specific parallel optimization settings in which many algorithms have been proposed, formulate them as graph-based oracle parallel optimization problems, instantiate our lower bounds, and compare them with the performance guarantees of specific algorithms. We highlight gaps between the lower bound and the best known upper bound and also situations where we can devise an optimal algorithm that matches the lower bound, but where this is not the “natural” and typical algorithm used in this settings. The latter indicates either a gap in our understanding of the “natural” algorithm or a need to depart from it.

Previously suggested models Previous work studied communication lower bounds for parallel convex optimization where there are  $M$  machines each containing a local function (e.g. a collection of samples from a distribution). Each machine can perform computation on its own function, and then periodically every machine is allowed to transmit information to the others. In order to prove meaningful lower bounds based on the number of rounds of communication, it is necessary to prevent the machines from simply transmitting their local function to a central machine, or else any objective could be optimized in one round. There are two established ways of doing this. First, one can allow arbitrary computation on the local machines, but restrict the number of bits that can be transmitted in each round. There is work focusing on specific statistical estimation problems that establishes communication lower bounds via information-theoretic arguments [8, 12, 28]. Alternatively, one can allow the machines to communicate real-valued vectors, but restrict the types of computation they are allowed to perform. For instance, Arjevani and Shamir [3] present communication complexity lower bounds for algorithms which can only compute vectors that lie in a certain subspace, which includes e.g. linear combinations of gradients of their local function. Lee et al. [15] assume a similar restriction, but allow the data defining the local functions to be allocated to the different machines in a strategic manner. Our framework applies to general stochastic optimization problems and does not impose any restrictions on what computation the algorithm may perform, and is thus a more direct generalization of the oracle model of optimization.

## 2 The graph-based oracle model

We consider the following stochastic optimization problem

$$\min_{x \in \mathbb{R}^d : \|x\| \leq B} F(x) := \mathbb{E}_{z \sim P} [f(x; z)] \tag{1}$$

The problem (1) captures many important tasks, such as supervised learning, in which case  $f(x; z)$  is the loss of a model parametrized by  $x$  on data instance  $z$  and the goal is to minimize the population risk  $\mathbb{E}[f(x; z)]$ . We assume that  $f(\cdot; z)$  is convex,  $L$ -Lipschitz, and  $H$ -smooth for all  $z$ . We also allow  $f$  to be non-smooth, which corresponds to  $H = \infty$ . A function  $g$  is  $L$ -Lipschitz when  $\|g(x) - g(y)\| \leq L\|x - y\|$  for all  $x, y$ , and it is  $H$ -smooth when it is differentiable and its gradient is  $H$ -Lipschitz. We consider optimization algorithms that use either a stochastic gradient or stochastic prox oracle ( $O_{\text{grad}}$  and  $O_{\text{prox}}$  respectively):

$$O_{\text{grad}}(x, z) = (f(x; z), \nabla f(x; z)) \tag{2}$$

$$\text{prox}_{\text{set}}(x, z) = \text{prox}_{\text{set}}(x, z) \tag{3}$$

where  $\text{prox}_{f(\cdot; z)}(x, \beta) = \underset{y}{\text{argmin}}(4)$

The prox oracle is quite powerful and provides global rather than local information about  $f$ . In particular, querying the prox oracle with  $\beta = 0$  fully optimizes  $f(\cdot; z)$ .

As stated,  $z$  is an argument to the oracle, however there are two distinct cases. In the standard stochastic oracle setting, the algorithm receives an oracle answer corresponding to a random and unknown  $z \sim P$ . We also consider a setting in which the algorithm is allowed to “actively query” the oracle. In this case, the algorithm may either sample  $z \sim P$  or choose a desired  $z$  and receive an oracle answer for that  $z$ . Our lower bounds hold for either type of oracle. Most optimization algorithms only use the fully stochastic oracle, but some require more powerful active queries.

We capture the structure of a parallel optimization algorithm with a directed, acyclic dependency graph  $G$ . Its depth,  $D$ , is the length of the longest directed path, and the size,  $N$ , is the number of nodes. Each node in the graph represents a single stochastic oracle access, and the edges in the graph indicate where the results of that oracle access may be used: only the oracle accesses from *ancestors* of each node are available when issuing a new query. These limitations might arise e.g. due to parallel computation delays or the expense of communicating between disparate machines.

Let  $Q$  be the set of possible oracle queries, with the exact form of queries (e.g.,  $q = x$  vs.  $q = (x, \beta, z)$ ) implied by the setting. Formally, a randomized optimization algorithm that accesses the stochastic oracle  $O$  as prescribed by the graph  $G$  is specified by associating with each node  $v_t$  a query rule  $R_t: (Q, O(Q))^* \times \Xi \rightarrow Q$ , plus a single output rule  $X: (Q, O(Q))^* \times \Xi \rightarrow X$ . We grant all of the nodes access to a source of shared randomness  $\xi \in \Xi$  (e.g. an infinite stream of random bits). The mapping  $R_t$  selects a query  $q_t$  to make at node  $v_t$  using the set of queries and oracle responses in ancestors of  $v_t$ , namely

$$q_t = R_t(q_i, O(q_i) : i \in \text{Ancestors}(v_t)) \quad (5)$$

Similarly, the output rule  $X$  maps from all of the queries and oracle responses to the algorithm’s output as  $x^* = X((q_i, O(q_i) : i \in [N]), \xi)$ . The essential question is: for a class of optimization problems  $(G, O, F)$  specified by a dependency graph  $G$ , a stochastic oracle  $O$ , and a function class  $F$ , what is the best possible guarantee on the expected suboptimality of an algorithm’s output, i.e.

$$\inf_{(R_1, \dots, R_N, X)} \sup_{x, z^*} \mathbb{E}_{x, z^*} [f(x^*; z)] - \min_x \mathbb{E}_z [f(x; z)] \quad (6)$$

In this paper, we consider optimization problems  $(G, O, F_{L, H, B})$  where  $F_{L, H, B}$  is the class of convex,  $L$ -Lipschitz, and  $H$ -smooth functions on the domain  $\{x \in \mathbb{R}^m : \|x\| \leq B\}$  and parametrized by  $z$ , and  $O$  is either a stochastic gradient oracle  $O_{\text{grad}}$  (2) or a stochastic prox oracle  $O_{\text{prox}}$  (3). We consider this function class to contain Lipschitz but non-smooth functions too, which corresponds to  $H = \infty$ . Our function class does not bound the dimension  $m$  of the problem, as we seek to understand the best possible guarantees in terms of Lipschitz and smoothness constants that hold in any dimension. Indeed, there are (typically impractical) algorithms such as center-of-mass methods [7], which might use the dimension in order to significantly reduce the oracle complexity, but at a potentially huge computational cost. Nemirovski [19] studied non-smooth optimization in the case that the dimension is bounded, proving lower bounds in this setting that scale with the  $1/3$ -power of the dimension but have only logarithmic dependence on the suboptimality. We do not analyze strongly convex functions, but the situation is similar and lower bounds can be established via reduction [27].

### 3 Lower bounds

We now provide lower bounds for the optimization problems  $(G, O_{\text{grad}}, F_{L,H,B})$  and  $(G, O_{\text{prox}}, F_{L,H,B})$  in terms of  $L, H, B$ , and the depth and size of  $G$ .

**Theorem 1 (Gradient Oracle Lower Bound).** *Let  $L, B \in (0, \infty)$ ,  $H \in [0, \infty]$ ,  $N \geq D \geq 1$ , let  $G$  be any dependency graph of depth  $D$  and size  $N$  and consider the optimization problem  $(G, O_{\text{grad}}, F_{L,H,B})$ . For any randomized algorithm*

$A = (R_1, \dots, R_N, \hat{X})$ , there exists a distribution  $P$  and a convex,  $L$ -Lipschitz, and  $H$ -smooth function  $f$  on a  $B$ -bounded domain in  $\mathbb{R}^N$  such that

$$\mathbb{E} [f(\hat{X})] \geq \frac{1}{N} \sum_{i=1}^N \mathbb{E} [f(R_i)] - \frac{L}{N} \sqrt{\frac{D}{N}}$$

**Theorem 2 (Prox Oracle Lower Bound).** *Let  $L, B \in (0, \infty)$ ,  $H \in [0, \infty]$ ,  $N \geq D \geq 1$ , let  $G$  be any dependency graph of depth  $D$  and size  $N$  and consider the optimization problem  $(G, O_{\text{prox}}, F_{L,H,B})$ . For any randomized algorithm*

$A = (R_1, \dots, R_N, \hat{X})$ , there exists a distribution  $P$  and a convex,  $L$ -Lipschitz, and  $H$ -smooth function  $f$  on a  $B$ -bounded domain in  $\mathbb{R}^N$  such that

$$\mathbb{E} [f(\hat{X})] \geq \frac{1}{N} \sum_{i=1}^N \mathbb{E} [f(R_i)] - \frac{L}{N} \sqrt{\frac{D}{N}}$$

These are the tightest possible lower bounds in terms of  $D$  and  $N$  in the sense that for all  $D, N$  there are graphs  $G$  and associated algorithms which match the lower bound. Of course, for specific, mostly degenerate graphs they might not be tight. For instance, our lower bound for the graph consisting of a short sequential chain plus a very large number of disconnected nodes might be quite loose due to the artificial inflation of  $N$ . Nevertheless, for many interesting graphs they are tight, as we shall see in Section 4.

Each lower bound has two components: an “optimization” term and a “statistical” term. The statistical term  $\Omega(\frac{LB}{N})$  is well known, although we include a proof of this portion of the bound in Appendix D for completeness. The optimization term depends on the depth  $D$ , and indicates, intuitively, the best suboptimality guarantee that can be achieved by an algorithm using unlimited parallelism but only  $D$  rounds of communication. Arjevani and Shamir [3] also obtain lower bounds in terms of rounds of communication, which are similar to how our lower bounds depend on depth. However they restricted the type of computations that are allowed to the algorithm to a specific class of operations, while we only limit the number of oracle queries and the dependency structure between them, but allow forming the queries in any arbitrary way.

Similar to Arjevani and Shamir [3], to establish the optimization term in the lower bounds, we construct functions that require multiple rounds of sequential oracle accesses to optimize. In the gradient case, we use a single, deterministic function which resembles a standard construction for first order optimization lower bounds. For the prox case, we construct two functions inspired by previous lower bounds for round-based and finite sum optimization [3, 27]. In order to account for randomized algorithms that might leave the span of gradients or proxs returned by the oracle, we use a technique that was proposed by Woodworth and Srebro [26, 27] and refined by Carmon et al. [9]; this portion of the analysis is detailed in Appendix A.

Our lower bounds also have a novel technical element, which is that they apply when both the Lipschitz constant and smoothness are bounded *concurrently*. Consequently, “non-smooth” in the subsequent discussion can be read as simply identifying the case where the  $L$  term achieves the minimum as opposed to the  $H$  term (even if  $H < \infty$ ). This is particularly important in for studying stochastic parallel optimization, since obtaining non-trivial guarantees in a purely stochastic setting requires some sort of control on the magnitude of the gradients (smoothness by itself is not sufficient),

while obtaining parallelization speedups often requires smoothness, and so we would like to ask what is the best that can be done when both Lipschitz and smoothness are controlled. Interestingly, the dependence on both  $L$  and  $H$  in our bounds is tight, even when the other is constrained, which shows that the optimization term cannot be substantially reduced by using both conditions together.

In order to be able to handle Lipschitz and smoothness concurrently, we use only a single construction to handle both smooth and non-smooth objectives. In contrast, previous lower bounds typically treat Lipschitz functions and smooth functions separately, constructing Lipschitz functions that are not at all smooth and vice versa. Our unified construction is thus also simpler in a way (one construction instead of two), and might be of independent interest more broadly.

In the case of the gradient oracle, we “smooth out” a standard non-smooth lower bound construction—this is similar to the approach taken by Agarwal and Hazan [2] when studying higher order oracle complexities. For  $\ell \leq L$  and  $\eta \leq H$ , and orthonormal  $v_1, \dots, v_{D+1}$  drawn uniformly at random, we define the  $\ell$ -Lipschitz but non-smooth function  $\tilde{f}$ , and its  $\ell$ -Lipschitz,  $\eta$ -smooth “ $\eta$ -Moreau envelope” [5]:


$$\tilde{f}(x) = \max_{\|y\| \leq \ell} \langle y, x \rangle + \frac{\eta}{2} \|x\|^2 \tag{7}$$

This defines a distribution over  $f$ 's based on the randomness in the draw of  $v_1, \dots, v_{D+1}$ , and we apply Yao’s minimax principle. In Appendix B, we prove Theorem 1 using this construction.

In the case of the prox oracle, we “straighten out” the smooth construction of Woodworth and Srebro [27]. For fixed constants  $c, \gamma$ , we define the following Lipschitz and smooth scalar function  $\varphi_c$ :

Graph example	With gradient oracle	With gradient and prox oracle
path( $T$ ) (Section 4.1)		
layer( $T, M$ ) (Section 4.2)		$\frac{L}{T}$
delay( $T, \tau$ ) (Section 4.3)	$\left( \frac{L}{\sqrt{T\tau}} \wedge \frac{H\tau^2}{T^2} \right) + \frac{L}{\sqrt{T}}$	$\left( \frac{L\tau}{T} \wedge \frac{H\tau^2}{T^2} \right) + \frac{L}{\sqrt{T}}$
intermittent( $T, K, M$ ) (Section 4.4)	$\geq \left( \frac{L}{\sqrt{KT}} \wedge \frac{H}{K^2T^2} \right) + \frac{L}{\sqrt{MKT}}$ $\leq \frac{L}{\sqrt{KT}} \wedge \left( \frac{H}{T^2} + \frac{L}{\sqrt{MKT}} \right)$ $\wedge \left( \frac{H}{TK} + \frac{L}{\sqrt{MKT}} \right) \log \left( \frac{MKT}{L} \right)$	$\geq \left( \frac{L}{KT} \wedge \frac{H}{K^2T^2} \right) + \frac{L}{\sqrt{MKT}}$ $\leq \frac{L}{\sqrt{KT}} \wedge \left( \frac{L}{T} \wedge \frac{H}{T^2} \right) + \frac{L}{\sqrt{MKT}}$ $\wedge \left( \frac{H}{TK} + \frac{L}{\sqrt{MKT}} \right) \log \left( \frac{MKT}{L} \right)$

Table 1: Summary of upper and lower bounds for stochastic convex optimization of  $L$ -Lipschitz and  $H$ -smooth functions with  $T$  iterations,  $M$  machines, and  $K$  sequential steps per machine. Green indicates lower bounds matched only by “unnatural” methods, red and blue indicates a gap between the upper and lower bounds.



$$\begin{aligned} & \mathbb{1}_{|z| \leq c} \\ & \mathbb{1}_{c < |z| \leq 2c} \end{aligned}$$

<sup>1</sup>. These lower bounds are matched by sequential stochastic gradient descent, yielding a tight complexity of  $\Theta(L/T)$  and the familiar conclusion that SGD is (worst case) optimal in this setting.

$$\varphi^c(z) = \frac{1}{2} \frac{1}{\gamma} z^2 \gamma^2 - |z| - 2c^2 \gamma^2 - 2c^2 2|z c| > \gamma |z| \leq \gamma \quad (8)$$

For  $P = \text{Uniform}\{1,2\}$  and orthonormal  $v_1, \dots, v_{2D}$  drawn uniformly at random, we define

$$\text{[Redacted]$$


$$\text{[Redacted]} \quad (10)$$

Again, this defines a distribution over  $f$ 's based on the randomness in the draw of  $v_1, \dots, v_{2D}$  and we apply Yao's minimax principle. In Appendix C, we prove Theorem 2 using this construction.


## 4 Specific dependency graphs

We now use our framework to study four specific parallelization structures. The main results (tight complexities and gaps between lower and upper bounds) are summarized in Table 1. For simplicity and without loss of generality, we set  $B = 1$ , i.e. we normalize the optimization domain to be  $\{x \in \mathbb{R}^m : \|x\| \leq 1\}$ . All stated upper and lower bounds are for the expected suboptimality  $E[F(x^*)] - F(x^*)$  of the algorithm's output.

### 4.1 Sequential computation: the path graph



We begin with the simplest model, that of sequential computation captured by the path graph of length  $T$  depicted above. The ancestors of each vertex  $v_i$ ,  $i = 1 \dots T$  are all the preceding vertices  $\{v_1, \dots, v_{i-1}\}$ . The sequential model is of course well studied and understood. To see how it fits into our framework: A path graph of length  $T$  has a depth of  $D = T$  and size of  $N = T$ , thus with either gradient or prox oracles, the statistical term is dominant in Theorems 1 and  $\sqrt{\dots}$



### 4.2 Simple parallelism: the layer graph

We now turn to a model in which  $M$  oracle queries can be made in parallel, and the results are broadcast for use in making the next batch of  $M$  queries. This corresponds to synchronized parallelism and fast communication between processors. The model is captured by a layer graph of width  $M$ , depicted above for  $M = 3$ . The graph consists of  $T$  layers  $i = 1, \dots, T$  each with  $M$  nodes  $v_{t,1}, \dots, v_{t,m}$  whose ancestors include  $v_{t^0,i}$  for all  $t^0 < t$  and  $i \in [M]$ . The graph has a depth of  $D = T$  and size of  $N = MT$ . With a stochastic gradient oracle, Theorem 1 yields a lower bound of:


$$\text{[Redacted]} \quad (11)$$

which is matched by accelerated mini-batch SGD (A-MB-SGD) [10, 14], establishing the optimality of A-MB-SGD in this setting. For sufficiently smooth objectives, the same algorithm is also optimal even if prox access is allowed, since Theorem 2 implies a lower bound of:

$$\text{[Redacted]} \quad (12)$$

That is, for smooth objectives, having access to a prox oracle does not improve the optimal complexity over just using gradient access. However, for non-smooth or insufficiently smooth objectives, there is a gap between (11)

and (12). An optimal algorithm, smoothed A-MB-SGD, uses the prox oracle in order to calculate gradients of the Moreau envelope of  $f(x; z)$  (cf. Proposition 12.29 of [5]), and then performs A-MB-SGD on the smoothed objectives. This yields a suboptimality guarantee that precisely matches (12), establishing that the lower bound from Theorem (2) is tight for the layer graph, and that smoothed A-MB-SGD is optimal. An analysis of the smoothed A-MB-SGD algorithm is provided in Appendix E.1.



### 4.3 Delayed updates

We now turn to a delayed computation model that is typical in many asynchronous parallelization and pipelined computation settings, e.g. when multiple processors or machines are working asynchronously, reading iterates, taking some time to perform the oracle accesses and computation, then communicating the results back (or updating the iterate accordingly) [1, 6, 16, 18, 24]. This is captured by a “delay graph” with  $T$  nodes  $v_1, \dots, v_T$  and delays  $\tau_t$  for the response to the oracle query performed at  $v_t$  to become available. Hence,  $\text{Ancestors}(v_t) = \{v_s \mid s + \tau_s \leq t\}$ . Analysis is typically based on the delays being bounded, i.e.  $\tau_t \leq \tau$  for all  $t$ . The depiction above corresponds to  $\tau_t = 2$ , and  $\tau_t = 1$  is simply the path graph. With constant delays  $\tau_t = \tau$ , the delay graph has depth  $D \leq T/\tau$  and size  $N = T$ , so Theorem 1 gives the following lower bound when using a gradient oracle:

$$\text{[Redacted Equation (13)]} \tag{13}$$

Delayed SGD, with updates  $x_t \leftarrow x_{t-1} - \eta_t \nabla f(x_{t-\tau}; z)$ , is a natural algorithm in this setting. Under the bounded delay assumption the best guarantee we are aware of for delayed update SGD is (see [11] improving over [1])

$$\text{[Redacted Equation (14)]} \tag{14}$$

This result is significantly worse than the lower bound 13 and quite disappointing. It does not provide for a  $1/T^2$  accelerated optimization rate, but even worse, compared to non-accelerated SGD it suffers a slowdown quadratic in the delay, compared to the linear slowdown we would expect. In particular, the guarantee  $\sqrt{\text{[Redacted]}}$  (14) only allows maximum delay

of  $\tau = O(T^{1/4})$  in order to attain the optimal statistical rate  $\Theta(L/T)$ , whereas the lower bound allows a delay up to  $\tau = O(T^{3/4})$ .


This raises the question of whether a different algorithm can match the lower bound (13). The answer is affirmative, but it requires using an “unnatural” algorithm, which simulates a mini-batch approach in what seems an unnecessarily wasteful way. We refer to this as a “wait-and-collect” approach: it works in  $T/(2\tau)$  stages, each stage consisting of  $2\tau$  iterations (i.e. nodes or oracle accesses). In stage  $i$ ,  $\tau$  iterations are used to obtain  $\tau$  stochastic gradient estimates  $\nabla f(x_i; Z_{2\tau i+j})$ ,  $j = 1.. \tau$  at the same point  $x_i$ . For the remaining  $\tau$  iterations, we wait for all the preceding oracle computations to become available and do not even use our allowed oracle access. We can then finally update the  $x_{i+1}$  using the minibatch of  $\tau$  gradient estimates. This approach is also specified formally as Algorithm 2 in Appendix E.2. Using this approach, we can perform  $T/(2\tau)$  A-MB-SGD updates with a minibatch size of  $\tau$ , yielding a suboptimality guarantee that precisely matches the lower bound (13).

Thus (13) indeed represents the tight complexity of the delay graph with a stochastic gradient oracle, and the wait-and-collect approach is optimal. However, this answer is somewhat disappointing and leaves an intriguing open question: can a more natural, and seemingly more efficient (no wasted oracle accesses) delayed update SGD algorithm also match the lower bound? An answer to this question has two parts: first, does the delayed update

SGD truly suffer from a  $\tau^2$  slowdown as indicated by (14), or does it achieve linear degradation and a speculative guarantee of

$$\text{[Redacted]}$$
(15)

Second, can delayed update SGD be accelerated to achieve the optimal rate (13). We note that concurrent with our work there has been progress toward closing this gap: Arjevani et al. [4] showed an improved bound matching the non-accelerated (15) for delayed updates (with a fixed delay) on quadratic objectives. It still remains to generalize the result to smooth non-quadratic objectives, handle non-constant bounded delays, and accelerate the procedure so as to improve the rate to  $(\tau/T)^2$ .



#### 4.4 Intermittent communication

We now turn to a parallel computation model which is relevant especially when parallelizing across disparate machines: in each of  $T$  iterations, there are  $M$  machines that, instead of just a single oracle access, perform  $K$  sequential oracle accesses before broadcasting to all other machines synchronously. This communication pattern is relevant in the realistic scenario where local computation is plentiful relative to communication costs (i.e.  $K$  is large). This may be the case with fast processors distributed across different machines, or in the setting of federated learning, where mobile devices collaborate to train a shared model while keeping their respective training datasets local [17].

This is captured by a graph consisting of  $M$  parallel chains of length  $TK$ , with cross connections between the chains every  $K$  nodes. Indexing the nodes as  $v_{t,m,k}$ , the nodes  $v_{t,m,1} \rightarrow \dots \rightarrow v_{t,m,K}$  form a chain, and  $v_{t,m,K}$  is connected to  $v_{t+1,m^0,1}$  for all  $m^0 = 1..M$ . This graph generalizes the layer graph by allowing  $K$  sequential oracle queries between each complete synchronization;  $K = 1$  recovers the layer graph, and the depiction above corresponds to  $K = M = 3$ . We refer to the computation between each synchronization step as a (communication) round.

The depth of this graph is  $D = TK$  and the size is  $N = TKM$ . Focusing on the stochastic gradient oracle (the situation is similar for the prox oracle, except with the potential of smoothing a non-smooth objective, as discussed in Section 4.2), Theorem 1 yields the lower bound:

$$\text{[Redacted]}$$
(16)

A natural algorithm for this graph is parallel SGD, where we run an SGD chain on each machine and average iterates during communication rounds, e.g. [17]. The updates are then given by:

$$\text{[Redacted]}$$
(17)

(note that  $x_{t,m,0}$  does not correspond to any node in the graph, and is included for convenience of presentation). Unfortunately, we are not aware of any satisfying analysis of such a parallel SGD approach. Instead, we consider two other algorithms in an attempt to match the lower bound (16). First, we can combine all  $KM$  oracle accesses between communication rounds in order to form a single mini-batch, giving up on the possibility of sequential computation along the “local”  $K$  node sub-paths. Using all  $KM$  nodes to obtain stochastic gradient estimates at the same point, we can perform  $T$  iterations of A-MB-SGD with a mini-batch size of  $KM$ , yielding an upper bound of



$$\dots$$

(18)

This is a reasonable and common approach, and it is optimal (up to constant factors) when  $\dots$  so that the statistical term is limiting. However, comparing (18) to the lower bound (16) we see a gap by a factor of  $K^2$  in the optimization term, indicating the possibility for significant gains when  $K$  is large (i.e. when we can process a large number of examples on each machine at each round). Improving the optimization term by this  $K^2$  factor would allow statistical optimality as long as  $M = O(T^3K^3)$ —this is a very significant difference. In many scenarios we would expect a modest number of machines, but the amount of data on each machine could easily be much more than the number of communication rounds, especially if communication is across a wide area network.

In fact, when  $K$  is large, a different approach is preferable: we can ignore all but a single chain and simply execute  $KT$  iterations of sequential SGD, offering an upper bound of

$$\dots$$

(19)

Although this approach seems extremely wasteful, it actually yields a better guarantee than (18) when  $K \geq \Omega(T^3L^2/H)$ . This is a realistic regime, e.g. in federated learning when computation is distributed across devices, communication is limited and sporadic and so only a relatively small number of rounds  $T$  are possible, but each device already posses a large amount of data. Furthermore, for non-smooth functions, (19) matches the lower bound (16).

Our upper bound on the complexity is therefore obtained by selecting either A-MB-SGD or single-machine sequential SGD, yielding a combined upper bound of

$$\dots$$

(20)

For smooth functions, there is still a significant gap between this upper bound and the lower bound (16). Furthermore, this upper bound is not achieved by a single algorithm, but rather a combination of two separate algorithms, covering two different regimes. This raises the question of whether there is a single, natural algorithm, perhaps an accelerated variant of the parallel SGD updates (17), that at the very least matches (20), and preferably also improves over them in the intermediate regime or even matches the lower bound (16).

Active querying and SVRG All methods discussed so far used fully stochastic oracles, requesting a gradient (or prox computation) with respect to an independently and randomly drawn  $z \sim P$ . We now turn to methods that also make active queries, i.e. draw samples from  $P$  and then repeatedly query the oracle, at different points  $x$ , but on the same samples  $z$ . Recall that all of our lower bounds are valid also in this setting.

With an active query gradient oracle, we can implement SVRG [13, 15] on an intermittent communication graph. More specifically, for an appropriate choice of  $n$  and  $\lambda$ , we apply SVRG to the regularized empirical objective:

$$\dots$$

(21)

---

Algorithm 1 SVRG

---

---

```

Parameters:  $n, S, I$ 
Sample  $z_1, \dots, z_n \sim P$ 
[redacted]
do
Initialize  $x_0 = 0$  for [redacted] (*)
  for  $i = 1, 2, \dots, I = \frac{n}{K}$  do [redacted]
     $j \sim$  [redacted] (**
    [redacted] end for
    Uniform $\{1, \dots, I\}$ 
  end for
Return
 $x_S$ 

```

---

To do so, we first pick a sample  $\{z_1, \dots, z_n\}$  (without actually querying the oracle). As indicated by Algorithm 1, we then alternate between computing full gradients on  $\{z_1, \dots, z_n\}$  in parallel (\*), and sequential variance-reduced

stochastic gradient updates in between (\*\*). The full gradient  $g^*$  is computed using  $n$  active queries to the gradient oracle. Since all of these oracle accesses are made at the same point  $x^*$ , this can be fully parallelized across the  $M$  parallel chains of length  $K$  thus requiring  $n/KM$  rounds. The sequential variance-reduced stochastic gradient updates *cannot* be parallelized in this way, and must be performed using queries to the gradient oracle in just one of the  $M$  available parallel chains, requiring  $I/K$  rounds of synchronization. Consequently, each outer iteration of SVRG requires [redacted] rounds. We analyze this method using [redacted], and [redacted]. Using the analysis of Johnson and Zhang [13], SVRG guarantees that, with an appropriate stepsize, we have  $F_\lambda(x_S) - \min_x F_\lambda(x) \leq 2^{-S}$ ; the value of  $x_S$  on the empirical objective also generalizes to the population, so [redacted] (see [22]). With our choice of parameters, this implies upper bound (see Appendix E.3)

$$[redacted] \tag{22}$$

$\sqrt{\quad}$

These guarantees improve over sequential SGD (17) as soon as  $M > \log^2(TKM/L)$  and  $K > H^2/L^2$ , i.e.  $L/TK < L^2/H$ . This is a very wide regime: we require only a moderate number of machines, and the second condition will typically hold for a smooth loss. Intuitively, SVRG does roughly the same number (up to a factor of two) of sequential updates as in the sequential SGD approach but it uses better, variance reduced updates. The price we pay is in the smaller total sample size since we keep calling the oracle on the same samples. Nevertheless, since SVRG only needs to calculate the “batch” gradient a logarithmic number of times, this incurs only an additional logarithmic factor.

Comparing (18) and (22), we see that SVRG also improves over A-MB-SGD as soon as  $K > T \log(TKM/L)$ , that is if the number of points we are processing on each machine each round is slightly more than the total number of rounds,

which is also a realistic scenario.

To summarize, the best known upper bound for optimizing with intermittent communication using a pure stochastic oracle is (20), which combines two different algorithms. However, with active oracle accesses, SVRG is also possible and the upper bound becomes:

$$\text{[Redacted Equation]} \tag{23}$$

## 5 Summary

Our main contributions in this paper are: (1) presenting a precise formal oracle framework for studying parallel stochastic optimization; (2) establishing tight oracle lower bounds in this framework that can then be easily applied to particular instances of parallel optimization; and (3) using the framework to study specific settings, obtaining optimality guarantees, understanding where additional assumptions would be needed to break barriers, and, perhaps most importantly, identifying gaps in our understanding that highlight possibilities for algorithmic improvement.

Specifically, we highlight the following gaps:

- For non-smooth objectives and a stochastic prox oracle, it is possible to improve performance in the layer graph setting using smoothing and acceleration. It is not clear whether one can match the same optimal performance using a more direct algorithm, e.g. one that averages the answers from the prox oracle.
- In the delay graph setting, delayed update SGD’s guarantee is not optimal. We suggest an alternative optimal algorithm, but it would be interesting and beneficial to understand the true behavior of delayed update SGD and to improve it as necessary to attain optimality.
- With intermittent communication, we show how different methods are better in different regimes, but even combining these methods does not match our lower bound. This sets the stage for studying the optimal complexity in this setting: is our lower bound achievable? Are current methods optimal? Is the true optimal complexity somewhere in between? Even just obtaining a single method that matches the current best performance in all regimes would be a significant advance here.
- With intermittent communication, active queries allow us to obtain better performance in a certain regime. Can we match this performance using pure stochastic queries or is there a real gap between active and pure stochastic queries?

The investigation into optimizing over  $F_{L,H,B}$  in our framework indicates that there is no advantage to the prox oracle for optimizing (sufficiently) smooth functions. This raises the question of what additional assumptions might allow us to leverage the prox oracle, which is intuitively much stronger as it allows global access to  $f(\cdot; z)$ . One option is to assume a bound on the variance of the stochastic oracle i.e.  $E_z[k\nabla f(x; z) - \nabla F(x)]^2 \leq \sigma^2$  which captures the notion that the functions  $f(\cdot; z)$  are somehow related and not arbitrarily different. In particular, if each stochastic oracle access, in each node, is based on a sample of  $b$  data points (thus, a prox operation optimizes a sub-problem of size  $b$ ), we have that  $\sigma^2 \leq L^2/b$ . Initial investigation into the complexity of optimizing over the restricted class  $F_{L,H,B,\sigma}$  (where we also require the above variance bound), reveals a significant theoretical advantage for the prox oracle over the gradient oracle, even for smooth functions. This is an example of how formalizing the optimization

problem gives insight into additional assumptions, in this case low variance, that are necessary for realizing the benefits of a stronger oracle.

## Acknowledgements

The seed for the generic lower bounds of Section 3 stems from a discussion initiated by Adam Smith, for which we are thankful. We would also like to thank Ohad Shamir for helpful discussions.

This work was partially funded by NSF-BSF award 1718970 (“Convex and Non-Convex Distributed Learning”) and a Google Research Award. BW is supported by the NSF Graduate Research Fellowship under award 1754881.

## References

- [1] Alekh Agarwal and John C Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems*, pages 873–881, 2011.
- [2] Naman Agarwal and Elad Hazan. Lower bounds for higher-order convex optimization. *arXiv preprint arXiv:1710.10329*, 2017.
- [3] Yossi Arjevani and Ohad Shamir. Communication complexity of distributed convex learning and optimization. In *Advances in Neural Information Processing Systems*, pages 1756–1764, 2015.
- [4] Yossi Arjevani, Ohad Shamir, and Nathan Srebro. A tight convergence analysis for stochastic gradient descent with delayed updates. 2018.
- [5] Heinz H Bauschke, Patrick L Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 2011. Springer, 2017.
- [6] Dimitri P Bertsekas. *Parallel and distributed computation: numerical methods*, volume 23. Prentice hall Englewood Cliffs, NJ, 1989.
- [7] Robert G Bland, Donald Goldfarb, and Michael J Todd. The ellipsoid method: A survey. *Operations research*, 29(6):1039–1091, 1981.
- [8] Mark Braverman, Ankit Garg, Tengyu Ma, Huy L Nguyen, and David P Woodruff. Communication lower bounds for statistical estimation problems via a distributed data processing inequality. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 1011–1020. ACM, 2016.
- [9] Yair Carmon, John C Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points i. *arXiv preprint arXiv:1710.11606*, 2017.
- [10] Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. Better mini-batch algorithms via accelerated gradient methods. In *Advances in neural information processing systems*, pages 1647–1655, 2011.
- [11] Hamid Reza Feyzmahdavian, Arda Aytekin, and Mikael Johansson. An asynchronous mini-batch algorithm for regularized stochastic optimization. *IEEE Transactions on Automatic Control*, 61(12):3740–3754, 2016.
- [12] Ankit Garg, Tengyu Ma, and Huy Nguyen. On communication cost of distributed statistical estimation and dimensionality. In *Advances in Neural Information Processing Systems*, pages 2726–2734, 2014.

- [13] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, 2013.
- [14] Guanghui Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1-2): 365–397, 2012.
- [15] Jason D Lee, Qihang Lin, Tengyu Ma, and Tianbao Yang. Distributed stochastic variance reduced gradient methods by sampling extra data with replacement. *The Journal of Machine Learning Research*, 18(1):4404–4446, 2017.
- [16] Brendan McMahan and Matthew Streeter. Delay-tolerant algorithms for asynchronous distributed online learning. In *Advances in Neural Information Processing Systems*, pages 2915–2923, 2014.
- [17] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, 2017.
- [18] A Nedic, Dimitri P Bertsekas, and Vivek S Borkar. Distributed asynchronous incremental subgradient methods. *Studies in Computational Mathematics*, 8(C):381–407, 2001.
- [19] Arkadi Nemirovski. On parallel complexity of nonsmooth convex optimization. *Journal of Complexity*, 10(4): 451–463, 1994.
- [20] Arkadii Nemirovski, David Borisovich Yudin, and Edgar Ronald Dawson. Problem complexity and method efficiency in optimization. 1983.
- [21] Yurii Nesterov. A method of solving a convex programming problem with convergence rate  $o(1/k^2)$ . 1983.
- [22] Shai Shalev-Shwartz and Nathan Srebro. Svm optimization: inverse dependence on training set size. In *International Conference on Machine Learning*, pages 928–935, 2008.
- [23] Eric V Slud et al. Distribution inequalities for the binomial law. *The Annals of Probability*, 5(3):404–412, 1977.
- [24] Suvrit Sra, Adams Wei Yu, Mu Li, and Alex Smola. Adadelay: Delay adaptive distributed stochastic optimization. In *Artificial Intelligence and Statistics*, pages 957–965, 2016.
- [25] Jialei Wang, Weiran Wang, and Nathan Srebro. Memory and communication efficient distributed stochastic optimization with minibatch prox. In *Conference on Learning Theory*, 2017.
- [26] Blake Woodworth and Nathan Srebro. Lower bound for randomized first order convex optimization. *arXiv preprint arXiv:1709.03594*, 2017.
- [27] Blake Woodworth and Nati Srebro. Tight complexity bounds for optimizing composite objectives. In *Advances in Neural Information Processing Systems*, pages 3639–3647, 2016.
- [28] Yuchen Zhang, John Duchi, Michael I Jordan, and Martin J Wainwright. Information-theoretic lower bounds for distributed statistical estimation with communication constraints. In *Advances in Neural Information Processing Systems*, pages 2328–2336, 2013.

## A Main lower bound lemma

This analysis closely follows that of previous work, specifically the proof of Theorem 1 in [26] and the proof of Lemma 4 in [9]. There are slight differences in the problem setup between this work and that of previous papers, thus we include the following analysis for completeness and to ensure that all of our results can be verified. We do not claim any technical novelty within this section.

Let  $V = \{v_1, \dots, v_k\}$  be a uniformly random orthonormal set of vectors in  $\mathbb{R}^m$ . All of the probabilities referred to in Appendix A will be over the randomness in the selection of  $V$ . Let  $X = \{x_1, x_2, \dots, x_N\}$  be a set of vectors in  $\mathbb{R}^m$  where  $\|x_i\| \leq 1$  for all  $i \leq N$ . Let these vectors be organized into disjoint subsets  $X_1 \cup X_2 \cup \dots \cup X_k = X$ . Furthermore, suppose that for each  $t \leq k$ , the set  $X_t$  is a deterministic function  $X_t = X_t(X_{<t}, V)$ , so it can also be expressed as  $X_t = X_t(V)$ .

Let  $S_t = X_{\leq t} \cup V_{\leq t}$ , let  $P_t$  be the projection operator onto the span of  $S_t$  and let  $\Pi_t$  be the projection onto the orthogonal complement of the span of  $S_t$ . As in [9, 26], define

$$\begin{aligned} & \left[ \text{[Redacted]} \right] \\ & \left[ \text{[Redacted]} \right] \end{aligned} \quad | \quad (24)$$

Finally, suppose that for each  $t$ ,  $X_t$  is of the form:

$$X_t(V) = X_t(V_{<t} \mathbf{1}_{G_{<t}} + V \mathbf{1}_{\neg G_{<t}}) \quad (25)$$

i.e. conditioned on the event  $G_{<t}$ , it is a deterministic function of  $V_{<t}$  only (and not  $v_t, \dots, v_k$ ). We say that  $\mathbb{P}[G_{<1}] = 1$ , so  $X_1$  is always independent of  $V$ .

First, we connect the events  $G_t$  to a more immediately useful condition

Lemma 1. [cf. Lemma 9 [9], Lemma 1 [26]] For any  $c, k, N, V$ , and [Redacted], let [Redacted] then for each  $t \leq k$  [Redacted] z

The proof of Lemma 1 involves straightforward linear algebra, and we defer it to Appendix A.1. By Lemma 1,  $G_{<t} \subseteq G^0_{<t}$ , therefore the property (25) is implied by

$$\left[ \text{[Redacted]} \right] \quad (26)$$

Now, we state the main result which allows us to prove our lower bounds:

Lemma 2. [cf. Lemma 4 [9], Lemma 4 [26]] For any  $k \geq 1, N \geq 1, c \in (0, 1)$ , and dimension

$$\left[ \text{[Redacted]} \right]$$

if the sets  $X_1, \dots, X_k$  satisfy the property (26) then

$$\left[ \text{[Redacted]} \right]$$

The proof of Lemma 2 relies upon the following, whose proof we defer to Appendix A.1.

Lemma 3. [cf. Lemma 11 [9], Lemma 3 [26]] Let  $R$  be any rotation operator,  $R^T R = I$ , that preserves  $S_{t-1}$ , that is  $Rw = w$  for any  $w \in \text{Span}(S_{t-1})$ . Then the following conditional densities are equal

$$p(V_{\geq t} | G_{<t}, V_{<t}) = p(RV_{\geq t} | G_{<t}, V_{<t})$$

*Proof of Lemma 2.* This closely follows the proof of Lemma 4 [9] and Lemma 4 [26], with small modifications to account for the different setting.

Set  $\dots$ . Then by Lemma 1, since  $X_1, \dots, X_k$  satisfy the property (26)

$$P[\forall t \leq k \forall x \in X_t \forall j \geq t | \langle x, v_j \rangle] \geq P[G_{\leq k}] = \prod_{t \leq k} P[G_t | G_{<t}] \tag{27}$$

Focus on a single term in this product,

$$\dots \tag{28}$$

For any particular  $V_{<t}$ ,

$$\dots \tag{29}$$

$$\dots \tag{30}$$

$$\dots \tag{31}$$

Conditioned on  $G_{<t}$  and  $V_{<t}$ , the set  $X_t = X_t(V_{<t})$  is fixed, as is the set  $S_{t-1}$  and therefore  $\dots$ , so the first term in the inner product is a fixed unit vector. By Lemma 3, the conditional density of  $v_j | G_{<t}, V_{<t}$  is spherically symmetric within the span onto which  $\dots$  projects. Therefore,  $\dots$  is distributed uniformly on the unit sphere in  $\text{Span}(S_{t-1})^\perp$ , which has dimension at least  $\dots$ .

The probability of a fixed vector and a uniform random vector on the unit sphere in  $\mathbb{R}^{m_0}$  having inner product more than  $\sqrt{\alpha}$

$\alpha$  is proportional to the surface area of the "end caps" of the sphere lying above and below circles of radius  $\sqrt{1 - \alpha^2}$ ,

which is strictly smaller than the surface area of a full sphere of radius  $\sqrt{1 - \alpha^2}$ . Therefore, for a given  $x, v_j$

$$\frac{\dots}{\text{SurfaceArea}_{m_0}(1)} = \frac{\dots \sqrt{1 - \alpha^2}}{\text{SurfaceArea}_{m_0}(1 - \alpha^2)} \tag{32}$$

$$\text{[Redacted]} \tag{33}$$

$$\text{[Redacted]} \tag{34}$$

where we used that  $1 - x \leq \exp(-x)$ . Finally, this holds for each  $t, x \in X_t$ , and  $j \geq t$ , so

$$\text{[Redacted]} \tag{35}$$

$$\text{[Redacted]} \tag{36}$$

$$\text{[Redacted]} \tag{37}$$

$$\text{[Redacted]} \tag{38}$$

Where we used that  $\text{[Redacted]}$  for (37). For (38), recall that we chose  $\text{[Redacted]}$ . □

A.1 Proof of Lemmas 1 and 3

Lemma 1. [cf. Lemma 9 [9], Lemma 1 [26]] For any  $c, k, N, V$ , and  $\text{[Redacted]}$ , let  $\text{[Redacted]}$  then for each  $t \leq k$

$$\text{[Redacted]}_z$$

*Proof.* This closely follows the proof of Lemma 9 [9], with slight modification to account for the different problem setup.

For  $t \leq k$  assume  $G_{\leq t}$ . For any  $x \in X_t$  and  $j \geq t$

$$\text{[Redacted]} \tag{39}$$

$$\text{[Redacted]} \tag{40}$$

$$\text{[Redacted]} \tag{41}$$

$$\text{[Redacted]} \tag{42}$$



First, we decomposed  $v_j$  into its  $S_{t-1}$  and  $\perp$  components and applied the triangle inequality. Next we used that  $\|x\| \leq 1$  and that the orthogonal projection operator  $P_t$  is self-adjoint. Finally, we used that the projection operator is non-expansive and the definition of  $G_t$ .

Next, we prove by induction on  $t$  that for all  $t \leq k$  and  $j \geq t$ , the event  $G_{\leq t}$  implies that  $\|P_t v_j\| \leq \alpha$ . As a base case ( $t = 1$ ), observe that, trivially,  $\|P_1 v_j\| \leq \alpha$ . For the inductive step, fix any  $t \leq k$  and  $j \geq t$  and suppose that  $\|P_{t'} v_{j'}\| \leq \alpha$  for all  $t' < t$  and  $j' \geq t'$ . Let  $\hat{P}_t$  project onto  $\text{Span}(S_t \cup X_{t+1})$  (this includes  $X_{t+1}$  in contrast with  $P_t$ ) and let  $\tilde{P}_t$  project onto the orthogonal subspace. Since  $\text{Span}(X_1 \cup X_2 \cup \dots \cup X_{t-1} \cup V_{\leq t-1}) = S_{t-1}$ ,

$$\tilde{P}_t = P_t - \hat{P}_t \tag{43}$$

is a (potentially over-complete) basis for  $S_{t-1}$ . Using the triangle inequality and  $G_{< t}$ , we can therefore expand

$$\|P_t v_j\|^2 = \|\hat{P}_t v_j\|^2 + \|\tilde{P}_t v_j\|^2 \tag{44}$$

$$\leq \alpha^2 + \|\tilde{P}_t v_j\|^2 \tag{45}$$

We must now bound the second term of (45). Focusing on the inner product in the numerator for one particular  $r < t$ :

$$\langle \tilde{P}_t v_j, v_r \rangle = \langle P_t v_j, v_r \rangle - \langle \hat{P}_t v_j, v_r \rangle \tag{46}$$

$$\leq \alpha + \langle \hat{P}_t v_j, v_r \rangle \tag{47}$$

$$\leq \alpha + \sum_{i=1}^{r-1} \langle X_i v_j, v_r \rangle \tag{48}$$

$$\leq \alpha + \sum_{i=1}^{r-1} \|X_i v_j\| \|v_r\| \tag{49}$$

$$\leq 2\alpha^2 \sum_{i=1}^{r-1} |X_i| + |X_r| \alpha^2 \tag{50}$$

$$\leq 2\alpha^2 \sum_{i=1}^{r-1} |X_i| + |X_r| \alpha^2 \tag{51}$$

First, we used that  $\|v_r\| \leq 1$ , then that  $v_r \perp v_j$ . Next, we applied the definition of  $\hat{P}_{r-1}$  and the triangle inequality. To get (49) we use the Cauchy-Schwarz inequality on the first term, and the definition of  $G_r$  for the second. Finally, we use the inductive hypothesis and that  $\alpha \leq 4^{1-N}$ .

We have now upper bounded the inner products in the second term of (45), and it remains to lower bound the norm in the denominator. We can rewrite

$$\|P_t v_j\|^2 = \langle P_t v_j, P_t v_j \rangle \tag{52}$$

$$= \langle v_j, v_j \rangle - \langle \hat{P}_t v_j, v_j \rangle \tag{53}$$

$$\text{[Redacted]} \tag{54}$$

$$\text{[Redacted]} \tag{55}$$

$$\text{[Redacted]} \tag{56}$$

$$\text{[Redacted]} \tag{57}$$

Here we again used [Redacted] followed by an (over)expansion of  $\hat{P}_{r-1}$ . The remaining steps follow from the inductive hypothesis and fact that  $\alpha \leq 4^{1/N}$ . Combining (57) with (51) and returning to (45), we have that

$$\text{[Redacted]} \tag{58}$$

$$\text{[Redacted]} \tag{59}$$

$$\text{[Redacted]} \tag{60}$$

Therefore, for each  $t \leq k$  and  $j \geq t$  an upper bound [Redacted]. Returning now to (42), we have that for any  $t \leq k$ ,  $x \in X_t$ , and  $j \geq t$  the event  $G_{\leq t}$  implies

$$\text{[Redacted]} \tag{61}$$

$$\text{[Redacted]} \tag{62}$$

$$\text{[Redacted]} \tag{63}$$

where we used that [Redacted] □

**Lemma 3.** [cf. Lemma 11 [9], Lemma 3 [26]] *Let  $R$  be any rotation operator,  $R^2 R = I$ , that preserves  $S_{t-1}$ , that is  $Rw = R^2 w = w$  for any  $w \in \text{Span}(S_{t-1})$ . Then the following conditional densities are equal*

$$p(V_{\geq t} | G_{< t}, V_{< t}) = p(RV_{\geq t} | G_{< t}, V_{< t})$$

*Proof.* This closely follows the proof of Lemma 11 [9].

First, we apply Bayes' rule to each density and use the fact that  $RV_{< t} = V_{< t}$ :

$$\text{[Redacted]} \tag{64}$$

$$\text{[Redacted]} \tag{65}$$

Since  $V$  has a spherically symmetric marginal distribution,  $p(V) = p(RV)$ . Therefore, it only remains to show that  $P[G_{<t} | V] = P[G_{<t} | RV]$ . The event  $G_{<t}$  is determined by  $V$  or by  $RV$ , thus both probabilities are either 0 or 1, so it suffices to show  $P[G_{<t} | V] = 1 \iff P[G_{<t} | RV] = 1$ .

Assume first  $P[G_{<t} | V] = 1$ . Then for each  $r < t$ ,  $x \in X_r$ , and  $\text{Span}\{x, \dots, x\}$ , and each set  $X_r$  is a deterministic function of  $V_{<r}$ . Also, observe that for any  $x \in X_r$  and  $j \geq r$ ,

$$\text{Span}\{x, \dots, x\} \subseteq \text{Span}\{x, \dots, x\} \quad (66)$$

where we used that  $\text{Span}\{x, \dots, x\} = \text{Span}\{x, \dots, x\}$ . Therefore, it suffices to show that the sequence  $X_1(RV), \dots, X_t(RV) = X_1(V), \dots, X_t(V)$  when  $P[G_{<t} | V] = 1$ . We prove this by induction.

For the base case, by definition  $X_1(RV) = X_1 = X_1(V)$ . For the inductive step, suppose now that  $X_{r_0}(RV) = X_{r_0}(V)$  for each  $r^0 < r$ . This, plus the fact that  $P[G_{<t} | V] = 1 \implies P[G_{<r} | V] = 1$  together imply that  $P[G_{<r} | RV] = 1$ . Thus,  $X_r(RV) = X_r(RV_{<r}) = X_r(V_{<r})$ . Therefore, we conclude that  $P[G_{<t} | V] = 1 \implies P[G_{<t} | RV] = 1$ , the reverse implication can be proven with a similar argument.  $\square$

## B Proof of Theorem 1

**Theorem 1 (Gradient Oracle Lower Bound).** *Let  $L, B \in (0, \infty)$ ,  $H \in [0, \infty]$ ,  $N \geq D \geq 1$ , let  $G$  be any dependency graph of depth  $D$  and size  $N$  and consider the optimization problem  $(G, O_{\text{grad}}, F_{L, H, B})$ . For any randomized algorithm*

$A = (R_1, \dots, R_N, \hat{X})$ , *there exists a distribution  $P$  and a convex,  $L$ -Lipschitz, and  $H$ -smooth function  $f$  on a  $B$ -bounded domain in  $\mathbb{R}^m$  such that*

$$\text{Span}\{x, \dots, x\} \subseteq \text{Span}\{x, \dots, x\}$$

*Proof.* Assume for now that  $B = 1$ , the lower bound can be established for other values of  $B$  by scaling inputs to our construction. Let

$$\eta = 10(D + 1)^{1.5} \quad (67)$$

and consider the following  $\eta$ -Lipschitz function:

$$\text{Span}\{x, \dots, x\} \subseteq \text{Span}\{x, \dots, x\} \quad (68)$$

where the vectors  $v_1, \dots, v_{D+1}$  are an orthonormal set drawn uniformly at random from the unit sphere in  $\mathbb{R}^m$ . We use the  $\eta$ -Moreau envelope [5] of this function in order to prove our lower bound:

$$\text{Span}\{x, \dots, x\} \subseteq \text{Span}\{x, \dots, x\} \quad (69)$$

The random draw of  $V$  defines a distribution over functions  $f$ . We will lower bound the expected suboptimality of any *deterministic* optimization algorithm's output and apply Yao's minimax principle at the end of the proof.

This function has the following properties:

Lemma 4. The function  $f$  is convex,  $\beta$ -Lipschitz, and  $\eta$ -smooth, with  $\beta \leq L$  and  $\eta \leq H$ .

Furthermore, optimizing  $f$  is equivalent to “finding” the vectors  $v_1, \dots, v_{D+1}$ . In particular, until a point that has a substantial inner product with all of  $v_1, \dots, v_{D+1}$  is found, the algorithm will remain far from the minimum:

Lemma 5. For any  $H, L > 0$ ,  $D \geq 1$ , and orthonormal  $v_1, \dots, v_{D+1}$ , for any  $x$  with

$$\langle x, v_i \rangle \leq \frac{1}{\sqrt{D+1}} \quad \forall i \in \{1, \dots, D+1\}$$

The function also has the property that if  $x$  has a small inner product with  $v_1, \dots, v_{D+1}$ , then the gradient oracle will reveal little information about  $f$  when queried at  $x$ :

Lemma 6. For any  $x$  with  $\langle x, v_i \rangle \leq \frac{1}{\sqrt{D+1}}$  for all  $i \in \{1, \dots, D+1\}$ , for all  $r \geq t$ , both the function value  $f(x)$  and gradient  $\nabla f(x)$  can be calculated from  $v_1, \dots, v_t$  only.

In Appendix A, we studied the situation where orthonormal  $v_1, \dots, v_{D+1}$  are chosen uniformly at random and a sequence of sets of vectors  $X_1, \dots, X_{D+1}$  are generated as

$$X_t = \{x \in \mathbb{R}^D : \langle x, v_i \rangle \leq \frac{1}{\sqrt{D+1}} \quad \forall i \in \{1, \dots, t\}\} \quad (70)$$

$$\text{where } \langle x, v_i \rangle \leq \frac{1}{\sqrt{D+1}} \quad \forall i \in \{1, \dots, D+1\} \quad (71)$$

Take  $X_1$  and consider the dependency graph. Let  $X_1$  be the set of queries made in vertices at depth 1 in the graph (i.e. they have no parents). Let  $X_2$  be the set of queries made in vertices at depth 2 in the graph (i.e. their parents correspond to the queries in  $X_1$ ). Continue in this way for each  $t \leq D$ , and let  $X_{D+1} = \{x^*\}$  correspond to the algorithm’s output, which is allowed to depend on all queries and oracle responses in the graph, and thus has depth  $D + 1$ .

Supposing  $G^{0 \leq t}$ , for all queries  $x \in X_1 \cup \dots \cup X_{t-1}$  and for all  $r \geq t - 1$  we have  $\langle x, v_i \rangle \leq \frac{1}{\sqrt{D+1}}$ . Therefore, by Lemma 6 all of the function evaluations and gradients returned by the stochastic gradient oracle are calculable from  $v_1, \dots, v_{t-1}$  only. Therefore, all of the queries in  $X_t$  are a deterministic function of  $V_{<t}$  (since we are currently considering only deterministic optimization algorithms), so  $X_t$  satisfies the required decomposition property (70). Finally, the queries are required to be in the domain of  $f$ , thus they will have norm bounded by 1.

Therefore, by Lemma 2, when the dimension

$$D \geq \frac{2 \ln(2/\epsilon)}{\epsilon^2} \quad (72)$$

with probability  $1/2$ , all  $x \in X_1 \cup \dots \cup X_{t+1}$  including the algorithm’s output  $x^*$  satisfy  $\langle x, v_i \rangle \leq \frac{1}{\sqrt{D+1}}$  so by Lemma 5

(73) Therefore, by Yao’s minimax principle for any randomized algorithm A

$$\max_{V \sim A} \mathbb{E} [f(x^*)] - \min_{x: \|x\| \leq 1} f(x) \geq \mathbb{E} [f(x^*)] - \min_{x: \|x\| \leq 1} f(x)$$

$$\geq \frac{1}{2} \epsilon$$

The statistical term  $\epsilon$  follows from Lemma 10. □

## B.1 Deferred proofs

Lemma 4. *The function  $f$  is convex,  $\ell$ -Lipschitz, and  $\eta$ -smooth, with  $\ell \leq L$  and  $\eta \leq H$ .*

*Proof.* Since  $\tilde{f}$  is the maximum of  $\ell$ -Lipschitz affine functions, it is convex and  $\ell$ -Lipschitz. Furthermore, by Proposition 12.29 [5],  $f$ , the  $\eta$ -Moreau Envelope of  $\tilde{f}$  is  $\eta$ -smooth and

$$\| \nabla f(x) - \nabla f(y) \| \leq \eta \| x - y \| \tag{75}$$

The minimizing  $y$  satisfies that  $\eta(x - y) \in \partial \tilde{f}(y)$  (where  $\partial \tilde{f}(y)$  denotes the set of subgradients of  $\tilde{f}$  at  $y$ ), and since  $\tilde{f}$  is  $\ell$ -Lipschitz this implies that  $\| \nabla f(x) \| \leq \ell$ . □

Lemma 5. *For any  $H, L > 0$ ,  $D \geq 1$ , and orthonormal  $v_1, \dots, v_{D+1}$ , for any  $x$  with*

$$\| x - \sum_{i=1}^D \langle x, v_i \rangle v_i \| \leq \epsilon$$

*Proof.* First

$$\| \nabla f(x) - \nabla f(y) \| \leq \eta \| x - y \| \tag{76}$$

Now, for an arbitrary point  $x$  such that  $\| x - \sum_{i=1}^D \langle x, v_i \rangle v_i \| \leq \epsilon$ , consider

$$y^* = \operatorname{prox}_{\frac{1}{\eta} \partial \tilde{f}}(x) \tag{77}$$

Since  $y^*$  is the minimizer,  $\eta(x - y^*) \in \partial \tilde{f}(y^*)$  and since  $\tilde{f}$  is  $\ell$ -Lipschitz,  $\| \eta(x - y^*) \| \leq \ell \| x - y^* \|$ . Thus  $\| x - y^* \| \leq \frac{\ell}{\eta} \| x - y^* \|$  and

$$\| x - y^* \| \leq \frac{\ell}{\eta - \ell} \epsilon \tag{78}$$

$$\| \nabla f(x) - \nabla f(y^*) \| \leq \eta \| x - y^* \| \tag{79}$$

$$\| \nabla f(x) - \nabla f(y^*) \| \leq \frac{\eta \ell}{\eta - \ell} \epsilon \tag{80}$$

$$\| \nabla f(x) - \nabla f(y^*) \| \leq \frac{\eta \ell}{\eta - \ell} \epsilon \tag{81}$$

$$\| \nabla f(x) - \nabla f(y^*) \| \leq \frac{\eta \ell}{\eta - \ell} \epsilon \tag{82}$$

Combining (76) and (82), for any  $x$  such that  $\| x - \sum_{i=1}^D \langle x, v_i \rangle v_i \| \leq \epsilon$

$$\| \nabla f(x) - \nabla f(y^*) \| \leq \frac{\eta \ell}{\eta - \ell} \epsilon \tag{83}$$

□

Lemma 6. For any  $x$  with [redacted] for all  $r \geq t$ , both the function value  $f(x)$  and gradient  $\nabla f(x)$  can be calculated from  $v_1, \dots, v_t$  only.

*Proof.* Let  $x$  be a point such that [redacted] for all  $r \geq t$ . By Proposition 12.29 [5]

$$[redacted] \text{prox} [redacted] \tag{84}$$

Since  $f$  is  $L$ -Lipschitz (Lemma 4), [redacted]  $\text{prox}_r(x, \eta)$ . Consequently, for  $y^* = \text{prox}_r(x, \eta)$  we have [redacted] for all  $r \geq t$ . Furthermore,

$$\nabla f(x) = \eta(x - y^*) \in \text{conv} [redacted] \tag{85}$$

For any  $r > t$

$$[redacted] \tag{86}$$

Whereas

$$[redacted] \tag{87}$$

For any  $r > t$  (86) is less than (87), thus no  $r > t$  can be in the argmax in (85). Therefore, using only  $v_1, \dots, v_t$  we can calculate

$$[redacted] \tag{88}$$

$$[redacted] \tag{89}$$

$$[redacted] \tag{90}$$

and

$$\text{prox} [redacted] \tag{91}$$

$$[redacted] \tag{92}$$

from which we get  $\nabla f(x) = \eta(x - \text{prox}_r(x, \eta))$ . □

## C Proof of Theorem 2

Theorem 2 (Prox Oracle Lower Bound). Let  $L, B \in (0, \infty)$ ,  $H \in [0, \infty]$ ,  $N \geq D \geq 1$ , let  $G$  be any dependency graph of depth  $D$  and size  $N$  and consider the optimization problem  $(G, O_{\text{prox}}, F_{L, H, B})$ . For any randomized algorithm

$A = (R_1, \dots, R_N, \hat{X})$ , there exists a distribution  $P$  and a convex,  $L$ -Lipschitz, and  $H$ -smooth function  $f$  on a  $B$ -bounded domain in [redacted] such that

$$[redacted]$$

*Proof.* Without loss of generality, assume  $B = 1$ , the lower bound can be proven for other values of  $B$  by scaling inputs to our construction by  $1/B$ . Let

$$\eta = \min\{H, 2LD\} \tag{93}$$

Define the following scalar function

$$\varphi_c(z) = \begin{cases} 0 & |z| \leq c \\ 2(|z| - c)^2 & |z| \leq 2c \\ 2\gamma^2 - |z| - 2c^2\gamma^2 - 2c^2|z|c & |z| > \gamma \end{cases} \tag{94}$$

It is straightforward to confirm that  $\varphi_c$  is convex,  $2\gamma$ -Lipschitz continuous, and 4-smooth. Let  $P$  be the uniform distribution over  $\{1,2\}$ . Let  $v_1, v_2, \dots, v_{2D}$  be a set of orthonormal vectors drawn uniformly at random and define

$$\tag{95}$$

$$\tag{96}$$

$$\tag{97}$$

$$\tag{98}$$

The random choice of  $V$  determines a distribution over functions  $f(\cdot;1)$  and  $f(\cdot;2)$ . We will lower bound the expectation (over  $V$ ) of the suboptimality of any deterministic algorithm's output, and then apply Yao's minimax principle.

First, we show that the functions  $f(\cdot;1)$  and  $f(\cdot;2)$  are convex,  $L$ -Lipschitz, and  $H$ -smooth:

*Lemma 7.* For any  $H, L \geq 0$ ,  $D \geq 1$ , and orthonormal  $v_1, \dots, v_{2D}$ , and with  $\eta, \gamma, a$ , and  $c$  chosen as in (93),  $f(\cdot;1)$  and  $f(\cdot;2)$  are convex,  $L$ -Lipschitz, and  $H$ -smooth.

Next, we show that optimizing  $F$  is equivalent to "finding" a large number of the vectors  $v_1, \dots, v_{2D}$ :

*Lemma 8.* For any  $H, L \geq 0$ ,  $D \geq 1$ , and orthonormal  $v_1, \dots, v_{2D}$ , and with  $\eta, \gamma, a$ , and  $c$  chosen as in (93), for any  $x$  such that  $\dots$  for all  $r > D$

$$\dots$$

Next, we show that at any point  $x$  such that  $\dots$  for all  $r \geq t$ , the function value, gradient, and prox of  $f(\cdot;1)$  and  $f(\cdot;2)$  at  $x$  are calculable using  $v_1, \dots, v_t$  only:

*Lemma 9.* For any  $x$  such that  $\dots$ , and any  $\beta \geq 0$  the function values, gradients, and proxs  $f(x;1)$ ,  $f(x;2)$ ,  $\nabla f(x;1)$ ,  $\nabla f(x;2)$ ,  $\text{prox}_{f(\cdot;1)}(x, \beta)$ , and  $\text{prox}_{f(\cdot;2)}(x, \beta)$  are calculable using  $\beta, x, v_1, \dots, v_t$  only.

In Appendix A, we studied the situation where orthonormal  $v_1, \dots, v_{2D}$  are chosen uniformly at random and a sequence of sets of vectors  $X_1, \dots, X_{2D}$  are generated as

$$\text{[Redacted]} \tag{99}$$

where

$$\text{[Redacted]} \tag{100}$$

Consider the dependency graph, and let  $X_1$  be the set of queries made in vertices at depth 1 in the graph (i.e. they have no parents). Let  $X_2$  be the set of queries made in vertices at depth 2 in the graph (i.e. their parents correspond to the queries in  $X_1$ ). Continue in this way for each  $t \leq D$ , and then let  $X_{D+1} = \{x^*\}$  correspond to the output of the optimization algorithm, which for now is deterministic.

Suppose  $G^{0 \leq t}$ . Then for all of the queries  $x \in X_1 \cup \dots \cup X_{t-1}$  and for all  $r \geq t - 1$  we have  $\text{[Redacted]}$ . Therefore, by Lemma 9 the function values, gradients, and prox of  $f(\cdot; 1)$  and  $f(\cdot; 2)$  are calculable based only on the query points and  $v_1, \dots, v_{t-1}$ . Therefore, all of the queries in  $X_t$  are a deterministic function of  $V_{<t}$  only so  $X_t$  satisfies the required decomposition property (99). Finally, the queries are required to be in the domain of  $f$ , thus they will have norm bounded by  $B$ .

Therefore, by Lemma 2 for

$$\text{[Redacted]} \tag{101}$$

with probability  $1/2$  for every  $x \in X_1 \cup \dots \cup X_{D+1}$  which includes  $\text{[Redacted]}$ , so by Lemma 8

$$\text{[Redacted]} \tag{102}$$

Therefore,

$$\text{[Redacted]} \tag{103} \text{ deterministic so by Yao's minimax principle, for any randomized algorithm A}$$

$$\text{[Redacted]} \tag{104}$$

The statistical term  $\text{[Redacted]}$  follows from Lemma 10. □

### C.1 Deferred proof

Lemma 7. For any  $H, L \geq 0$ ,  $D \geq 1$ , and orthonormal  $v_1, \dots, v_{2D}$ , and with  $\eta$ ,  $\gamma$ ,  $a$ , and  $c$  chosen as in (93),  $f(\cdot; 1)$  and  $f(\cdot; 2)$  are convex,  $L$ -Lipschitz, and  $H$ -smooth.

*Proof.* The functions  $f(\cdot; 1)$  and  $f(\cdot; 2)$  are a sum of linear functions and  $\varphi_c$ , which is convex; therefore both are convex. Also, the scalar function  $\varphi_c$  is  $2\gamma$ -Lipschitz, so

$$\text{[Redacted]} \tag{105}$$

$$\text{[Redacted]} \tag{106}$$



where we used that [redacted]. Similarly,

$$[redacted]$$

(107) Therefore,  $f(\cdot;1)$  and

$f(\cdot;2)$  are  $L$ -Lipschitz. Furthermore, since  $\varphi_c$  is 4-smooth,

$$[redacted]$$

and

$$[redacted]$$

(108)

therefore, the maximum eigenvalue of  $\nabla^2 f(\cdot;1)$  and  $\nabla^2 f(\cdot;2)$  is at most  $\eta \leq H$ . □

Lemma 8. For any  $H, L \geq 0$ ,  $D \geq 1$ , and orthonormal  $v_1, \dots, v_{2D}$ , and with  $\eta, \gamma, a$ , and  $c$  chosen as in (93), for any  $x$  such that [redacted] for all  $r > D$

$$[redacted]$$

Proof. First, we upper bound  $\min_{x: \|x\| \leq 1} F(x)$ . Recalling that [redacted], define

$$[redacted]$$

(109)

(110)

For this

$$[redacted]$$

[redacted] and with our choice of parameters (93),  $2c = a < \gamma$ , so that [redacted], thus

!

(111)

thus,

$$[redacted]$$

(112)

$$[redacted] \quad 2 \leq r \leq 2D - 1$$

(113)

$$[redacted]$$

(114)

Since  $\|x^*\| \leq 1$  and  $\nabla F(x^*) = 0$ , we conclude

$$[redacted]$$

(115)

Let [redacted]. We will now lower bound

$$\min F(x) = \min_{x \in X_D, \|x\| \leq 1} F(x) \text{ s.t. (116)} \quad [redacted]$$

Introducing dual variables  $\lambda_{D+1}, \dots, \lambda_{2D} \geq 0$ , solving (116) amounts to finding an  $x \in X_D$  and a set of non-negative

$\lambda$ s such that [redacted] and such that [redacted] for each  $r$ . Let

$$[redacted]$$

(117)

Since [redacted] and  $kx^*k \leq 1$  it follows that  $kx_{Dk} \leq 1$ . Furthermore, since [redacted] and  $2c = a < \gamma$ , the gradient

$$[redacted]$$

(118)

$D$  (119)

$$) = 0 \quad 2 \leq r \leq$$

(120)

$$D + 2 \leq r \leq 2D \quad (121)$$

Therefore,

$$[redacted]$$

(122) Combining (115) and (122) we have that

$$[redacted]$$

□

Lemma 9. For any  $x$  such that [redacted], and any  $\beta \geq 0$  the function values, gradients, and proxes  $f(x;1)$ ,  $f(x;2)$ ,  $\nabla f(x;1)$ ,  $\nabla f(x;2)$ ,  $\text{prox}_{f(\cdot,1)}(x,\beta)$ , and  $\text{prox}_{f(\cdot,2)}(x,\beta)$  are calculable using  $\beta, x, v_1, \dots, v_t$  only.

*Proof.* Suppose that  $x$  is a point such that [redacted] for all  $r \geq t$ , and  $\beta \geq 0$ . Therefore, [redacted] for  $r > t$  so

$$[redacted]$$

!

(124)

(125)

(126)

(127)

Thus both  $f(x;1)$  and  $f(x;2)$  can be calculated from  $x, v_1, \dots, v_t$  only. Similarly, [redacted] so

$$\begin{aligned} & \dots \\ & \dots \\ & \dots \end{aligned} \tag{128}$$

$$\dots \tag{129}$$

$$\dots \tag{130}$$

$$\dots \tag{131}$$

Thus  $\nabla f(x;1)$  and  $\nabla f(x;2)$  can also be calculated from  $x, v_1, \dots, v_t$  only.

Now, we consider the prox at such a point  $x$ . Let  $t^0 = t$  if  $t$  is odd, and  $t^0 = t - 1$  if  $t$  is even. Let  $P$  be the projection operator onto  $S = \text{Span}(v_1, \dots, v_{t^0})$  and let  $P^\perp$  be the projection onto the orthogonal subspace,  $S^\perp$ . Then, since  $f(x;1) = f(Px;1) + f(P^\perp x;1)$ , we can decompose the prox:

$$\begin{aligned} & \text{prox}_{f(\cdot;1)}(x, \beta) \\ & = \underset{y}{\text{argmin}} \left( \dots \right) \end{aligned} \tag{133}$$

$$\dots \tag{134}$$

$$\dots \tag{135}$$

$$\dots \tag{136}$$

Where we used that  $\dots$  for all  $r > t^0$ , so setting  $y_2 = P^\perp x$  achieves the minimum since every term in the expression is zero and function is non-negative. The vector  $P^\perp x$  is calculable from  $x, v_1, \dots, v_{t^0} \subseteq x, v_1, \dots, v_t$ , and similarly the second term is a minimization depends only on  $\beta, x, v_1, \dots, v_{t^0} \subseteq \beta, x, v_1, \dots, v_t$ . A nearly identical argument shows that  $\text{prox}_{f(\cdot;2)}(x, \beta)$  has the same property.  $\square$

### D Statistical term

Lemma 10. For any  $L, B > 0$ , there exists a distribution  $\mathcal{P}$ , and an  $L$ -Lipschitz,  $0$ -smooth function  $f$  defined on  $[-B, B]$  such that the output  $\hat{x}$  of any potentially randomized optimization algorithm which accesses a stochastic gradient or prox oracle at most  $N$  times satisfies

$$\dots$$

Proof. Let  $\theta$  and  $p \sim \text{Uniform}\{p_1, p_{-1}\}$  where  $\theta$  and  $\theta$ . Define  $P_p$  as

$$P_{p_p}[Z = 1] = 1 - P_{p_p}[Z = -1] = p \tag{137}$$

Then, let  $f(x; z) = zLx$ , so  $E_{z \sim P_p}[f(x; z)] = (2p - 1)Lx$ . When  $p = p_1$ ,  $(2p - 1) > 0$  so the minimizer is  $x = -B$  with value  $\theta$ , and when  $p = p_{-1}$ ,  $(2p - 1) < 0$  so the minimizer is  $x = B$ , also with value  $\theta$ . Furthermore, if  $p = p_1$  and  $x \geq 0$  then it is at least  $LB$ -suboptimal, and if  $p = p_{-1}$  and  $x \leq 0$  then it is also at least  $LB$ -suboptimal.

Now consider any deterministic optimization algorithm which accesses the gradient or prox oracle  $N$  times. Each gradient or prox oracle response can be simulated using a single  $z \sim P_p$ , so the algorithm's output is  $x^* = x^*(z_1, \dots, z_N) \in [-B, B]$ . Consider

$$E_{p \sim \text{Uniform}\{p_1, p_{-1}\}, z \sim P_p} [(2p - 1)Lx^*(z_1, \dots, z_N) | z_1, \dots, z_N] \\ \theta \text{Uniform}\{p_1, p_{-1}\}, z \sim P_p [\text{sign}(x^*(z_1, \dots, z_N)) = \text{sign}(2p - 1) | z_1, \dots, z_N] \tag{138}$$

Furthermore, the Bayes optimal estimate  $x^*$  of  $\text{sign}(2p - 1)$  is

$$\theta \tag{139}$$

so

$$P_{p \sim \text{Uniform}} \theta \\ \geq P_{p \sim \text{Uniform}} \theta \tag{140}$$

$$\theta \tag{141}$$

This simply bounding Binomial  $N$ ,  $\theta$  requires lower the tail of the  $\theta$  distribution.

Using Theorem 2.1 in [23],

$$\theta \tag{142}$$

where  $\Phi$  is the distribution function of the standard normal. Let  $\theta$ , then  $\theta$  and

$$\theta \tag{143}$$

Therefore, we conclude that

$$E_{p \sim \text{Uniform}\{p_1, p_{-1}\}, z \sim P_p} [(2p - 1)Lx^*(z_1, \dots, z_N) | z_1, \dots, z_N] \geq \theta \sqrt{\frac{1}{4N}} \tag{144}$$

Therefore, by Yao's minimax principle, for any randomized algorithm A

$$\blacksquare \tag{145}$$

□

## E Supplement to Section 4

### E.1 Smoothed accelerated mini-batch SGD

Smoothed accelerated mini-batch SGD is the composition of two ingredients. First, we approximate the non-smooth  $f$  with a smooth surrogate, and then perform accelerated mini-batch SGD on the surrogate [10, 14]. In particular, we use the  $\beta$ -Moreau envelope  $f^{(\beta)}$  of  $f$ :

$$\blacksquare \tag{146}$$

Since  $f$  is  $L$ -Lipschitz,  $f^{(\beta)}$  has the following properties (Proposition 12.29 [5]):

1.  $f^{(\beta)}$  is  $\beta$ -smooth
2.  $\nabla f^{(\beta)}(x; z) = \beta(x - \text{prox}_{f^{(\cdot; z)}}(x, \beta))$
3.  $\blacksquare$  for all  $x$

We use the prox oracle to execute A-MB-SGD on the  $L$ -Lipschitz and  $\beta$ -smooth  $f^{(\beta)}$ , with updates

$$w_t = \alpha y_t + (1 - \alpha)x_t \tag{147}$$

$$\blacksquare \text{prox} \blacksquare \tag{148}$$

$$x_{t+1} = \alpha y_{t+1} + (1 - \alpha)x_t \tag{149}$$

The A-MB-SGD algorithm will converge on  $f^{(\beta)}$  at a rate (see [10, 14])

$$\blacksquare \tag{150}$$

Choosing  $\beta = \min\{LT, H\}$  we conclude

$$\blacksquare \tag{151}$$

$$\blacksquare \tag{152}$$

$$\blacksquare \tag{153}$$

which matches the lower bound in Theorem 2.

### E.2 Wait-and-collect accelerated mini-batch SGD

### E.3 Analysis of technical results in Section 4.4

Applying SVRG under intermittent synchronization graph To apply SVRG method to solve stochastic convex optimization problems under intermittent synchronization graph. We adopt the approach by [15, 25], first we sample  $n$

---

Algorithm 2 "Wait-and-collect" accelerated minibatch SGD

---

```

Initialize  $\hat{x} = \tilde{x} = x_0 = 0$ , parameter  $\alpha$ .
for  $t = 1, 2, \dots, T$  do
  if  $\text{mod}(t, 2\tau + 1) \leq \tau$  then
    Query stochastic gradient at  $\tilde{x}$ .
    Update  $x_t \leftarrow x_{t-1}, \tilde{g} = 0$ .
  else if  $\text{mod}(t, 2\tau + 1) > \tau$  and  $\text{mod}(t, 2\tau + 1) \leq 2\tau$  then
    Update  $x_t \leftarrow x_{t-1}$ .
    Receive noisy gradient  $g_{t-1-\tau}$ , update  $\tilde{g} \leftarrow \tilde{g} + (1/\tau) * g_{t-1-\tau}$  else
  if  $\text{mod}(t, 2\tau + 1) = 0$  then Update  $x_t \leftarrow \tilde{x} - \eta \tilde{g}$ .
  Update  $\hat{x} \leftarrow \alpha \hat{x} + (1 - \alpha)x_t, \tilde{x} \leftarrow \alpha \tilde{x} + (1 - \alpha)x_t$ . end if
end for

```

---

instances  $\{z_1, \dots, z_n\}$  and solve a regularized empirical risk minimization problem based on  $\{z_1, \dots, z_n\}$ :

$$\min_x \frac{1}{n} \sum_{i=1}^n \ell(x; z_i) + \frac{\lambda}{2} \|x\|^2, \quad (154)$$

where  $\lambda$  is the regularization parameter will specified later. We will apply SVRG algorithm on the intermittent synchronization graph to solve above empirical objective (154) to certain sub-optimality. The SVRG method works in stages, at each stage, we first use  $n/KM$  communication rounds to calculate the full gradient of (154) at a reference point  $\tilde{x}$ , and then using a single chain to perform stochastic gradient updates, equipped with  $\nabla^{F^*} \sqrt{(w^*)}$  to reduce the variance. We choose  $\tilde{x}$ , which will makes the objective  $\sqrt{(154)}$  to be at least  $\sqrt{L/(nB)}$ -strongly

convex, thus the condition number of (154) will be bounded by  $O(H/(L/(nB))) = O(HnB/L)$ . The SVRG analysis [13] requires the number of stochastic gradient updates to be scales as the condition number, so here we  $\sqrt{}$

will use  $O(HnB/(LK))$  communication rounds to perform the stochastic updates, since one chain within each communication round has length  $K$ . Let  $\hat{x}^* = \text{argmin}_x \hat{F}_\lambda(x)$ , and let  $\hat{x}_s$  to be the iterate after running the SVRG algorithm for  $s$ -stages. By the standard results of SVRG (Theorem 1 in [13]), we have

$$\mathbb{E} \|\hat{x}_s - \hat{x}^*\|^2 \leq \frac{1}{s} \left( \frac{HnB}{L} + \frac{1}{s} \right)$$

By standard estimation-optimization error decomposition (e.g. Section 4 in [22]), we have

$$\text{[Redacted Equation]}$$

(155)

given [Redacted]. Thus to implement SVRG successfully, we need to choose  $n$  such that the following two conditions are satisfied:

$$\text{[Redacted]} \quad \text{and} \quad \text{[Redacted]}$$

Thus we know by choosing  $n$  below will satisfy above condition:

$$\text{[Redacted]}$$

substitute the scale of  $n$  to (155) we get

$$\text{[Redacted]}$$

and we obtain the desired result.