

Inexact SARAH Algorithm for Stochastic Optimization

Lam M. Nguyen^a, Katya Scheinberg^b, and Martin Takáč^c

^aIBM Research, Thomas J. Watson Research Center, Yorktown Heights, NY, USA; ^bSchool of Operations Research and Information Engineering, Cornell University, Ithaca, NY, USA;

^cDepartment of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA, USA

ARTICLE HISTORY

Compiled August 27, 2020

ABSTRACT

We develop and analyze a variant of the SARAH algorithm, which does not require computation of the exact gradient. Thus this new method can be applied to general expectation minimization problems rather than only finite sum problems. While the original SARAH algorithm, as well as its predecessor, SVRG, require an exact gradient computation on each outer iteration, the inexact variant of SARAH (iSARAH), which we develop here, requires only stochastic gradient computed on a mini-batch of sufficient size. The proposed method combines variance reduction via sample size selection and iterative stochastic gradient updates. We analyze the convergence rate of the algorithms for strongly convex and non-strongly convex cases, under smooth assumption with appropriate mini-batch size selected for each case. We show that with an additional, reasonable, assumption iSARAH achieves the best known complexity among stochastic methods in the case of non-strongly convex stochastic functions.

KEYWORDS

Stochastic gradient algorithms; variance reduction; stochastic optimization; smooth convex

1. Introduction

We consider the problem of stochastic optimization

$$\min_{w \in \mathbb{R}^d} \{F(w) = \mathbb{E}[f(w; \xi)]\}, \quad (1)$$

where ξ is a random variable and f has a Lipschitz continuous gradient. One of the most popular applications of this problem is expected risk minimization in supervised learning. In this case, random variable ξ represents a random data sample (x, y) , or a set of such samples $\{(x_i, y_i)\}_{i \in I}$. We can consider a set of realizations $\{\xi_{[i]}\}_{i=1}^n$ of ξ corresponding to a set of random samples $\{(x_i, y_i)\}_{i=1}^n$, and define $f_i(w) := f(w; \xi_{[i]})$. Then the sample average approximation of $F(w)$, known as empirical risk in supervised

learning, is written as

$$\min_{w \in \mathbb{R}^d} \left\{ F(w) = \frac{1}{n} \sum_{i=1}^n f_i(w) \right\}. \quad (2)$$

Throughout the paper, we assume the existence of unbiased gradient estimator, that is $\mathbb{E}[\nabla f(w; \xi)] = \nabla F(w)$ for any fixed $w \in \mathbb{R}^d$. In addition we assume that there exists a lower bound of function F .

In recent years, a class of variance reduction methods [1, 3, 5, 6, 11, 18] has been proposed for problem (2) which have smaller computational complexity than both, the full gradient descent method and the stochastic gradient method. All these methods rely on the finite sum form of (2) and are, thus, not readily extendable to (1). In particular, SVRG [3] and SARAH [11] are two similar methods that consist of an outer loop, which includes one exact gradient computation at each outer iteration and an inner loop with multiple iterative stochastic gradient updates. The only difference between SVRG and SARAH is how the iterative updates are performed in the inner loop. The advantage of SARAH is that the inner loop itself results in a convergent stochastic gradient algorithm. Hence, it is possible to apply only one-loop of SARAH with sufficiently large number of steps to obtain an approximately optimal solution (in expectation). The convergence behavior of one-loop SARAH is similar to that of the standard stochastic gradient method [11]. The multiple-loop SARAH algorithm matches convergence rates of SVRG in the strongly convex case, however, due to its convergent inner loop, it has an additional practical advantage of being able to use an adaptive inner loop size, as was demonstrated in [11] for details).

A version of SVRG algorithm, SCSG, has been recently proposed and analyzed in [7, 8]. While this method has been developed for (2) it can be directly applied to (1) because the exact gradient computation is replaced with a mini-batch stochastic gradient. The size of the inner loop of SCSG is then set to a geometrically distributed random variable with distribution dependent on the size of the mini-batch used in the outer iteration. In this paper, we propose and analyze an inexact version of SARAH (iSARAH) which can be applied to solve (1). Instead of exact gradient computation, a mini-batch gradient is computed using a sufficiently large sample size. We develop total sample complexity analysis for this method under various convexity assumptions on $F(w)$. These complexity results are summarized in Tables 1-2 and are compared to the result for SCSG from [7, 8] when applied to (1). We also list the complexity bounds for SVRG, SARAH and SCSG when applied to finite sum problem (2).

As SVRG, SCSG and SARAH, iSARAH algorithm consists of the *outer loop*, which performs variance reduction by computing sufficiently accurate gradient estimate, and the *inner loop*, which performs the stochastic gradient updates. If only one outer iteration of SARAH is performed and then followed by sufficiently many inner iterations, we refer to this algorithm as one-loop SARAH. In [11] one-loop SARAH is analyzed and shown to match the complexity of stochastic gradient descent. Here along with multiple-loop iSARAH we analyze one-loop iSARAH, which is obtained from one-loop SARAH by replacing the first full gradient computation with a stochastic gradient based on a sufficiently large mini-batch.

All our complexity results present the bound on the number of iterations it takes to achieve $\mathbb{E}[\|\nabla F(w)\|^2] \leq \epsilon$. These complexity results are developed under the assumption that $f(w, \xi)$ is L -smooth for every realization of the random variable ξ . Table 1 shows the complexity results in the case when $F(w)$, but not necessarily every realization

$f(w, \xi)$, is μ -strongly convex, with $\kappa = L/\mu$ denoting the condition number. Notice that results for one-loop iSARAH (and SARAH) are the same in the strongly convex case as in the non-strongly convex case. The convergence rate of the multiple-loop iSARAH, on the other hand, is better, in terms of dependence on κ , than the rate achieved by SCSG, which is the only other variance reduction method of the type we consider here, that applies to (1). For example, if a small strongly convex perturbation is added to a convex objective function, then $\kappa = \mathcal{O}(\frac{1}{\epsilon})$, and the total complexity of SCSG is $\mathcal{O}((\frac{\kappa}{\epsilon} + \kappa) \log(\frac{1}{\epsilon})) = \mathcal{O}((\frac{1}{\epsilon^2}) \log(\frac{1}{\epsilon}))$ while iSARAH with multiple-loops has $\mathcal{O}(\max\{\frac{1}{\epsilon}, \kappa\} \log(\frac{1}{\epsilon})) = \mathcal{O}((\frac{1}{\epsilon}) \log(\frac{1}{\epsilon}))$. The non-strongly convex case is summarized in Table 2. In this case, the multiple-loop iSARAH achieves the best convergence rate with respect to ϵ among the compared stochastic methods, but this rate is derived under an additional assumption (Assumption 5 – see (34)) which is discussed in detail in Section 3. We note here that Assumption 5 is weaker than the, common in the analysis of stochastic algorithms, assumption that the iterates remain in a bounded set. In recent paper [15] a stochastic line search method is shown to have expected iteration complexity of $\mathcal{O}(\frac{1}{\epsilon})$ under the assumption that the iterates remain in a bounded set, however, no total sample complexity is derived in [15].

For non-strongly convex and non-convex problems, we derive convergence rate for one-loop iSARAH with only requiring $\mathbb{E}[\|\nabla f(w; \xi)\|^2]$ finite at w_* (non-strongly convex) and w_0 (non-convex). This convergence rate matches that of the general stochastic gradient algorithm, i.e. $\mathcal{O}(\frac{1}{\epsilon^2})$, while the one-loop iSARAH can be viewed as a type of stochastic gradient method with momentum. In order to derive the convergence results in the non-strongly convex and non-convex cases, the learning rate of one-loop iSARAH needs to be chosen in the order of $\mathcal{O}(\epsilon)$, which could be practically undesirable. This is a limitation of one-loop stochastic gradient algorithms.

Multiple-loop iSARAH method achieves the state-of-the-art complexity result of $\mathcal{O}(\frac{\sigma}{\epsilon^{3/2}} + \frac{\sigma^2}{\epsilon})$, under bounded variance assumption, which requires that $\mathbb{E}[\|\nabla f(w; \xi) - \nabla F(w)\|^2] \leq \sigma^2$, for some $\sigma > 0$ and $\forall w \in \mathbb{R}^d$. The analysis is a special case of the analysis ProxSARAH [16] and hence we do not derive it here. This results matches those for two new variants of SARAH, SPIDER and SpiderBoost, that have been proposed in [2, 19]. SPIDER [2] uses the SARAH update rule as was originally proposed in [11] and the mini-batch version of SARAH in [14]. SPIDER and SARAH are different in terms of updated iteration, which are $w_{t+1} = w_t - \eta(v_t/\|v_t\|)$ and $w_{t+1} = w_t - \eta v_t$, respectively. Also, SPIDER does not divide into outer loop and inner loop as SARAH does although SPIDER does also perform a full gradient update after a certain fixed number of iterations. A recent technical report [19] provides an improved version of SPIDER called SpiderBoost which allows a larger learning rate.

Note that, in the original SARAH paper, the result for non-strongly convex case is also derived under the assumption $F(w) - F(w_*) \leq M\|\nabla F(w)\|^2 + N$, for some $M > 0$ and $N > 0$, however this assumption was not explicitly stated.

2. The Algorithm

Like SVRG and SARAH, iSARAH consists of the outer loop and the inner loop. The inner loop performs recursive stochastic gradient updates, while the outer loop of SVRG and SARAH compute the exact gradient. Specifically, given an iterate w_0 at the beginning of each outer loop, SVRG and SARAH compute $v_0 = \nabla F(w_0)$. The only difference between SARAH and iSARAH is that the latter replaces the exact gradient computation by a

Table 1.: Comparison results (Strongly convex). Criteria: $\|\nabla F(w_{\mathcal{T}})\|^2 \leq \epsilon$ (*), $F(w_{\mathcal{T}}) - F(w_*) \leq \epsilon$ (**), and $\|w_{\mathcal{T}} - w_*\|^2 \leq \epsilon$ (***) . Assumption: $\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] \leq \sigma_*^2$ (A). Note $\kappa = L/\mu$; “Ad. Asm.” is Additional Assumption

Method	Bound	Problem	Ad. Asm.
SARAH (one-loop)	N/A	N/A	N/A
SARAH (multiple-loop)	$\mathcal{O}\left((n + \kappa) \log\left(\frac{1}{\epsilon}\right)\right)$	(2) (*)	None
SVRG	$\mathcal{O}\left((n + \kappa) \log\left(\frac{1}{\epsilon}\right)\right)$	(2) (*)	None
SCSG	$\mathcal{O}\left(\left(\min\left\{\frac{\sigma_*^2}{\mu\epsilon}, n\right\} + \kappa\right) \log\left(\frac{\max\{n, \kappa\}}{n} \cdot \frac{1}{\epsilon}\right)\right)$	(2) (**)	(A)
SCSG	$\mathcal{O}\left(\left(\frac{\sigma_*^2}{\mu\epsilon} + \kappa\right) \log\left(\frac{1}{\epsilon}\right)\right)$	(1) (**)	(A)
iSARAH (one-loop)	N/A	N/A	N/A
iSARAH (multiple-loop)	$\mathcal{O}\left(\max\left\{\frac{\sigma_*^2}{\epsilon}, \kappa\right\} \log\left(\frac{1}{\epsilon}\right)\right)$	(1) (*)	(A)

Table 2.: Comparison results (Non-strongly convex). Criteria: $\|\nabla F(w_{\mathcal{T}})\|^2 \leq \epsilon$ (*), $F(w_{\mathcal{T}}) - F(w_*) \leq \epsilon$ (**). Assumption: $\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] \leq \sigma_*^2$ (A), $F(w) - F(w_*) \leq M\|\nabla F(w)\|^2 + N$ (B). “Ad. Asm.” is Additional Assumption

Method	Bound	Problem	Ad. Asm.
SARAH (one-loop)	$\mathcal{O}\left(n + \frac{L}{\epsilon^2}\right)$	(2) (*)	None
SARAH (multiple-loop)	$\mathcal{O}\left((n + \max\{LM, \frac{LN}{\epsilon}\}) \log\left(\frac{1}{\epsilon}\right)\right)$	(2) (*)	(B)
SVRG	$\mathcal{O}\left(n + \frac{L\sqrt{n}}{\epsilon}\right)$	(2) (*)	None
SCSG	$\mathcal{O}\left(\min\left\{\frac{\sigma_*^2}{\epsilon^2}, \frac{nL}{\epsilon}\right\}\right)$	(2) (**)	(A)
SCSG	$\mathcal{O}\left(\frac{\sigma_*^2}{\epsilon^2}\right)$	(1) (**)	(A)
iSARAH (one-loop)	$\mathcal{O}\left(\frac{\max\{L, \sigma_*^2\}}{\epsilon} + \frac{\max\{L^2, \sigma_*^4\}}{\epsilon^2}\right)$	(1) (*)	(A)
iSARAH (multiple-loop)	$\mathcal{O}\left(\max\left\{LM, \frac{\max\{LN, \sigma_*^2\}}{\epsilon}\right\} \log\left(\frac{1}{\epsilon}\right)\right)$	(1) (*)	(A) and (B)

gradient estimate based on a sample set of size b .

In other words, given a iterate w_0 and a sample set size b , v_0 is a random vector computed as

$$v_0 = \frac{1}{b} \sum_{i=1}^b \nabla f(w_0; \zeta_i), \quad (3)$$

where $\{\zeta_i\}_{i=1}^b$ are i.i.d.¹ and $\mathbb{E}[\nabla f(w_0; \zeta_i)|w_0] = \nabla F(w_0)$. We have

$$\mathbb{E}[v_0|w_0] = \frac{1}{b} \sum_{i=1}^b \nabla F(w_0) = \nabla F(w_0).$$

The larger b is, the more accurately the gradient estimate v_0 approximates $\nabla F(w_0)$. The key idea of the analysis of iSARAH is to establish bounds on b which ensure sufficient accuracy for recovering original SARAH convergence rate.

The key step of the algorithm is a recursive update of the stochastic gradient estimate

¹Independent and identically distributed random variables. We note from probability theory that if ζ_1, \dots, ζ_b are i.i.d. random variables then $g(\zeta_1), \dots, g(\zeta_b)$ are also i.i.d. random variables for any measurable function g .

Algorithm 1 Inexact SARAH (iSARAH)

Parameters: the learning rate $\eta > 0$ and the inner loop size m , the sample set size b .

Initialize: \tilde{w}_0 .

Iterate:

for $s = 1, 2, \dots, \mathcal{T}$, **do**

$\tilde{w}_s = \text{iSARAH-IN}(\tilde{w}_{s-1}, \eta, m, b)$.

end for

Output: $\tilde{w}_{\mathcal{T}}$.

Algorithm 2 iSARAH-IN(w_0, η, m, b)

Input: $w_0 (= \tilde{w}_{s-1})$ the learning rate $\eta > 0$, the inner loop size m , the sample set size b .

Generate random variables $\{\zeta_i\}_{i=1}^b$ i.i.d.

Compute $v_0 = \frac{1}{b} \sum_{i=1}^b \nabla f(w_0; \zeta_i)$.

$w_1 = w_0 - \eta v_0$.

Iterate:

for $t = 1, \dots, m - 1$, **do**

Generate a random variable ξ_t

$v_t = \nabla f(w_t; \xi_t) - \nabla f(w_{t-1}; \xi_t) + v_{t-1}$.

$w_{t+1} = w_t - \eta v_t$.

end for

Set $\tilde{w} = w_t$ with t chosen uniformly at random from $\{0, 1, \dots, m\}$

Output: \tilde{w}

(SARAH update)

$$v_t = \nabla f(w_t; \xi_t) - \nabla f(w_{t-1}; \xi_t) + v_{t-1}, \quad (4)$$

followed by the iterate update

$$w_{t+1} = w_t - \eta v_t. \quad (5)$$

Let $\mathcal{F}_t = \sigma(w_0, w_1, \dots, w_t)$ be the σ -algebra generated by w_0, w_1, \dots, w_t . We note that ξ_t is independent of \mathcal{F}_t and that v_t is a *biased estimator* of the gradient $\nabla F(w_t)$.

$$\mathbb{E}[v_t | \mathcal{F}_t] = \nabla F(w_t) - \nabla F(w_{t-1}) + v_{t-1}.$$

In contrast, SVRG update is given by

$$v_t = \nabla f(w_t; \xi_t) - \nabla f(w_0; \xi_t) + v_0 \quad (6)$$

which implies that v_t is an *unbiased estimator* of $\nabla F(w_t)$.

The outer loop of iSARAH algorithm is summarized in Algorithm 1 and the inner loop is summarized in Algorithm 2.

As a variant of SARAH, iSARAH inherits the special property that a one-loop iSARAH, which is the variant of Algorithm 1 with $\mathcal{T} = 1$ and $m \rightarrow \infty$, is a convergent

algorithm. In the next section we provide the analysis for both one-loop and multiple-loop versions of i-SARAH.

Convergence criteria. Our iteration complexity analysis aims to bound the total number of stochastic gradient evaluations needed to achieve a desired bound on the gradient norm. For that we will need to bound the number of outer iterations \mathcal{T} which is needed to guarantee that $\|\nabla F(w_{\mathcal{T}})\|^2 \leq \epsilon$ and also to bound m and b . Since the algorithm is stochastic and $w_{\mathcal{T}}$ is random the ϵ -accurate solution is only achieved in expectation. i.e.,

$$\mathbb{E}[\|\nabla F(w_{\mathcal{T}})\|^2] \leq \epsilon. \quad (7)$$

3. Convergence Analysis of iSARAH

3.1. Basic Assumptions

The analysis of the proposed algorithm will be performed under appropriate subset of the following key assumptions.

Assumption 1 (L -smooth). $f(w; \xi)$ is L -smooth for every realization of ξ , i.e., there exists a constant $L > 0$ such that

$$\|\nabla f(w; \xi) - \nabla f(w'; \xi)\| \leq L\|w - w'\|, \quad \forall w, w' \in \mathbb{R}^d. \quad (8)$$

Note that this assumption implies that $F(w) = \mathbb{E}[f(w; \xi)]$ is also L -smooth. The following strong convexity assumption will be made for the appropriate parts of the analysis, otherwise, it would be dropped.

Assumption 2 (μ -strongly convex). The function $F : \mathbb{R}^d \rightarrow \mathbb{R}$, is μ -strongly convex, i.e., there exists a constant $\mu > 0$ such that $\forall w, w' \in \mathbb{R}^d$,

$$F(w) \geq F(w') + \nabla F(w')^T(w - w') + \frac{\mu}{2}\|w - w'\|^2.$$

Under Assumption 2, let us define the (unique) optimal solution of (2) as w_* . Then strong convexity of F implies that

$$2\mu[F(w) - F(w_*)] \leq \|\nabla F(w)\|^2, \quad \forall w \in \mathbb{R}^d. \quad (9)$$

Under strong convexity assumption we will use κ to denote the condition number $\kappa \stackrel{\text{def}}{=} L/\mu$.

Finally, as a special case of the strong convexity with $\mu = 0$, we state the general convexity assumption, which we will use for some of the convergence results.

Assumption 3 (Convex). $f(w; \xi)$ is convex for every realization of ξ , i.e., $\forall w, w' \in \mathbb{R}^d$,

$$f(w; \xi) \geq f(w'; \xi) + \nabla f(w'; \xi)^T(w - w').$$

We note that Assumption 2 does not imply Assumption 3, because the latter applies to all realizations, while the former applied only to the expectation.

Hence in our analysis, depending on the result we aim at, we will require Assumption 3 to hold by itself, or Assumption 2 and Assumption 3 to hold together. We will always

use Assumption 1.

We also assume that, the $\mathbb{E}[\|\nabla f(w; \xi)\|^2]$ at w_* is bounded by some constant.

Assumption 4. *There exists some $\sigma_* > 0$ such that*

$$\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] \leq \sigma_*^2, \quad (10)$$

where w_* is any optimal solution of $F(w)$; and ξ is some random variable.

3.2. Existing Results

We provide some well-known results from the existing literature that support our theoretical analysis as follows. First, we start introducing two standard lemmas in smooth convex optimization ([9]) for a general function f .

Lemma 1 (Theorem 2.1.5 in [9]). *Suppose that f is convex and L -smooth. Then, for any $w, w' \in \mathbb{R}^d$,*

$$f(w) \leq f(w') + \nabla f(w')^T(w - w') + \frac{L}{2}\|w - w'\|^2, \quad (11)$$

$$f(w) \geq f(w') + \nabla f(w')^T(w - w') + \frac{1}{2L}\|\nabla f(w) - \nabla f(w')\|^2, \quad (12)$$

$$(\nabla f(w) - \nabla f(w'))^T(w - w') \geq \frac{1}{L}\|\nabla f(w) - \nabla f(w')\|^2. \quad (13)$$

Note that (11) does not require the convexity of f .

Lemma 2 (Theorem 2.1.11 in [9]). *Suppose that f is μ -strongly convex and L -smooth. Then, for any $w, w' \in \mathbb{R}^d$,*

$$(\nabla f(w) - \nabla f(w'))^T(w - w') \geq \frac{\mu L}{\mu + L}\|w - w'\|^2 + \frac{1}{\mu + L}\|\nabla f(w) - \nabla f(w')\|^2. \quad (14)$$

The following existing results are more specific properties of component functions $f(w; \xi)$.

Lemma 3 ([3]). *Suppose that Assumptions 1 and 3 hold. Then, $\forall w \in \mathbb{R}^d$,*

$$\mathbb{E}[\|\nabla f(w; \xi) - \nabla f(w_*; \xi)\|^2] \leq 2L[F(w) - F(w_*)], \quad (15)$$

where w_* is any optimal solution of $F(w)$.

Lemma 4 (Lemma 1 in [13]). *Suppose that Assumptions 1 and 3 hold. Then, for $\forall w \in \mathbb{R}^d$,*

$$\mathbb{E}[\|\nabla f(w; \xi)\|^2] \leq 4L[F(w) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2], \quad (16)$$

where w_* is any optimal solution of $F(w)$.

Lemma 5 (Lemma 1 in [12]). *Let ξ and $\{\xi_i\}_{i=1}^b$ be i.i.d. random variables with*

$\mathbb{E}[\nabla f(w; \xi_i)] = \nabla F(w)$, $i = 1, \dots, b$, for all $w \in \mathbb{R}^d$. Then,

$$\mathbb{E} \left[\left\| \frac{1}{b} \sum_{i=1}^b \nabla f(w; \xi_i) - \nabla F(w) \right\|^2 \right] = \frac{\mathbb{E}[\|\nabla f(w; \xi)\|^2] - \|\nabla F(w)\|^2}{b}. \quad (17)$$

The proof of this Lemma is in [12].

Lemmas 4 and 5 clearly imply the following result.

Corollary 1. *Suppose that Assumptions 1 and 3 hold. Let ξ and $\{\xi_i\}_{i=1}^b$ be i.i.d. random variables with $\mathbb{E}[\nabla f(w; \xi_i)] = \nabla F(w)$, $i = 1, \dots, b$, for all $w \in \mathbb{R}^d$. Then,*

$$\begin{aligned} \mathbb{E} \left[\left\| \frac{1}{b} \sum_{i=1}^b \nabla f(w; \xi_i) - \nabla F(w) \right\|^2 \right] \\ \leq \frac{4L[F(w) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] - \|\nabla F(w)\|^2}{b}, \end{aligned} \quad (18)$$

where w_* is any optimal solution of $F(w)$.

Based on the above lemmas, we will show in detail how to achieve our main results in the following subsections.

3.3. Special Property of SARAH Update

The most important property of the SVRG algorithm is the variance reduction of the steps. This property holds as the number of outer iteration grows, but it does not hold, if only the number of inner iterations increases. In other words, if we simply run the inner loop for many iterations (without executing additional outer loops), the variance of the steps does not reduce in the case of SVRG, while it goes to zero in the case of SARAH with large learning rate in the strongly convex case. We recall the SARAH update as follows.

$$v_t = \nabla f(w_t; \xi_t) - \nabla f(w_{t-1}; \xi_t) + v_{t-1}, \quad (19)$$

followed by the iterate update:

$$w_{t+1} = w_t - \eta v_t. \quad (20)$$

We will now show that $\|v_t\|^2$ is going to zero in expectation in the *strongly convex* case. These results substantiate our conclusion that SARAH uses more stable stochastic gradient estimates than SVRG.

Proposition 1. *Suppose that Assumptions 1, 2 and 3 hold. Consider v_t defined by (19) with $\eta < 2/L$ and any given v_0 . Then, for any $t \geq 1$,*

$$\begin{aligned} \mathbb{E}[\|v_t\|^2] &\leq \left[1 - \left(\frac{2}{\eta L} - 1 \right) \mu^2 \eta^2 \right] \mathbb{E}[\|v_{t-1}\|^2] \\ &\leq \left[1 - \left(\frac{2}{\eta L} - 1 \right) \mu^2 \eta^2 \right]^t \|v_0\|^2. \end{aligned}$$

The proof of this Proposition can be derived directly from Theorem 1a in [11]. This result implies that by choosing $\eta = \mathcal{O}(1/L)$, we obtain the linear convergence of $\|v_t\|^2$ in expectation with the rate $(1 - 1/\kappa^2)$.

We will provide our convergence analysis in detail in next sub-section. We will divide our results into two parts: the *one-loop* results corresponding to iSARAH-IN (Algorithm 2) and the *multiple-loop* results corresponding to iSARAH (Algorithm 1).

3.4. One-loop (iSARAH-IN) Results

We begin with providing two useful lemmas that do not require convexity assumption. Lemma 6 bounds the sum of expected values of $\|\nabla F(w_t)\|^2$; and Lemma 7 expands the value of $\mathbb{E}[\|\nabla F(w_t) - v_t\|^2]$.

Lemma 6. *Suppose that Assumption 1 holds. Consider iSARAH-IN (Algorithm 2). Then, we have*

$$\begin{aligned} \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t)\|^2] &\leq \frac{2}{\eta} \mathbb{E}[F(w_0) - F(w_*)] \\ &\quad + \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] - (1 - L\eta) \sum_{t=0}^m \mathbb{E}[\|v_t\|^2], \end{aligned} \quad (21)$$

where $w_* = \arg \min_w F(w)$.

Proof. By Assumption 1 and $w_{t+1} = w_t - \eta v_t$, we have

$$\begin{aligned} \mathbb{E}[F(w_{t+1})] &\stackrel{(11)}{\leq} \mathbb{E}[F(w_t)] - \eta \mathbb{E}[\nabla F(w_t)^T v_t] + \frac{L\eta^2}{2} \mathbb{E}[\|v_t\|^2] \\ &= \mathbb{E}[F(w_t)] - \frac{\eta}{2} \mathbb{E}[\|\nabla F(w_t)\|^2] + \frac{\eta}{2} \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] \\ &\quad - \left(\frac{\eta}{2} - \frac{L\eta^2}{2} \right) \mathbb{E}[\|v_t\|^2], \end{aligned}$$

where the last equality follows from the fact $a_1^T a_2 = \frac{1}{2} [\|a_1\|^2 + \|a_2\|^2 - \|a_1 - a_2\|^2]$.

By summing over $t = 0, \dots, m$, we have

$$\begin{aligned} \mathbb{E}[F(w_{m+1})] &\leq \mathbb{E}[F(w_0)] - \frac{\eta}{2} \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t)\|^2] + \frac{\eta}{2} \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] \\ &\quad - \left(\frac{\eta}{2} - \frac{L\eta^2}{2} \right) \sum_{t=0}^m \mathbb{E}[\|v_t\|^2], \end{aligned}$$

which is equivalent to ($\eta > 0$):

$$\begin{aligned} \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t)\|^2] &\leq \frac{2}{\eta} \mathbb{E}[F(w_0) - F(w_m)] \\ &\quad + \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] - (1 - L\eta) \sum_{t=0}^m \mathbb{E}[\|v_t\|^2] \end{aligned}$$

$$\leq \frac{2}{\eta} \mathbb{E}[F(w_0) - F(w_*)] + \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] - (1 - L\eta) \sum_{t=0}^m \mathbb{E}[\|v_t\|^2],$$

where the second inequality follows since $w_* = \arg \min_w F(w)$. \square

Lemma 7. *Suppose that Assumption 1 holds. Consider v_t defined by (4) in iSARAH-IN (Algorithm 2). Then for any $t \geq 1$,*

$$\begin{aligned} \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] &= \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] \\ &\quad + \sum_{j=1}^t \mathbb{E}[\|v_j - v_{j-1}\|^2] - \sum_{j=1}^t \mathbb{E}[\|\nabla F(w_j) - \nabla F(w_{j-1})\|^2]. \end{aligned}$$

Proof. Let $\mathcal{F}_j = \sigma(w_0, w_1, \dots, w_j)$ be the σ -algebra generated by w_0, w_1, \dots, w_j . We note that ξ_j is independent of \mathcal{F}_j . For $j \geq 1$, we have

$$\begin{aligned} &\mathbb{E}[\|\nabla F(w_j) - v_j\|^2 | \mathcal{F}_j] \\ &= \mathbb{E}[\|[\nabla F(w_{j-1}) - v_{j-1}] + [\nabla F(w_j) - \nabla F(w_{j-1})] - [v_j - v_{j-1}]\|^2 | \mathcal{F}_j] \\ &= \|\nabla F(w_{j-1}) - v_{j-1}\|^2 + \|\nabla F(w_j) - \nabla F(w_{j-1})\|^2 + \mathbb{E}[\|v_j - v_{j-1}\|^2 | \mathcal{F}_j] \\ &\quad + 2(\nabla F(w_{j-1}) - v_{j-1})^T (\nabla F(w_j) - \nabla F(w_{j-1})) \\ &\quad - 2(\nabla F(w_{j-1}) - v_{j-1})^T \mathbb{E}[v_j - v_{j-1} | \mathcal{F}_j] \\ &\quad - 2(\nabla F(w_j) - \nabla F(w_{j-1}))^T \mathbb{E}[v_j - v_{j-1} | \mathcal{F}_j] \\ &= \|\nabla F(w_{j-1}) - v_{j-1}\|^2 - \|\nabla F(w_j) - \nabla F(w_{j-1})\|^2 + \mathbb{E}[\|v_j - v_{j-1}\|^2 | \mathcal{F}_j], \end{aligned}$$

where the last equality follows from

$$\mathbb{E}[v_j - v_{j-1} | \mathcal{F}_j] \stackrel{(4)}{=} \mathbb{E}[\nabla f(w_j; \xi_j) - \nabla f(w_{j-1}; \xi_j) | \mathcal{F}_j] = \nabla F(w_j) - \nabla F(w_{j-1}).$$

By taking expectation for the above equation, we have

$$\begin{aligned} \mathbb{E}[\|\nabla F(w_j) - v_j\|^2] &= \mathbb{E}[\|\nabla F(w_{j-1}) - v_{j-1}\|^2] \\ &\quad - \mathbb{E}[\|\nabla F(w_j) - \nabla F(w_{j-1})\|^2] + \mathbb{E}[\|v_j - v_{j-1}\|^2]. \end{aligned}$$

By summing over $j = 1, \dots, t$ ($t \geq 1$), we have

$$\begin{aligned} \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] &= \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] \\ &\quad + \sum_{j=1}^t \mathbb{E}[\|v_j - v_{j-1}\|^2] - \sum_{j=1}^t \mathbb{E}[\|\nabla F(w_j) - \nabla F(w_{j-1})\|^2]. \end{aligned}$$

\square

² \mathcal{F}_j contains all the information of w_0, \dots, w_j as well as v_0, \dots, v_{j-1}

3.4.1. Non-Strongly Convex Case

In this subsection, we analyze one-loop results of Inexact SARAH (Algorithm 2) in the non-strongly convex case. We first derive the bound for $\mathbb{E}[\|\nabla F(w_t) - v_t\|^2]$.

Lemma 8. *Suppose that Assumptions 1 and 3 hold. Consider v_t defined as (4) in SARAH (Algorithm 1) with $\eta < 2/L$. Then we have that for any $t \geq 1$,*

$$\mathbb{E}[\|\nabla F(w_t) - v_t\|^2] \leq \frac{\eta L}{2 - \eta L} \left[\mathbb{E}[\|v_0\|^2] - \mathbb{E}[\|v_t\|^2] \right] + \mathbb{E}[\|\nabla F(w_0) - v_0\|^2]. \quad (22)$$

Proof. For $j \geq 1$, we have

$$\begin{aligned} & \mathbb{E}[\|v_j\|^2 | \mathcal{F}_j] \\ &= \mathbb{E}[\|v_{j-1} - (\nabla f(w_{j-1}; \xi_j) - \nabla f(w_j; \xi_j))\|^2 | \mathcal{F}_j] \\ &= \|v_{j-1}\|^2 \\ & \quad + \mathbb{E}[\|\nabla f(w_{j-1}; \xi_j) - \nabla f(w_j; \xi_j)\|^2 - \frac{2}{\eta} (\nabla f(w_{j-1}; \xi_j) - \nabla f(w_j; \xi_j))^T (w_{j-1} - w_j) | \mathcal{F}_j] \\ & \stackrel{(13)}{\leq} \|v_{j-1}\|^2 + \mathbb{E}[\|\nabla f(w_{j-1}; \xi_j) - \nabla f(w_j; \xi_j)\|^2 - \frac{2}{L\eta} \|\nabla f(w_{j-1}; \xi_j) - \nabla f(w_j; \xi_j)\|^2 | \mathcal{F}_j] \\ &= \|v_{j-1}\|^2 + \left(1 - \frac{2}{\eta L}\right) \mathbb{E}[\|\nabla f(w_{j-1}; \xi_j) - \nabla f(w_j; \xi_j)\|^2 | \mathcal{F}_j] \\ & \stackrel{(4)}{=} \|v_{j-1}\|^2 + \left(1 - \frac{2}{\eta L}\right) \mathbb{E}[\|v_j - v_{j-1}\|^2 | \mathcal{F}_j], \end{aligned}$$

which, if we take expectation, implies that

$$\mathbb{E}[\|v_j - v_{j-1}\|^2] \leq \frac{\eta L}{2 - \eta L} \left[\mathbb{E}[\|v_{j-1}\|^2] - \mathbb{E}[\|v_j\|^2] \right],$$

when $\eta < 2/L$.

By summing the above inequality over $j = 1, \dots, t$ ($t \geq 1$), we have

$$\sum_{j=1}^t \mathbb{E}[\|v_j - v_{j-1}\|^2] \leq \frac{\eta L}{2 - \eta L} \left[\mathbb{E}[\|v_0\|^2] - \mathbb{E}[\|v_t\|^2] \right]. \quad (23)$$

By Lemma 7, we have

$$\begin{aligned} \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] &\leq \sum_{j=1}^t \mathbb{E}[\|v_j - v_{j-1}\|^2] + \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] \\ &\stackrel{(23)}{\leq} \frac{\eta L}{2 - \eta L} \left[\mathbb{E}[\|v_0\|^2] - \mathbb{E}[\|v_t\|^2] \right] + \mathbb{E}[\|\nabla F(w_0) - v_0\|^2]. \quad \square \end{aligned}$$

Lemma 9. *Suppose that Assumptions 1 and 3 hold. Consider v_0 defined as (3) in*

iSARAH (Algorithm 1). Then we have,

$$\begin{aligned}
& \frac{\eta L}{2 - \eta L} \mathbb{E}[\|v_0\|^2] + \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] \\
& \leq \frac{2}{2 - \eta L} \left(\frac{4L\mathbb{E}[F(w_0) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] - \mathbb{E}[\|\nabla F(w_0)\|^2]}{b} \right) \\
& \quad + \frac{\eta L}{2 - \eta L} \mathbb{E}[\|\nabla F(w_0)\|^2]. \tag{24}
\end{aligned}$$

Proof. By Corollary 1, we have

$$\begin{aligned}
& \frac{\eta L}{2 - \eta L} \mathbb{E}[\|v_0\|^2 | w_0] - \frac{\eta L}{2 - \eta L} \|\nabla F(w_0)\|^2 + \mathbb{E}[\|\nabla F(w_0) - v_0\|^2 | w_0] \\
& = \frac{2}{2 - \eta L} \left[\mathbb{E}[\|v_0\|^2 | w_0] - \|\nabla F(w_0)\|^2 \right] \\
& = \frac{2}{2 - \eta L} \left[\mathbb{E}[\|v_0 - \nabla F(w_0)\|^2 | w_0] \right] \\
& \stackrel{(18)}{\leq} \frac{2}{2 - \eta L} \left(\frac{4L[F(w_0) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] - \|\nabla F(w_0)\|^2}{b} \right).
\end{aligned}$$

Taking the expectation and adding $\frac{\eta L}{2 - \eta L} \mathbb{E}[\|\nabla F(w_0)\|^2]$ for both sides, the desired result is achieved. \square

We then derive this basic result for the convex case by using Lemmas 8 and 9.

Lemma 10. Suppose that Assumptions 1 and 3 hold. Consider *iSARAH-IN* (Algorithm 2) with $\eta \leq 1/L$. Then, we have

$$\begin{aligned}
\mathbb{E}[\|\nabla F(\tilde{w})\|^2] & \leq \frac{2}{\eta(m+1)} \mathbb{E}[F(w_0) - F(w_*)] + \frac{\eta L}{2 - \eta L} \mathbb{E}[\|\nabla F(w_0)\|^2] \\
& \quad + \frac{2}{2 - \eta L} \left(\frac{4L\mathbb{E}[F(w_0) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] - \mathbb{E}[\|\nabla F(w_0)\|^2]}{b} \right), \tag{25}
\end{aligned}$$

where w_* is any optimal solution of $F(w)$; and ξ is the random variable.

Proof. By Lemma 8, we have

$$\sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] \leq \frac{m\eta L}{2 - \eta L} \mathbb{E}[\|v_0\|^2] + (m+1) \mathbb{E}[\|\nabla F(w_0) - v_0\|^2]. \tag{26}$$

Hence, by Lemma 6 with $\eta \leq 1/L$, we have

$$\begin{aligned}
\sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t)\|^2] & \leq \frac{2}{\eta} \mathbb{E}[F(w_0) - F(w_*)] + \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] \\
& \stackrel{(26)}{\leq} \frac{2}{\eta} \mathbb{E}[F(w_0) - F(w_*)] + \frac{m\eta L}{2 - \eta L} \mathbb{E}[\|v_0\|^2]
\end{aligned}$$

$$+ (m+1)\mathbb{E}[\|\nabla F(w_0) - v_0\|^2]. \quad (27)$$

Since $\tilde{w} = w_t$, where t is picked uniformly at random from $\{0, 1, \dots, m\}$. The following holds,

$$\begin{aligned} \mathbb{E}[\|\nabla F(\tilde{w})\|^2] &= \frac{1}{m+1} \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t)\|^2] \\ &\stackrel{(27)}{\leq} \frac{2}{\eta(m+1)} \mathbb{E}[F(w_0) - F(w_*)] + \frac{\eta L}{2 - \eta L} \mathbb{E}[\|v_0\|^2] + \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] \\ &\stackrel{(24)}{\leq} \frac{2}{\eta(m+1)} \mathbb{E}[F(w_0) - F(w_*)] + \frac{\eta L}{2 - \eta L} \mathbb{E}[\|\nabla F(w_0)\|^2] \\ &\quad + \frac{2}{2 - \eta L} \left(\frac{4L\mathbb{E}[F(w_0) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] - \mathbb{E}[\|\nabla F(w_0)\|^2]}{b} \right). \quad \square \end{aligned}$$

This expected bound for $\|\nabla F(\tilde{w})\|^2$ will be used for deriving both one-loop and multiple-loop results in the convex case.

Lemma 10 can be used to get the following result for non-strongly convex.

Theorem 1. *Suppose that Assumptions 1, 3, and 4 hold. Consider iSARAH-IN (Algorithm 2) with $\eta = \frac{1}{L\sqrt{m+1}} \leq \frac{1}{L}$, $b = 2\sqrt{m+1}$ and a given w_0 . Then we have,*

$$\mathbb{E}[\|\nabla F(\tilde{w})\|^2] \leq \frac{1}{\sqrt{m+1}} \left[6L[F(w_0) - F(w_*)] + 2\sigma_*^2 \right].$$

Proof. By Lemma 10, for any given w_0 , we have

$$\begin{aligned} \mathbb{E}[\|\nabla F(\tilde{w})\|^2] &\leq \frac{2}{\eta(m+1)} [F(w_0) - F(w_*)] + \frac{\eta L}{2 - \eta L} \|\nabla F(w_0)\|^2 \\ &\quad + \frac{2}{2 - \eta L} \left(\frac{4L[F(w_0) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] - \|\nabla F(w_0)\|^2}{b} \right) \\ &\leq \frac{2L}{\sqrt{m+1}} [F(w_0) - F(w_*)] + \frac{1}{2 - \eta L} \frac{4L}{\sqrt{m+1}} [F(w_0) - F(w_*)] \\ &\quad + \frac{2}{2 - \eta L} \frac{\mathbb{E}[\|\nabla f(w_*; \xi)\|^2]}{\sqrt{m+1}} \\ &\leq \frac{2L}{\sqrt{m+1}} [F(w_0) - F(w_*)] \\ &\quad + \frac{1}{\sqrt{m+1}} \left[4L[F(w_0) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] \right] \\ &\stackrel{(10)}{\leq} \frac{1}{\sqrt{m+1}} \left[6L[F(w_0) - F(w_*)] + 2\sigma_*^2 \right]. \end{aligned}$$

The second inequality follows since $\eta = \frac{1}{L\sqrt{m+1}}$ and $b = 2\sqrt{m+1}$. The third inequality follows since $\eta \leq \frac{1}{L}$, which implies $\frac{1}{2 - \eta L} \leq 1$. \square

Based on Theorem 1, we are able to derive the following total complexity for iSARAH-IN in the non-strongly convex case.

Corollary 2. Suppose that Assumptions 1, 3, and 4 hold. Consider *iSARAH-IN* (Algorithm 2) with the learning rate $\eta = \frac{1}{L\sqrt{m+1}}$ and the number of samples $b = 2\sqrt{m+1}$, where m is the total number of iterations, then $\|\nabla F(\tilde{w})\|^2$ converges sublinearly in expectation with a rate of $\mathcal{O}\left(\frac{\max\{L, \sigma_*^2\}}{\sqrt{m+1}}\right)$, and therefore, the total complexity to achieve an ϵ -accurate solution is

$$\mathcal{O}\left(\frac{\max\{L, \sigma_*^2\}}{\epsilon} + \frac{\max\{L^2, \sigma_*^4\}}{\epsilon^2}\right).$$

Proof. It is easy to see that to achieve $\mathbb{E}[\|\nabla F(\tilde{w})\|^2] \leq \epsilon$ we need

$$m + 1 = \frac{(6L[F(w_0) - F(w_*)] + 2\sigma_*^2)^2}{\epsilon^2},$$

and hence the total work is

$$b + 2m = 2\sqrt{m+1} + 2m = \mathcal{O}\left(\frac{\max\{L, \sigma_*^2\}}{\epsilon} + \frac{\max\{L^2, \sigma_*^4\}}{\epsilon^2}\right). \quad \square$$

3.4.2. Non-Convex Case

We now move to the non-convex case. We begin by stating and proving a lemma similar to Lemma 8, bounding $\mathbb{E}[\|\nabla F(w_t) - v_t\|^2]$, but without Assumption 3.

Lemma 11. Suppose that Assumption 1 holds. Consider v_t defined as (4) in *iSARAH-IN* (Algorithm 2). Then for any $t \geq 1$,

$$\mathbb{E}[\|\nabla F(w_t) - v_t\|^2] \leq \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] + L^2\eta^2 \sum_{j=1}^t \mathbb{E}[\|v_{j-1}\|^2]. \quad (28)$$

Proof. We have, for $t \geq 1$,

$$\|v_t - v_{t-1}\|^2 \stackrel{(4)}{=} \|\nabla f(w_t; \xi_t) - \nabla f(w_{t-1}; \xi_t)\|^2 \stackrel{(8)}{\leq} L^2\|w_t - w_{t-1}\|^2 = L^2\eta^2\|v_{t-1}\|^2. \quad (29)$$

Hence, by Lemma 7,

$$\begin{aligned} \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] &\leq \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] + \sum_{j=1}^t \mathbb{E}[\|v_j - v_{j-1}\|^2] \\ &\stackrel{(29)}{\leq} \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] + L^2\eta^2 \sum_{j=1}^t \mathbb{E}[\|v_{j-1}\|^2]. \quad \square \end{aligned}$$

Lemma 12. Suppose that Assumption 1 holds. Consider v_t defined as (4) in *iSARAH-IN* (Algorithm 2) with $\eta \leq \frac{2}{L(\sqrt{1+4m+1})}$. Then we have

$$L^2\eta^2 \sum_{t=0}^m \sum_{j=1}^t \mathbb{E}[\|v_{j-1}\|^2] - (1 - L\eta) \sum_{t=0}^m \mathbb{E}[\|v_t\|^2] \leq 0. \quad (30)$$

Proof. For $\eta \leq \frac{2}{L(\sqrt{1+4m+1})}$, we have

$$\begin{aligned}
L^2\eta^2 \sum_{t=0}^m \sum_{j=1}^t \mathbb{E}[\|v_{j-1}\|^2] - (1-L\eta) \sum_{t=0}^m \mathbb{E}[\|v_t\|^2] \\
&= L^2\eta^2 \left[m\mathbb{E}\|v_0\|^2 + (m-1)\mathbb{E}\|v_1\|^2 + \cdots + \mathbb{E}\|v_{m-1}\|^2 \right] \\
&\quad - (1-L\eta) \left[\mathbb{E}\|v_0\|^2 + \mathbb{E}\|v_1\|^2 + \cdots + \mathbb{E}\|v_m\|^2 \right] \\
&\leq [L^2\eta^2 m - (1-L\eta)] \sum_{t=1}^m \mathbb{E}[\|v_{t-1}\|^2] \leq 0,
\end{aligned}$$

since $\eta = \frac{2}{L(\sqrt{1+4m+1})}$ is a root of the equation $L^2\eta^2 m - (1-L\eta) = 0$. \square

With the help of the above lemmas, we are able to derive our result for non-convex.

Theorem 2. Suppose that Assumption 1 holds and $\mathbb{E}[\|\nabla f(w_0; \xi)\|^2]$ is finite. Consider iSARAH-IN (Algorithm 2) with $\eta \leq \frac{2}{L(\sqrt{1+4m+1})} \leq \frac{1}{L}$, $b = \sqrt{m+1}$ and a given w_0 . Then we have,

$$\mathbb{E}[\|\nabla F(\tilde{w})\|^2] \leq \frac{2}{\eta(m+1)} [F(w_0) - F^*] + \frac{1}{\sqrt{m+1}} \left(\mathbb{E}[\|\nabla f(w_0; \xi)\|^2] \right), \quad (31)$$

where F^* is any lower bound of F ; and ξ is some random variable.

Proof. Let F^* be any lower bound of F . By Lemma 6 and since $\tilde{w} = w_t$, where t is picked uniformly at random from $\{0, 1, \dots, m\}$, we have

$$\begin{aligned}
\mathbb{E}[\|\nabla F(\tilde{w})\|^2] &= \frac{1}{m+1} \sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t)\|^2] \\
&\leq \frac{2}{\eta(m+1)} \mathbb{E}[F(w_0) - F^*] \\
&\quad + \frac{1}{m+1} \left(\sum_{t=0}^m \mathbb{E}[\|\nabla F(w_t) - v_t\|^2] - (1-L\eta) \sum_{t=0}^m \mathbb{E}[\|v_t\|^2] \right) \\
&\stackrel{(28)}{\leq} \frac{2}{\eta(m+1)} \mathbb{E}[F(w_0) - F^*] + \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] \\
&\quad + \frac{1}{m+1} \left(L^2\eta^2 \sum_{t=0}^m \sum_{j=1}^t \mathbb{E}[\|v_{j-1}\|^2] - (1-L\eta) \sum_{t=0}^m \mathbb{E}[\|v_t\|^2] \right) \\
&\stackrel{(30)}{\leq} \frac{2}{\eta(m+1)} \mathbb{E}[F(w_0) - F^*] + \mathbb{E}[\|\nabla F(w_0) - v_0\|^2] \\
&\stackrel{(17)}{\leq} \frac{2}{\eta(m+1)} \mathbb{E}[F(w_0) - F^*] + \frac{1}{b} \mathbb{E}[\|\nabla f(w_0; \xi)\|^2].
\end{aligned}$$

For any given w_0 and $b = \sqrt{m+1}$, we could achieve the desired result. \square

Based on Theorem 2, we are able to derive the following total complexity for iSARAH-IN in the non-convex case.

Corollary 3. *Suppose that Assumption 1 holds and $\mathbb{E}[\|\nabla f(w_0; \xi)\|^2]$ is finite. Consider iSARAH-IN (Algorithm 2) with the learning rate $\eta = \mathcal{O}\left(\frac{1}{L\sqrt{m+1}}\right)$ and the number of samples $b = \sqrt{m+1}$, where m is the total number of iterations, then $\|\nabla F(\tilde{w})\|^2$ converges sublinearly in expectation with a rate of $\mathcal{O}\left(\sqrt{\frac{1}{m+1}}\right)$, and therefore, the total complexity to achieve an ϵ -accurate solution is $\mathcal{O}(1/\epsilon^2)$.*

Proof. Same as non-strongly convex case, since $\mathbb{E}[\|\nabla f(w_0; \xi)\|^2]$ is finite, to achieve $\mathbb{E}[\|\nabla F(\tilde{w})\|^2] \leq \epsilon$ we need $m = \mathcal{O}(1/\epsilon^2)$ and hence the total work is $\sqrt{m} + 2m = \mathcal{O}\left(\frac{1}{\epsilon} + \frac{1}{\epsilon^2}\right) = \mathcal{O}\left(\frac{1}{\epsilon^2}\right)$. \square

3.5. Multiple-loop iSARAH Results

In this section, we analyze multiple-loop results of Inexact SARAH (Algorithm 1).

3.5.1. Strongly Convex Case

We now turn to the discussion on the convergence of iSARAH under the strong convexity assumption on F .

Theorem 3. *Suppose that Assumptions 1, 2, 3, and 4 hold. Consider iSARAH (Algorithm 1) with the choice of η , m , and b such that*

$$\alpha = \frac{1}{\mu\eta(m+1)} + \frac{\eta L}{2 - \eta L} + \frac{4\kappa - 2}{b(2 - \eta L)} < 1.$$

(Note that $\kappa = L/\mu$.) Then, we have

$$\mathbb{E}[\|\nabla F(\tilde{w}_s)\|^2] - \Delta \leq \alpha^s (\|\nabla F(\tilde{w}_0)\|^2 - \Delta), \quad (32)$$

where

$$\Delta = \frac{\delta}{1 - \alpha} \text{ and } \delta = \frac{4}{b(2 - \eta L)} \sigma_*^2.$$

Proof. By Lemma 10, with $\tilde{w} = \tilde{w}_s$ and $w_0 = \tilde{w}_{s-1}$, we have

$$\begin{aligned} & \mathbb{E}[\|\nabla F(\tilde{w}_s)\|^2] \\ & \leq \frac{2}{\eta(m+1)} \mathbb{E}[F(\tilde{w}_{s-1}) - F(w_*)] + \frac{\eta L}{2 - \eta L} \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] \\ & \quad + \frac{2}{2 - \eta L} \left(\frac{4L\mathbb{E}[F(\tilde{w}_{s-1}) - F(w_*)] + 2\mathbb{E}[\|\nabla f(w_*; \xi)\|^2] - \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2]}{b} \right) \\ & \stackrel{(9)}{\leq} \left(\frac{1}{\mu\eta(m+1)} + \frac{\eta L}{2 - \eta L} + \frac{4\kappa - 2}{b(2 - \eta L)} \right) \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] \\ & \quad + \frac{4}{b(2 - \eta L)} \mathbb{E}[\|\nabla f(w_*; \xi)\|^2] \\ & \stackrel{(10)}{\leq} \left(\frac{1}{\mu\eta(m+1)} + \frac{\eta L}{2 - \eta L} + \frac{4\kappa - 2}{b(2 - \eta L)} \right) \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] \end{aligned}$$

$$\begin{aligned}
& + \frac{4}{b(2-\eta L)}\sigma_*^2 \\
& = \alpha\mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] + \delta \\
& \leq \alpha^s\|\nabla F(\tilde{w}_0)\|^2 + \alpha^{s-1}\delta + \dots + \alpha\delta + \delta \\
& \leq \alpha^s\|\nabla F(\tilde{w}_0)\|^2 + \delta\frac{1-\alpha^s}{1-\alpha} \\
& = \alpha^s\|\nabla F(\tilde{w}_0)\|^2 + \Delta(1-\alpha^s) \\
& = \alpha^s(\|\nabla F(\tilde{w}_0)\|^2 - \Delta) + \Delta.
\end{aligned} \tag{33}$$

By adding $-\Delta$ to both sides, we achieve the desired result. \square

Based on Theorem 3, we are able to derive the following total complexity for iSARAH in the strongly convex case.

Corollary 4. *Let $\eta = \frac{2}{5L}$, $m = 20\kappa - 1$, and $b = \max\left\{20\kappa - 10, \frac{20\sigma_*^2}{\epsilon}\right\}$ in Theorem 3. Then, the total work complexity to achieve $\mathbb{E}[\|\nabla F(\tilde{w}_s)\|^2] \leq \epsilon$ is $\mathcal{O}\left(\max\left\{\frac{\sigma_*^2}{\epsilon}, \kappa\right\} \log\left(\frac{1}{\epsilon}\right)\right)$.*

Proof. With $\eta = \frac{2}{5L}$, $m = 20\kappa - 1$, and $b = \max\left\{20\kappa - 10, \frac{20\sigma_*^2}{\epsilon}\right\}$, from (33), we have

$$\begin{aligned}
\mathbb{E}[\|\nabla F(\tilde{w}_s)\|^2] & \leq \left(\frac{1}{8} + \frac{1}{4} + \frac{1}{8}\right)\mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] + \frac{\epsilon}{8} \\
& \leq \frac{1}{2}\mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] + \frac{\epsilon}{8} \\
& \leq \frac{1}{2^s}\|\nabla F(\tilde{w}_0)\|^2 + \frac{\epsilon}{4}.
\end{aligned}$$

Since $\mathbb{E}[\|\nabla f(w_*; \xi)\|^2]$ is finite, to guarantee that $\mathbb{E}[\|\nabla F(\tilde{w}_s)\|^2] \leq \epsilon$, it is sufficient to make $\frac{1}{2^s}\|\nabla F(\tilde{w}_0)\|^2 = \frac{3}{4}\epsilon$ or equivalently $s = \log\left(\frac{\|\nabla F(\tilde{w}_0)\|^2}{\frac{3}{4}\epsilon}\right)$. This implies the total complexity to achieve an ϵ -accuracy solution is $(b+m)s = \mathcal{O}\left(\left(\max\left\{\frac{\sigma_*^2}{\epsilon}, \kappa\right\} + \kappa\right) \log\left(\frac{1}{\epsilon}\right)\right) = \mathcal{O}\left(\max\left\{\frac{\sigma_*^2}{\epsilon}, \kappa\right\} \log\left(\frac{1}{\epsilon}\right)\right)$. \square

3.5.2. Non-Strongly Convex Case

We turn to the analysis of the convergence rate of the multiple-loop iSARAH in the non-strongly convex case. As mentioned in the introduction, we are able to achieve best sample complexity rates of any stochastic algorithm, to the best of our knowledge, but under an additional reasonably mild assumption. We introduce this assumption below.

Assumption 5. *Let $\tilde{w}_0, \dots, \tilde{w}_{\mathcal{T}}$ be the (outer) iterations of Algorithm 1. We assume that there exist $M > 0$ and $N > 0$ such that, $k = 0, \dots, \mathcal{T}$,*

$$F(\tilde{w}_k) - F(w_*) \leq M\|\nabla F(\tilde{w}_k)\|^2 + N, \tag{34}$$

where $F(w_*)$ is the optimal value of F .

Let us discuss Assumption 5. First, we note that the assumption only requires to hold for the outer iterations $\tilde{w}_0, \dots, \tilde{w}_{\mathcal{T}}$ of Algorithm 1 instead of holding for all $w \in \mathbb{R}^d$

or for all of the inner iterates. Moreover, this assumption is clearly weaker than the Polyak-Lojasiewicz (PL) condition, which has been studied and discussed in [4, 10, 17] which itself is weaker than strong convexity assumption. Under PL condition, we simply have $N = 0$ in (34) and as we will discuss below we can recover better convergence rate in this case. On the other hand, if PL condition does not hold but if the sequence of iterates $\{\tilde{w}_k\}$ remains in a set, say \mathcal{W} , on which the objective function is bounded from above, that is for all $w \in \mathcal{W}$, $F(w) \leq F_{\max}$ for some finite value F_{\max} , then Assumption 5 is satisfied with $N = F_{\max} - F(w_*)$ and $M = 0$, where $F(w_*)$ is the optimal value of F . In other words, Assumption 5 is a relaxation of the boundedness assumption and the PL condition. As an example, consider the following modification of the logistic function (for some $\lambda > 0$)

$$F(w) = \begin{cases} \log(1 + e^{-w}) & w \geq -2 \\ \log(1 + e^{-w}) + \frac{\lambda}{2}(w + 2)^2 & w < -2 \end{cases}$$

which is not strongly convex and does not satisfy the PL condition. This function can be considered in practice instead of the usual logistic function as it simply adds somewhat larger penalty on solutions that are far away from the optimal. Notice that this function is continuous and continuously differentiable. When, $w < -2$, the function is strongly convex and satisfies the PL condition. When $w \geq -2$, the function is bounded in $(0, \log(1 + e^2)]$. Therefore, it satisfies Assumption 5. This is a simplified example to make derivations easy. However, it is possible to generalize this example by modifying standard logistic loss by adding strongly convex penalty outside some ball containing the optimal solution.

Theorem 4. *Suppose that Assumptions 1, 3, 4, and 5 hold. Consider iSARAH (Algorithm 1) with the choice of η , m , and b such that*

$$\alpha_c = \frac{2M}{\eta(m+1)} + \frac{\eta L}{2 - \eta L} + \frac{8LM - 1}{b(2 - \eta L)} < 1.$$

Then, we have

$$\mathbb{E}[\|\nabla F(\tilde{w}_s)\|^2] - \Delta_c \leq \alpha_c^s (\|\nabla F(\tilde{w}_0)\|^2 - \Delta_c), \quad (35)$$

where

$$\Delta_c = \frac{\delta_c}{1 - \alpha_c} \text{ and } \delta_c = \frac{2N}{\eta(m+1)} + \frac{8LN}{b(2 - \eta L)} + \frac{4\sigma_*^2}{b(2 - \eta L)}.$$

Proof. By Lemma 10, with $\tilde{w} = \tilde{w}_s$ and $w_0 = \tilde{w}_{s-1}$, we have

$$\begin{aligned} & \mathbb{E}[\|\nabla F(\tilde{w}_s)\|^2] \\ & \leq \frac{2}{\eta(m+1)} \mathbb{E}[F(\tilde{w}_{s-1}) - F(w_*)] + \frac{\eta L}{2 - \eta L} \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] \\ & + \frac{2}{2 - \eta L} \left(\frac{4L \mathbb{E}[F(\tilde{w}_{s-1}) - F(w_*)] + 2 \mathbb{E}[\|\nabla f(w_*; \xi)\|^2] - \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2]}{b} \right) \\ & \stackrel{(34)}{\leq} \left(\frac{2M}{\eta(m+1)} + \frac{\eta L}{2 - \eta L} + \frac{8LM - 1}{b(2 - \eta L)} \right) \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] \end{aligned}$$

$$\begin{aligned}
& + \frac{2N}{\eta(m+1)} + \frac{8LN}{b(2-\eta L)} + \frac{4\mathbb{E}[\|\nabla f(w_*; \xi)\|^2]}{b(2-\eta L)} \\
& \stackrel{(10)}{\leq} \left(\frac{2M}{\eta(m+1)} + \frac{\eta L}{2-\eta L} + \frac{8LM-1}{b(2-\eta L)} \right) \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] \\
& + \frac{2N}{\eta(m+1)} + \frac{8LN}{b(2-\eta L)} + \frac{4\sigma_*^2}{b(2-\eta L)} \\
& = \alpha_c \mathbb{E}[\|\nabla F(\tilde{w}_{s-1})\|^2] + \delta_c \\
& \leq \alpha_c^s (\|\nabla F(\tilde{w}_0)\|^2 - \Delta_c) + \Delta_c. \quad \square
\end{aligned}$$

Now choosing the appropriate values for η , m , and b , we can follow the proof of Corollary 4 to achieve the following complexity result.

Corollary 5. *In Theorem 4, let $\eta = \frac{2}{5L}$, $m = \max\left\{40LM - 1, \frac{120LN}{\epsilon} - 1\right\}$, and $b = \max\left\{40LM - 5, \frac{120LN}{\epsilon}, \frac{60\sigma_*^2}{\epsilon}\right\}$. Then, the total work complexity to achieve $\mathbb{E}[\|\nabla F(\tilde{w}_s)\|^2] \leq \epsilon$ is*

$$(b+m)s = \mathcal{O}\left(\max\left\{LM, \frac{\max\{LN, \sigma_*^2\}}{\epsilon}\right\} \log\left(\frac{1}{\epsilon}\right)\right).$$

We can observe that, with the help of Assumption 5, iSARAH could achieve the best known complexity among stochastic methods (those which do not have access to exact gradient computation) in the non-strongly convex case.

4. Conclusion

We have provided the analysis of the inexact version of SARAH, which requires only stochastic gradient information computed on a mini-batch of sufficient size. We provide the one-loop results (iSARAH-IN) in the non-strongly convex and non-convex cases. Moreover, we analyze the multiple-loop results (iSARAH) in the strongly convex case and with an additional assumption (Assumption 5) in the non-strongly convex case. With this Assumption 5, which we argue is reasonable, iSARAH achieves the best known complexity among stochastic methods.

Acknowledgement

Katya Scheinberg and Martin Takáč were partially supported by the U.S. National Science Foundation, under award number CCF-1618717 and CCF-1740796.

References

- [1] A. Defazio, F. Bach, and S. Lacoste-Julien, *SAGA: A Fast Incremental Gradient Method With Support for Non-Strongly Convex Composite Objectives*, in *NIPS*. 2014, pp. 1646–1654.
- [2] C. Fang, C.J. Li, Z. Lin, and T. Zhang, *Spider: Near-optimal non-convex optimization via stochastic path integrated differential estimator*, arXiv preprint arXiv:1807.01695 (2018).
- [3] R. Johnson and T. Zhang, *Accelerating Stochastic Gradient Descent Using Predictive Variance Reduction*, in *NIPS*. 2013, pp. 315–323.

- [4] H. Karimi, J. Nutini, and M. Schmidt, *Linear Convergence of Gradient and Proximal-Gradient Methods Under the Polyak-Łojasiewicz Condition*, in *Machine Learning and Knowledge Discovery in Databases*, P. Frasconi, N. Landwehr, G. Manco, and J. Vreeken, eds., Cham. Springer International Publishing, 2016, pp. 795–811.
- [5] J. Konečný and P. Richtárik, *Semi-stochastic gradient descent methods*, arXiv:1312.1666 (2013).
- [6] N. Le Roux, M. Schmidt, and F. Bach, *A Stochastic Gradient Method with an Exponential Convergence Rate for Finite Training Sets*, in *NIPS*. 2012, pp. 2663–2671.
- [7] L. Lei and M.I. Jordan, *Less than a Single Pass: Stochastically Controlled Stochastic Gradient*, in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, A. Singh and J. Zhu, eds., Proceedings of Machine Learning Research Vol. 54, 20–22 Apr, Fort Lauderdale, FL, USA. PMLR, 2017, pp. 148–156.
- [8] L. Lei, C. Ju, J. Chen, and M.I. Jordan, *Non-convex finite-sum optimization via SCSG methods*, in *Advances in Neural Information Processing Systems 30*, I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds., Curran Associates, Inc., 2017, pp. 2348–2358. Available at <http://papers.nips.cc/paper/6829-non-convex-finite-sum-optimization-via-scsg-methods.pdf>.
- [9] Y. Nesterov, *Introductory lectures on convex optimization: a basic course*, Applied optimization, Kluwer Academic Publ., Boston, Dordrecht, London, 2004.
- [10] Y. Nesterov and B.T. Polyak, *Cubic regularization of Newton method and its global performance*, *Mathematical Programming* 108 (2006), pp. 177–205.
- [11] L. Nguyen, J. Liu, K. Scheinberg, and M. Takáč, *SARAH: A novel method for machine learning problems using stochastic recursive gradient*, *ICML (2017)*, pp. 2613–2621.
- [12] L. Nguyen, N. Nguyen, D. Phan, J. Kalagnanam, and K. Scheinberg, *When does stochastic gradient algorithm work well?*, arXiv:1801.06159 (2018).
- [13] L. Nguyen, P.H. Nguyen, M. van Dijk, P. Richtárik, K. Scheinberg, and M. Takáč, *SGD and Hogwild! convergence without the bounded gradients assumption*, arXiv:1802.03801 (2018).
- [14] L.M. Nguyen, J. Liu, K. Scheinberg, and M. Takáč, *Stochastic recursive gradient algorithm for nonconvex optimization*, *CoRR abs/1705.07261 (2017)*.
- [15] C. Paquette and K. Scheinberg, *A stochastic line search method with convergence rate analysis*, arXiv preprint arXiv:1807.07994 (2018).
- [16] N.H. Pham, L.M. Nguyen, D.T. Phan, and Q. Tran-Dinh, *ProxSARAH: An efficient algorithmic framework for stochastic composite nonconvex optimization*, arXiv preprint arXiv:1902.05679 (2019).
- [17] B.T. Polyak, *Some methods of speeding up the convergence of iteration methods*, *USSR Computational Mathematics and Mathematical Physics* 4 (1964), pp. 1–17.
- [18] M. Schmidt, N. Le Roux, and F. Bach, *Minimizing finite sums with the stochastic average gradient*, *Mathematical Programming* 162 (2017), pp. 83–112.
- [19] Z. Wang, K. Ji, Y. Zhou, Y. Liang, and V. Tarokh, *Spiderboost: A class of faster variance-reduced algorithms for nonconvex optimization*, arXiv preprint arXiv:1810.10690 (2018).