# Momentum Schemes with Stochastic Variance Reduction for Nonconvex Composite Optimization

**Yi Zhou** [1] **Zhe Wang** [2] **Kaiyi Ji** [2] **Yingbin Liang** [2] **Vahid Tarokh** [1]

## Abstract

Two new stochastic variance-reduced algorithms named SARAH and SPIDER have been recently proposed, and SPIDER has been shown to achieve a near-optimal gradient oracle complexity for nonconvex optimization. However, the theoretical advantage of SPIDER does not lead to substantial improvement of practical performance over SVRG. To address this issue, momentum technique can be a good candidate to improve the performance of SPIDER. However, existing momentum schemes used in variance-reduced algorithms are designed specifically for convex optimization, and are not applicable to nonconvex scenarios. In this paper, we develop novel momentum schemes with flexible coefficient settings to accelerate SPIDER for nonconvex and nonsmooth composite optimization, and show that the resulting algorithms achieve the near-optimal gradient oracle complexity for achieving a generalized first-order stationary condition. Furthermore, we generalize our algorithm to online nonconvex and nonsmooth optimization, and establish an oracle complexity result that matches the state-of-the-art. Our extensive experiments demonstrate the superior performance of our proposed algorithm over other stochastic variance-reduced algorithms.

## 1. Introduction

In the era of machine learning, optimization problems associated with practical applications have a rapidly increasing data volume. In many scenarios, such optimization problems take the following composite form:

$$\min_{x\in\mathbb{R}^d} F(x) := f(x) + g(x), \qquad \text{(P)}$$
$$\text{where } f(x) := \frac{1}{n}\sum_{i=1}^{n} \ell_i(x),$$

where $x \in \mathbb{R}^d$ is the optimization variable, integer $n$ denotes the total sample size, $\ell_i : \mathbb{R}^d \to \mathbb{R}$ is a differentiable function that corresponds to the loss on the $i$-th data sample and $g : \mathbb{R}^d \to \mathbb{R}$ denotes a possibly nonsmooth regularizer function. In particular, solving the above problem (P) can be demanding due to the tremendous data size $n$ and complex machine learning models (e.g., neural networks) that result in highly nonconvex and nonsmooth loss landscape (Goodfellow et al., 2016). Therefore, stochastic gradient-like algorithms are commonly used in practice to leverage their sample efficiency and implementation simplicity while maintaining provable convergence guarantee in nonconvex optimization.

A variety of stochastic algorithms have been proposed in the literature for solving the problem (P) without the existence of the regularizer $g$ (i.e., smooth nonconvex optimization). The simplest algorithm is the stochastic gradient descent (SGD) algorithm (Robbins & Monro, 1951; Bottou, 2010) that approximates the full gradient by one mini-batch of stochastic samples. Although SGD has a low per-iteration complexity, its convergence rate can be significantly deteriorated by the intrinsic variance of its stochastic estimator. Such an issue has been successfully resolved by using more advanced stochastic variance-reduced gradient estimators that induce a smaller variance, leading to the design of a variety of stochastic variance-reduced algorithms such as SAG (Schmidt et al., 2017), SAGA (Defazio et al., 2014), SVRG (Johnson & Zhang, 2013), etc. To further handle the nonsmooth regularizer $g$, proximal versions of these advanced algorithms have been developed (Xiao & Zhang, 2014; Ghadimi et al., 2016; Reddi et al., 2016). However, these algorithms do not yield an optimal stochastic gradient oracle complexity for generic nonconvex optimization.

Recently, (Nguyen et al., 2017a;b) and (Fang et al., 2018) proposed a new type of stochastic variance-reduced algorithms called SARAH and SPIDER, respectively. In specific,

[1]Department of ECE, Duke University. [2]Department of ECE, The Ohio State University. Correspondence to: Yi Zhou <yi.zhou610@duke.edu>.

under an accuracy-dependent stepsize, it has been shown in (Fang et al., 2018) that a natural gradient descent step taken in SPIDER yields the optimal[1] stochastic gradient oracle complexity for solving the problem (P) without the regularizer $g$. In a subsequent work (Wang et al., 2018a), the authors further proposed an improved algorithm scheme called Proximal SpiderBoost that allows to use a much larger constant-level stepsize and achieves the same order-level stochastic gradient oracle complexity for solving the problem (P) under a convex regularizer $g$.

Although the aforementioned SPIDER-based algorithms achieve the optimal stochastic gradient oracle complexity in nonconvex optimization, their practical performance has been found in recent works (Nguyen et al., 2017b; Fang et al., 2018) to be hardly advantageous to that of the traditional SVRG. Therefore, it is of vital importance to exploit the structure of the SPIDER estimator in other algorithmic dimensions to further improve the practical performance of SPIDER-based algorithms. Momentum is such a promising and important perspective. In fact, there are still two major challenges ahead to design momentum schemes for variance-reduced algorithms in nonconvex optimization. First, while momentum scheme has been well studied for (stochastic) gradient algorithms (Ghadimi & Lan, 2016) in nonconvex optimization, the convergence guarantee of stochastic variance-reduced-like algorithms is only explored for SVRG in certain convex scenarios (Nitanda, 2016; Allen-Zhu, 2017; 2018; Shang et al., 2018). Therefore, it is not clear whether a certain momentum scheme can be applied to stochastic variance-reduced algorithms based on SPIDER and yield the optimal oracle gradient complexity for nonconvex optimization. Furthermore, the existing momentum scheme for stochastic algorithms to solve the nonconvex problem (P) has convergence guarantee only for convex regularizers $g$ that have a bounded domain (Ghadimi & Lan, 2016), which are not applicable to a variety of application scenarios where regularizers with unbounded domain (e.g., $\ell_1, \ell_2$) are commonly used.

In this paper, we explore momentum schemes for SPIDER-based variance reduction algorithms that can solve the nonconvex and nonsmooth problem (P) under a much broader choice of regularizers with convergence guarantee. We summarize our contributions as follows.

### Summary of Contributions

We consider solving the problem (P) with nonconvex loss functions and an arbitrary convex regularizer (possibly nonsmooth). We propose Proximal SPIDER-M, which is a proximal stochastic algorithm that exploits both the SPIDER variance-reduction scheme and a momentum scheme for solving the problem (P). We show that the output point generated by the Proximal SPIDER-M satisfies a generalized $\epsilon$-first-order stationary condition within $O(\epsilon^{-2})$ number of iterations, and the corresponding stochastic gradient oracle complexity is in the order of $O(n + \sqrt{n}\epsilon^{-2})$, matching the complexity lower bound for nonconvex optimization. To the best of our knowledge, this is the first known theoretical guarantee for stochastic variance-reduced type of algorithms with momentum in nonconvex optimization. We also note that the design of our momentum scheme is applicable to arbitrary convex regularizers, which significantly relaxes the constraint of the existing momentum scheme that requires the regularizer to have a bounded domain in order to have a convergence guarantee for nonconvex optimization (Ghadimi & Lan, 2016).

We further propose two variants of the momentum scheme, i.e., epochwise diminishing momentum and epochwise restart momentum, for Proximal SPIDER-M. We establish the same order-level oracle complexity result in nonconvex optimization as mentioned above. To the best of our knowledge, this is the first formal theoretical guarantee for epochwise diminishing and restart momentum schemes in nonconvex optimization. Moreover, we generalize the Proximal SPIDER-M to solve the problem (P) in an online setting, and show that the algorithm satisfies the generalized $\epsilon$-first-order stationary condition within $O(\epsilon^{-2})$ number of iterations, and the associated stochastic gradient oracle complexity is in the order $O(\epsilon^{-3})$, matching the state-of-the-art result. Our numerical experiments demonstrate that the momentum scheme does substantially improve the practical performance of SPIDER and outperform other momentum-based variance-reduced algorithms.

### Related Work

**Stochastic algorithms for nonconvex optimization:** For nonconvex optimization, SGD has been shown to achieve an $\epsilon$-first-order stationary condition with an overall stochastic gradient oracle complexity of $O(\epsilon^{-4})$ (Ghadimi et al., 2016). Convergence guarantee for various stochastic variance-reduced algorithms have been established in nonconvex optimization. In specific, SAGA and SVRG have been shown to yield an overall stochastic gradient oracle complexity of $O(n^{2/3}\epsilon^{-2})$ (Reddi et al., 2016; Allen-Zhu & Hazan, 2016) to achieve an $\epsilon$-first-order stationary condition. More recently, (Nguyen et al., 2017a;b) proposed a novel stochastic variance reduction algorithm named SARAH and showed that the corresponding stochastic gradient oracle complexity is $O(\epsilon^{-4})$ to attain an $\epsilon$-first-order stationary point. The SPIDER algorithm (Fang et al., 2018) is a variant of SARAH that uses the same gradient estimator as SARAH but adopts a natural gradient descent update. (Fang et al., 2018) showed that SPIDER achieves an overall $\min\{O(n^{1/2}\epsilon^{-2}), O(\epsilon^{-3})\}$ stochastic gradient oracle com-

[1] In the parameter regime $n \leq \Theta(\epsilon^{-4})$, where $\epsilon$ corresponds to the desired accuracy.

plexity, which is optimal within the regime $n \leq O(\epsilon^{-4})$. (Wang et al., 2018a) further proposed an improved SPIDER scheme that allows to use a constant-level stepsize and can solve composite nonconvex optimization problems. In (Zhou et al., 2018), the authors proposed a nested stochastic variance reduction scheme for nonconvex optimization and achieve the same order-level oracle complexity result as that of SPIDER. More recently, (Zhou et al., 2019; Zhang et al., 2018) further applied the SARAH and SPIDER estimators to nonconvex optimization problems over manifolds.

**Momentum schemes for nonconvex optimization:** Momentum scheme is originally designed for accelerating gradient algorithms to achieve an optimal convergence rate in convex optimization (Nesterov, 2014; Beck & Teboulle, 2009; Tseng, 2010; Ghadimi & Lan, 2016). For nonconvex optimization, (Ghadimi & Lan, 2016) established convergence of stochastic gradient algorithms with momentum to an $\epsilon$-first-order stationary point with an overall stochastic gradient oracle complexity of $O(\epsilon^{-4})$. The convergence guarantee of SVRG with momentum has been explored under a certain local gradient dominance geometry in nonconvex optimization (Li et al., 2017). However, the momentum scheme there requires to compare the objective function value (and hence calculate the total loss) at each iteration and hence is not sample efficient. Similar momentum scheme has also been explored in second-order algorithms for nonconvex optimization (Wang et al., 2018b).

## 2. Preliminaries

In this section, we introduce some definitions and assumptions that are used throughout the paper. Recall that we are interested in solving the following optimization problem with composite objective function

$$\min_{x\in\mathbb{R}^d} F(x) := f(x) + g(x), \tag{P}$$
$$\text{where } f(x) := \frac{1}{n}\sum_{i=1}^{n} \ell_i(x),$$

where the function $f$ denotes the total loss on the training data and the function $g$ corresponds to the regularizer that penalizes the violation of a desired structure (e.g., sparsity, low-rankness, etc). We adopt the following standard assumptions on the problem (P).

**Assumption 1.** *The objective function in the problem (P) satisfies:*

1. *Function $F$ is bounded below, i.e.,*
$$F^* := \inf_{x\in\mathbb{R}^d} F(x) > -\infty; \tag{1}$$
2. *The loss functions $\ell_i, i = 1, \ldots, n$ are $L$-smooth, i.e., for all $i = 1, \ldots, n$, there exists an $L > 0$ such that*
$$\|\nabla \ell_i(x) - \nabla \ell_i(y)\| \leq L\|x - y\|, \forall x, y \in \mathbb{R}^d; \tag{2}$$
3. *The regularizer function $g$ is proper*[2] *and convex.*

Intuitively, item 1 of Assumption 1 guarantees the feasibility of the optimization problem (P) and item 2 imposes smoothness on the individual loss functions. For the set of convex regularizers, many of them (e.g., $\ell_1$, elastic net, etc) are not differentiable and therefore one cannot use gradient to evaluate the first-order stationary condition for such a nonsmooth composite objective function. This motivates us to introduce a generalized notion of gradient as we elaborate below.

We first introduce the following proximal mapping that is useful to handle the nonsmoothness of a function.

**Definition 1** (Proximal mapping)**.** *For any proper and convex function $g$, its proximal mapping evaluated at $x \in \mathbb{R}^d$ with parameter $\eta > 0$ is the unique point defined as*

$$\text{prox}_{\eta g}(x) := \arg\min_{u\in\mathbb{R}^d} \left\{ g(u) + \frac{1}{2\eta}\|u - x\|^2 \right\}.$$

The proximal mapping is uniquely defined for convex functions. Particularly, in the special case where $g$ is the indicator function of a convex set, its proximal mapping reduces to the projection operator onto the convex set. More importantly, the proximal mapping can be used to characterize the first-order stationary condition of nonsmooth composite functions in the following way.

**Fact 1.** *(Bauschke & Combettes, 2011) Let $g$ be a proper and convex function. Define the following notion of generalized gradient*

$$G_\eta(x, \nabla f(x)) := \frac{1}{\eta}\Big(x - \text{prox}_{\eta g}\big(x - \eta\nabla f(x)\big)\Big). \tag{3}$$

*Then, $x$ is a critical point of the function $F = f + g$ (i.e., $\mathbf{0} \in \nabla f(x) + \partial g(x)$) if and only if $G_\eta(x, \nabla f(x)) = \mathbf{0}$.*

Intuitively, $G_\eta(x, \nabla f(x))$ can be understood as a generalized notion of gradient for composite objective function. In the special case where $g \equiv 0$, the generalized gradient reduces to the usual notion of the gradient $\nabla f(x)$.

Based on the above definition, throughout the paper, we say that a point $x$ satisfies an *$\epsilon$-first-order stationary condition* of the problem (P) if $\|G_\eta(x, \nabla f(x))\| \leq \epsilon$.

## 3. Proximal SPIDER-M for Nonconvex Composite Optimization

In this section, we propose a proximal SPIDER algorithm that incorporates a momentum scheme (referred to as Proximal SPIDER-M) for solving the composite problem (P), and study its theoretical guarantee as well as the oracle complexity.

[2] An extended-valued function $h : \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ is called proper if its domain $\{x : h(x) < +\infty\}$ is non-empty.

### 3.1. Algorithm Design

We present the detailed update rule of the Proximal SPIDER-M in Algorithm 1, where "Unif" denotes the uniform sampling scheme with replacement.

**Algorithm 1** Proximal SPIDER-M

**Input:** $q, K \in \mathbb{N}, \{\lambda_k\}_{k=1}^{K-1}, \{\beta_k\}_{k=1}^{K-1} > 0$.
**Set:** $\alpha_k = \frac{2}{k+1}$ for $k = 1, ..., K-1$.
**Initialize:** $y_0 = x_0 \in \mathbb{R}^d$.
**for** $k = 0, 1, \ldots, K-1$ **do**
  $z_k = (1-\alpha_{k+1})y_k + \alpha_{k+1}x_k$,
  **if** $mod(k, q) = 0$ **then**
    set $v_k = \nabla f(z_k)$,
  **else**
    sample $\xi_k \overset{\text{Unif}}{\sim} \{1, \ldots, n\}$, and compute
    $v_k = \sum_{i\in\xi_k} \frac{1}{|\xi_k|}(\nabla \ell_i(z_k) - \nabla \ell_i(z_{k-1})) + v_{k-1}$.
  **end**
  $x_{k+1} = \text{prox}_{\lambda_k g}\big(x_k - \lambda_k v_k\big)$,
  $y_{k+1} = z_k - \frac{\beta_k}{\lambda_k}x_k + \frac{\beta_k}{\lambda_k}\text{prox}_{\lambda_k g}\big(x_k - \lambda_k v_k\big)$.
**end**
**Output:** $z_\zeta$, where $\zeta \overset{\text{Unif}}{\sim} \{0, \ldots, K-1\}$.

To elaborate on the algorithm design, note that Proximal SPIDER-M generates a tuple of variable sequences $\{x_k, y_k, z_k\}_k$ according to the momentum scheme. In specific, the variables $x_k$, $y_k$ are updated via proximal gradient-like steps using the gradient estimate $v_k$ proposed for SARAH in (Nguyen et al., 2017a;b) and different stepsizes $\lambda_k, \beta_k$, respectively. Then, their convex combination with momentum coefficient $\alpha_{k+1}$ yields the variable $z_{k+1}$. We choose a standard diminishing momentum coefficient $\alpha_k = \frac{2}{k+1}$ that serves for proving convergence guarantee in nonconvex optimization. We also note that the two updates for $x_{k+1}$ and $y_{k+1}$ do not introduce extra computation overhead as compared to a single update, since they both depend on the same stochastic gradient $v_k$.

We want to highlight the difference between our momentum scheme design for Proximal SPIDER-M and the existing momentum scheme design for proximal SGD in (Ghadimi & Lan, 2016) and proximal SVRG in (Allen-Zhu, 2017). In these works, they use the following proximal gradient steps for updating the variables $x_{k+1}$ and $y_{k+1}$:

$$x_{k+1} = \text{prox}_{\lambda_k g}\big(x_k - \lambda_k v_k\big), \tag{4}$$
$$y_{k+1} = \text{prox}_{\beta_k g}\big(z_k - \beta_k v_k\big). \tag{5}$$

Note that eq. (4) and eq. (5) use different proximal gradient updates that are based on $x_k$ and $z_k$, respectively. As a comparison, our momentum scheme in Algorithm 1 applies the same proximal gradient term $\text{prox}_{\lambda_k g}\big(x_k - \lambda_k v_k\big)$ to update both variables $x_{k+1}$ and $y_{k+1}$, and therefore requires less computation for evaluating the proximal mapping. Moreover, our update for the variable $y_{k+1}$ is not a single proximal gradient update (as opposed to eq. (5)), and it couples with the variables $z_k$ and $x_k$.

The momentum scheme introduced in (Ghadimi & Lan, 2016) based on eq. (4) and eq. (5) was shown to have convergence guarantee in nonconvex composite optimization only for convex regularizers $g$ that have a bounded domain. Therefore, it cannot yield a provable convergence guarantee for regularizers with unbounded domain, which are commonly used in practical applications, e.g., $\ell_1$, $\ell_2$, elastic net, etc. On the other hand, the momentum scheme introduced in (Allen-Zhu, 2017) was not proven to have a convergence guarantee in nonconvex optimization. In the next subsection, we prove that our momentum scheme in Algorithm 1 has a provable convergence guarantee for nonconvex composite optimization with arbitrary convex regularizers, therefore eliminating the restriction on the regularizers in (Ghadimi & Lan, 2016).

### 3.2. Convergence and Complexity Analysis

In this subsection, we study the convergence guarantee as well as the stochastic gradient oracle complexity of Proximal SPIDER-M for solving the problem (P). We obtain the following main theorem.

**Theorem 1.** *Let Assumption 1 hold. Apply the Proximal SPIDER-M (see Algorithm 1) to solve the problem (P) with parameters* $\alpha_k = \frac{2}{k+1}$, $q = |\xi_k| \equiv \sqrt{n}$, $\beta_k \equiv \frac{1}{8L}$ *and* $\lambda_k \in [\beta_k, (1+\alpha_k)\beta_k]$. *Then, the output* $z_\zeta$ *produced by the algorithm satisfies* $\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\| \leq \epsilon$ *for any* $\epsilon > 0$ *provided that the total number* $K$ *of iterations satisfies*

$$K \geq \Theta\bigg(\frac{L(F(x_0) - F^*)}{\epsilon^2}\bigg). \tag{6}$$

*Moreover, the total number of stochastic gradient oracle calls is at most* $\Theta(n + \sqrt{n}\epsilon^{-2})$ *and the total number of proximal mapping oracle calls is at most* $\Theta(\epsilon^{-2})$.

Theorem 1 establishes the convergence rate of Proximal SPIDER-M to satisfy the generalized first-order stationary condition and the corresponding oracle complexity. Specifically, the iteration complexity to achieve the generalized $\epsilon$-first-order stationary condition is in the order of $\Theta(\epsilon^{-2})$, which matches the state-of-art result of stochastic nonsmooth nonconvex optimization (Wang et al., 2018a). Furthermore, the corresponding stochastic gradient oracle complexity $\Theta(n + \sqrt{n}\epsilon^{-2})$ matches the lower bound for nonconvex optimization (Fang et al., 2018). Therefore, Proximal SPIDER-M enjoys the same optimal convergence guarantee as that for the Proximal SpiderBoost (Wang et al., 2018a) in nonconvex optimization, and it further benefits from the momentum scheme that can lead to significant acceleration in practical applications (as we demonstrate

via experiments in Section 6). We also note that the design of Proximal SPIDER-M allows to use constant stepsizes $\lambda_k, \beta_k = \Theta(\frac{1}{L})$ as opposed to the accuracy-dependent stepsize $\Theta(\frac{\epsilon}{L})$ adopted by the original SPIDER (Fang et al., 2018). This also facilitates the convergence of the algorithm in practice.

***Outline of Proof for Theorem 1***. As the technical proof is involved, we briefly outline the key intermediate steps below to convey some intuition on the analysis. The detailed proof is provided in the supplementary materials.

Based on the definition of generalized gradient (see Fact 1), we can rewrite the updates for $\{x_k, y_k, z_k\}_k$ in Algorithm 1 as follows:

$$\begin{aligned} z_k &= (1-\alpha_{k+1})y_k + \alpha_{k+1}x_k, \\ x_{k+1} &= x_k - \lambda_k G_{\lambda_k}(x_k, v_k), \\ y_{k+1} &= z_k - \beta_k G_{\lambda_k}(x_k, v_k). \end{aligned}$$

It can be seen that the term $G_{\lambda_k}(x_k, v_k)$ serves as a generalized gradient in the updates. Then, under the above momentum scheme, we can characterize the per-iteration progress of Proximal SPIDER-M by bounding the progressive function value gap $F(x_{k+1}) - F(x_k)$ as

$$\begin{aligned} F(x_{k+1}) - F(x_k) \leq & -\Theta\bigg(\lambda_k \|G_{\lambda_k}(x_k, v_k)\|^2\bigg) \\ & + \Theta\bigg(\Gamma_k \sum_{t=0}^{k-1} \frac{\lambda_t - \beta_t}{\alpha_{t+1}\Gamma_{t+1}} \|G_{\lambda_t}(x_t, v_t)\|^2\bigg) \\ & + \Theta\bigg(\lambda_k \|\nabla f(z_k) - v_k\|^2\bigg), \end{aligned}$$

where $\Gamma_k = \frac{2}{k(k+1)}$ and we have hidden the constant factors for simplicity of presentation. The next key step is to bound the estimation error term $\|\nabla f(z_k) - v_k\|^2$ in terms of the generalized gradient term $\|G_{\lambda_k}(x_k, v_k)\|^2$ as

$$\begin{aligned} \mathbb{E}\big[\|\nabla f(z_k) - v_k\|^2\big] \leq & \sum_{i=(\tau(k)-1)q}^{k-1} \frac{L^2}{|\xi_i|} \bigg[2\beta_i^2 \mathbb{E}\|G_{\lambda_i}(x_i, v_i)\|^2 \\ & + 2\alpha_{i+2}^2 \Gamma_{i+1} \sum_{t=0}^{i} \frac{(\lambda_t - \beta_t)^2}{\alpha_{t+1}\Gamma_{t+1}} \mathbb{E}\|G_{\lambda_t}(x_t, v_t)\|^2\bigg], \end{aligned}$$

where $\tau(k)$ denotes the index of the period that iteration $k$ belongs to. Then, combining the above two inequalities, telescoping and simplifying with much effort yield that

$$\mathbb{E}[F(x_K)] \leq F(x_0) - \sum_{k=0}^{K-1} \frac{\beta_k}{16} \mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2.$$

Based on the above result, we further exploit the randomized output strategy and finally obtain that

$$\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta, v_\zeta)\| \leq \Theta\bigg(\sqrt{\frac{L(F(x_0) - F^*)}{K}}\bigg),$$

where $\zeta$ is selected from $\{0, ..., K-1\}$ uniformly at random. Then, the desired convergence rate and oracle complexity results follow. □

From a technical perspective, we highlight the following three major new developments in the proof of Theorem 1 that is different from the proof for the basic stochastic gradient algorithm with momentum (Ghadimi & Lan, 2016) for nonconvex optimization: 1) our proof exploits the martingale structure of the SPIDER estimate $v_k$ which allows to bound the mean-square error term $\mathbb{E}\|\nabla f(z_k) - v_k\|^2$ in a tight way under the momentum scheme. In traditional analysis of stochastic algorithms with momentum (Ghadimi & Lan, 2016), such an error term corresponds to the variance of the stochastic estimator and is assumed to be bounded by a universal constant. 2): Our proof requires a very careful manipulation of the bounding strategy to handle the accumulation of the mean-square error $\mathbb{E}\|\nabla f(z_k) - v_k\|^2$ over the entire optimization path. 3): Our design of the momentum scheme allows to prove the convergence under arbitrary convex regularizers, whereas the proof of (Ghadimi & Lan, 2016) requires the regularizer to have a bounded domain.

# 4. Other Momentum Scheduling Schemes for Proximal SPIDER-M

It turns out that the design of Proximal SPIDER-M in Algorithm 1 allows to use more flexible momentum schemes in nonconvex optimization. In this section, we explore two variant momentum schemes for Proximal SPIDER-M that can be useful in practice and study the corresponding convergence guarantees.

## 4.1. Epochwise-diminishing Momentum

The Proximal SPIDER-M in Algorithm 1 uses a momentum coefficient $\alpha_k$ that diminishes to zero iterationwisely. As the epoch length $q$ usually consists of many inner iterations (e.g., multiple passes over the data), the momentum coefficient $\alpha_k$ can be very small after several epochs and hence leads to limited acceleration. Therefore, one strategy to alleviate such a problem is to set the momentum coefficient $\alpha_k$ to diminish epochwisely, i.e., set

(Epochwise-diminish momentum):

$$\alpha_k = \frac{2}{\lceil k/q \rceil + 1},\ k = 1, ..., K-1,$$

where $q \in \mathbb{N}$ corresponds to the number of inner iterations within each epoch and '$\lceil \cdot \rceil$' denotes the ceiling function. Under such a momentum setting, the momentum coefficient $\alpha_k$ remains to be constant within each epoch and diminishes slowly along progressive epochs. We note that a similar momentum coefficient setting is adopted in (Allen-Zhu, 2017;

Shang et al., 2018) for accelerating SVRG. However, the focus there is to solve convex optimization problems and no convergence guarantee was established for nonconvex optimization.

### 4.2. Epochwise-restart Momentum

Another widely used momentum setting is to restart the momentum scheme after a fixed number of iterations. Specifically in the context of Proximal SPIDER-M, we synchronize the variables $x_{k+1}$ and $y_{k+1}$ to be the $x_k$ obtained in the previous iteration after every epoch, i.e., we add the following algorithmic code to the Proximal SPIDER-M in Algorithm 1

$$\textbf{If mod}(k, q-1) = 0 \textbf{ then}$$
$$\textbf{set } y_{k+1} = x_{k+1} = x_k.$$

This can be understood as a reinitialization of the variables epochwisely. On the other hand, we restart the momentum coefficient $\alpha_k$ after every epoch as

$$\alpha_k = \overline{\alpha}_{\mathrm{mod}(k,q)}, \text{ where } \overline{\alpha}_t = \frac{2}{t+1}, \ t = 1, ..., q-1.$$

Under such a momentum scheme, the momentum coefficient $\alpha_k$ reboots at the beginning of every epoch, injecting a periodic momentum into the algorithm dynamic consistently. Finally, the algorithm outputs the point $z_\zeta$ where $\zeta$ is selected from $\{k : 0 \le k \le K-1, \mathrm{mod}(k, q-1) \neq 0\}$ uniformly at random.

The momentum scheme with restart has been applied to the gradient descent algorithm in (O'donoghue & Candès, 2015). There, it has been justified that a proper restart scheme can significantly accelerate the practical convergence of the algorithm. However, it is unclear whether a restart momentum scheme can have a convergence guarantee in nonconvex and nonsmooth optimization, especially under the more sample-efficient SPIDER scheme. We establish such a theoretical result in the next subsection.

To further illustrate the differences among these three momentum schemes, we plot and compare the scheduling of the momentum coefficient $\alpha_k$ of these momentum schemes in Figure 1. The area below each curve roughly corresponds to the total momentum that is injected into the algorithm dynamic by the corresponding momentum scheme. One can see that the original momentum scheme that diminishes $\alpha_k$ iterationwisely has the smallest total momentum, whereas the epochwise-diminishing momentum scheme has the largest total momentum (within a considerable number of epochs). We further demonstrate that the practical performance of these momentum schemes is highly correlated with the accumulative momentum via numerical experiments in Section 6.

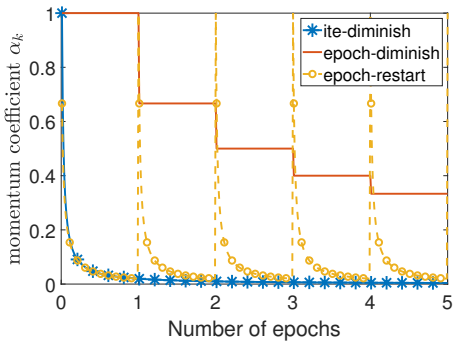


Figure 1: Momentum coefficient scheduling of different momentum schemes.

### 4.3. Convergence and Complexity Analysis

In this subsection, we present the convergence result and the corresponding oracle complexity of Proximal SPIDER-M under the variants of momentum schemes introduced in the previous subsections. We obtain the following main result.

**Theorem 2.** *Let Assumption 1 hold. Apply the Proximal SPIDER-M with either epochwise-diminishing momentum or epochwise-restart momentum to solve the problem (P). Set parameters* $q = |\xi_k| \equiv \sqrt{n}$, $\beta_k \equiv \frac{1}{8L}$ *and* $\lambda_k \in [\beta_k, (1+\alpha_k)\beta_k]$. *Then, the output* $z_\zeta$ *of the algorithm satisfies* $\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\| \le \epsilon$ *for any* $\epsilon > 0$ *under the same complexity requirements as those in Theorem 1.*

From Theorem 2, it can be seen that the Proximal SPIDER-M maintains the optimal stochastic gradient oracle complexity in nonconvex optimization under the more flexible epochwise-diminishing and the epochwise restart momentum schemes. Therefore, this demonstrates that the algorithmic structure of SPIDER provides much flexibility in designing compatible momentum schemes in the nonconvex regime.

## 5. Proximal SPIDER-M for Online Nonconvex Composite Optimization

The objective function $\frac{1}{n}\sum_{i=1}^n \ell_i(x)$ in the optimization problem (P) contains a finite number of data samples that are typically drawn from an underlying data distribution. Therefore, it can be viewed as a finite-sample approximation of the population risk $\mathbb{E}_{u\sim\mathbb{U}}[\ell_u(x)]$, where the data sample $u$ is generated from an underlying distribution $\mathbb{U}$. In this section, we study the following online composite optimization problem that involves the population risk:

$$\min_{x\in\mathbb{R}^d} F(x) := f(x) + g(x), \qquad \text{(R)}$$
$$\text{where } f(x) := \mathbb{E}_{u\sim\mathbb{U}}[\ell_u(x)],$$

where the function $g$ corresponds to the regularizer. As the problem (R) depends on the population risk that contains infinite samples, we propose a variant of Proximal SPIDER-M that can solve it in an online setting. We summarize the detailed steps of the algorithm in Algorithm 2, where we refer to it as Online Proximal SPIDER-M.

**Algorithm 2** Online Proximal SPIDER-M

**Input:** $q, K \in \mathbb{N}, \{\lambda_k\}_{k=1}^{K-1}, \{\beta_k\}_{k=1}^{K-1} > 0$.
**Set:** $\alpha_k = \frac{2}{k+1}$ for $k = 1, ..., K-1$.
**Initialize:** $y_0 = x_0 \in \mathbb{R}^d$.
**for** $k = 0, 1, \ldots, K-1$ **do**
  $z_k = (1-\alpha_{k+1})y_k + \alpha_{k+1}x_k$,
  **if** *mod*$(k, q) = 0$ **then**
    draw $\xi_1$ data samples from $\mathbb{U}$ and compute
    $v_k = \frac{1}{|\xi_1|}\sum_{i=1}^{|\xi_1|}\nabla \ell_{u_i}(x)$
  **else**
    draw $\xi_2$ data samples from $\mathbb{U}$ and compute
    $v_k = \frac{1}{|\xi_2|}\sum_{i=1}^{|\xi_2|}(\nabla \ell_{u_i}(z_k) - \nabla \ell_{u_i}(z_{k-1})) + v_{k-1}$.
  **end**
  $x_{k+1} = \mathrm{prox}_{\lambda_k g}\big(x_k - \lambda_k v_k\big)$,
  $y_{k+1} = z_k - \frac{\beta_k}{\lambda_k}x_k + \frac{\beta_k}{\lambda_k}\mathrm{prox}_{\lambda_k g}\big(x_k - \lambda_k v_k\big)$.
**end**
**Output:** $z_\zeta$, where $\zeta \overset{\text{Unif}}{\sim} \{0, \ldots, K-1\}$.

Note that unlike the Proximal SPIDER-M for the finite-sum case, the Online Proximal SPIDER-M keeps drawing new data samples from the underlying distribution $\mathbb{U}$ (uniformly at random) to construct the gradient estimate $v_k$. To study its convergence guarantee, we make the following standard assumption on the variance of the random sampling.

**Assumption 2.** *There exists a constant $\sigma > 0$ such that for all $x \in \mathbb{R}^d$ and all random samples $u \sim \mathbb{U}$, it holds that $\mathbb{E}_{u\sim\mathbb{U}}\|\nabla \ell_u(x) - \nabla f(x)\|^2 \leqslant \sigma^2$.*

Based on Assumption 2, we obtain the following convergence guarantee for Online Proximal SPIDER-M.

**Theorem 3.** *Let Assumptions 1 and 2 hold. Apply Online Proximal SPIDER-M (see Algorithm 2) to solve the problem (R). Choose any desired accuracy $\epsilon > 0$ and set parameters $\alpha_k = \frac{2}{k+1}, q = |\xi_2| = \sqrt{|\xi_1|} = \sqrt{\frac{2\sigma^2}{\epsilon^2}}$, $\beta_k \equiv \frac{1}{8L}$ and $\lambda_k \in [\beta_k, (1+\alpha_k)\beta_k]$. Then, the output $z_\zeta$ of the algorithm satisfies $\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\| \leq \epsilon$ provided that the total number of iterations $K$ satisfies*

$$K \geq \Theta\bigg(\frac{L(F(x_0) - F^*)}{\epsilon^2}\bigg). \tag{7}$$

*Moreover, the total number of stochastic gradient calls is at most $\Theta(\epsilon^{-3})$ and the total number of proximal mapping calls is at most $\Theta(\epsilon^{-2})$.*

The orders of the results in Theorem 3 match those of state-of-arts (Fang et al., 2018; Wang et al., 2018a). Our result demonstrates that the momentum scheme can be applied to facilitate the convergence of Proximal SPIDER for solving online nonsmooth and nonconvex problems with a provable convergence guarantee. Moreover, we obtain a similar complexity result for Online Proximal SPIDER-M under the other two momentum schemes proposed in Section 4 in the following theorem.

**Theorem 4.** *Let Assumptions 1 and 2 hold. Apply Online Prox-SPIDER-M with either epochwise-diminishing momentum or epochwise restart momentum to solve the problem (R). Choose any desired accuracy $\epsilon > 0$ and set parameters $q = |\xi_2| = \sqrt{|\xi_1|} = \sqrt{\frac{2\sigma^2}{\epsilon^2}}$, $\beta_k \equiv \frac{1}{8L}$ and $\lambda_k \in [\beta_k, (1+\alpha_k)\beta_k]$. Then, the output $z_\zeta$ of the algorithm satisfies $\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\| \leq \epsilon$ under the same complexity requirements as those in Theorem 3.*

## 6. Experiments

In this section, we compare the practical performance of the following stochastic variance-reduced algorithms: SVRG in (Johnson & Zhang, 2013), SpiderBoost in (Wang et al., 2018a), Katyusha$^{ns}$ in (Allen-Zhu, 2017), ASVRG in (Shang et al., 2018), RSAG in (Ghadimi & Lan, 2016), Proximal SPIDER-M (Algorithm 1 in this paper), Proximal SPIDER-MED (epochwise-diminishing momentum) and Proximal SPIDER-MER (epochwise restart momentum). We note that all algorithms use certain momentum schemes except for SVRG and SpiderBoost. For all algorithms considered, we set their learning rates to be $0.05$. For each experiment, we initialize all the algorithms at the same point that is generated randomly from the normal distribution. Also, we choose a fixed mini-batch size $256$ and set the epoch length $q$ to be $2n/256$ such that all algorithms pass over the entire dataset twice in each epoch.

### 6.1. Un-regularized Nonconvex Optimization

We first apply these algorithms to solve an un-regularized nonconvex optimization problem. The first problem is the following nonconvex logistic regression problem

$$\min_{w\in\mathbb{R}^d} f(w) := \frac{1}{n}\sum_{i=1}^{n}\ell(w^\intercal x_i, y_i) + \alpha\sum_{i=1}^{d}\frac{w_i^2}{1+w_i^2},$$

where $x_i \in \mathbb{R}^d$ denotes the features and $y_i \in \{\pm 1\}$ corresponds to the labels, and we set the loss $\ell$ to be the cross-entropy loss and $\alpha = 0.1$. For this problem, we use two different datasets from the LIBSVM (Chang & Lin, 2011): the a9a dataset ($n = 32561, d = 123$) and the w8a dataset ($n = 49749, d = 300$). We report the learning curves on the function value gap of these algorithms in Figure 2.

In this experiment, one can see from Figure 2 that our SPIDER-MED with epochwise diminishing momentum achieves the best performance and significantly outperforms other algorithms. Also, we note that the performances of both Katyusha$^{ns}$ and ASVRG do not achieve much acceleration in such a nonconvex case, as these algorithms are originally developed to accelerate solving convex problems. This demonstrates that our design of SPIDER-M has a stable performance in nonconvex optimization as well as provable theoretical guarantee. We note that the curve of SpiderBoost overlaps with that of SVRG (both algorithms have similar performance). On the other hand, among all SPIDER-M-type of algorithms, the one that uses the epochwise diminishing momentum (SPIDER-MED) has the best performance, whereas the one that uses the iterationwise diminish momentum (SPIDER-M) is the slowest. This corroborates the comparison of the total momentum that we illustrate in Figure 1.

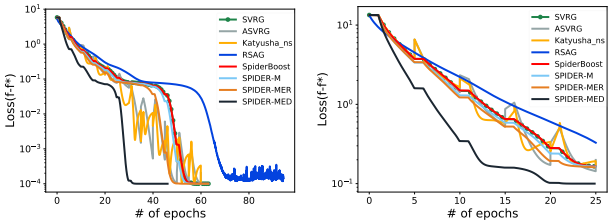


(a) Dataset: a9a (b) Dataset: w8a

Figure 2: Comparison among stochastic variance-reduced algorithms for solving nonconvex logistic regression.

Next, we compare these algorithms in solving the following nonconvex robust linear regression problem

$$\min_{w\in\mathbb{R}^d} f(w) := \frac{1}{n}\sum_{i=1}^{n} \ell(y_i - w^\intercal x_i),$$

where we use the nonconvex loss function $\ell(x) := \log(\frac{x^2}{2} + 1)$. We report the learning curves on the function value gap of these algorithms in Figure 3. One can see that our SPIDER-MED with epochwise diminishing momentum has a comparable performance to that of Katyusha$^{ns}$, and they both outperform other algorithms.

### 6.2. Nonsmooth and Nonconvex Optimization

We further add an $\ell_1$ nonsmooth regularizer with weight coefficient $0.1$ to the objective functions of the above two optimization problems, and apply the corresponding proximal versions of these algorithms to solve the nonconvex composite optimization problems. All the results are presented in Figures 4 and 5. One can see that our Proximal SPIDER-MED still significantly outperforms all the other algorithms in these nonconvex and nonsmooth scenarios. This demonstrates that our novel design of the coupled update for $\{y_k\}_k$ in the momentum scheme is efficient in the

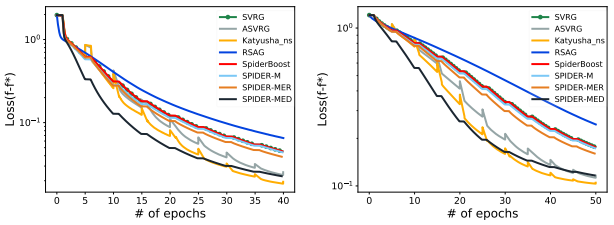


(a) Dataset: a9a (b) Dataset: w8a

Figure 3: Comparison among stochastic variance-reduced algorithms for solving robust linear regression.

nonsmooth and nonconvex setting. Also, it turns out that Katyusha$^{ns}$ and ASVRG are suffering from a slow convergence (their convergences occur at around 40 epochs). Together with the first two experiments, this implies that their performance is not stable and may not be generally suitable for solving nonsmooth and nonconvex problems.

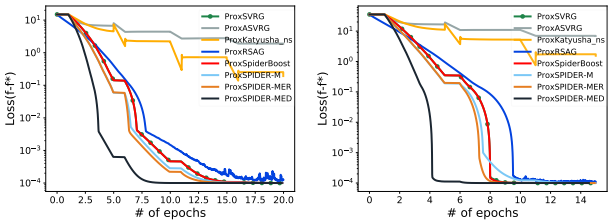


(a) Dataset: a9a (b) Dataset: w8a

Figure 4: Comparison among stochastic variance-reduced algorithms for solving logistic regression with $\ell_1$ regualarizer.

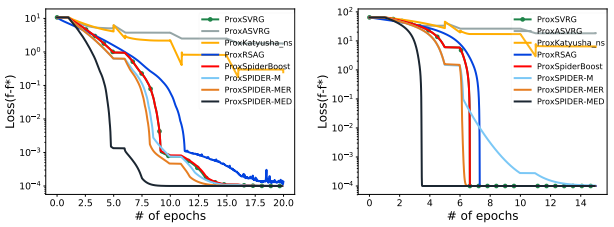


(a) Dataset: a9a (b) Dataset: w8a

Figure 5: Comparison among stochastic variance-reduced algorithms for solving robust linear regression with $\ell_1$ regualarizer.

## 7. Conclusion

In this paper, we design an efficient proximal stochastic variance-reduced algorithm with momentum to solve nonconvex composite optimization problems with provable convergence guarantee. Under a basic momentum scheme, we show that our Proximal SPIDER-M achieves the best possible stochastic gradient oracle complexity for nonconvex

optimization. Our algorithm design further allows to apply other momentum schemes and to solve online composite optimization problems with an optimal oracle complexity. We anticipate our algorithm design to inspire the development of more advanced momentum acceleration schemes for stochastic nonconvex optimization. On the other hand, it is also interesting to explore whether our algorithm can achieve the best possible convergence rate in convex optimization.

## References


Allen-Zhu, Z. Katyusha: The first direct acceleration of stochastic gradient methods. *Journal of Machine Learning Research*, 18(1):8194–8244, January 2017.

Allen-Zhu, Z. Katyusha x: Simple momentum method for stochastic sum-of-nonconvex optimization. In *Proc. International Conference on Machine Learning (ICML)*, volume 80, pp. 179–185, 10–15 Jul 2018.

Allen-Zhu, Z. and Hazan, E. Variance reduction for faster non-convex optimization. In *Proc. International Conference on Machine Learning (ICML)*, pp. 699–707, 2016.

Bauschke, H. and Combettes, P. L. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2011.

Beck, A. and Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, March 2009.

Bottou, L. Large-scale machine learning with stochastic gradient descent. In *Proc. International Conference on Computational Statistics (COMPSTAT)*, pp. 177–186, 2010.

Chang, C. and Lin, C. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):1–27, 2011.

Defazio, A., Bach, F., and Lacoste-Julien, S. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 1646–1654. 2014.

Fang, C., Li, C., Lin, Z., and Zhang, T. SPIDER: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 687–697. 2018.

Ghadimi, S. and Lan, G. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Mathematical Programming*, 156(1-2):59–99, March 2016.

Ghadimi, S., Lan, G., and Zhang, H. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1): 267–305, Jan 2016.

Goodfellow, I., Bengio, Y., and Courville, A. *Deep Learning*. MIT Press, 2016.

Johnson, R. and Zhang, T. Accelerating stochastic gradient descent using predictive variance reduction. In *Proc. International Conference on Neural Information Processing Systems (NIPS)*, pp. 315–323, 2013.

Li, Q., Zhou, Y., Liang, Y., and Varshney, P. K. Convergence analysis of proximal gradient with momentum for nonconvex optimization. In *Proc. International Conference on Machine Learning (ICML)*, volume 70, pp. 2111–2119, 2017.

Nesterov, Y. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer, 2014.

Nguyen, L., Liu, J., Scheinberg, K., and Takáč, K. Stochastic recursive gradient algorithm for nonconvex optimization. *ArXiv:1705.07261*, May 2017a.

Nguyen, L., Liu, J., Scheinberg, K., and Takáč, M. SARAH: A novel method for machine learning problems using stochastic recursive gradient. In *Proc. International Conference on Machine Learning (ICML)*, 2017b.

Nitanda, A. Accelerated stochastic gradient descent for minimizing finite sums. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 51, pp. 195–203, May 2016.

O'donoghue, B. and Candès, E. Adaptive restart for accelerated gradient schemes. *Foundations of Computational Mathematics*, 15(3):715–732, June 2015.

Reddi, S. J., Sra, S., Poczos, B., and Smola, A. Proximal stochastic methods for nonsmooth nonconvex finite-sum optimization. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 1145–1153. 2016.

Robbins, H. and Monro, S. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3): 400–407, 09 1951.

Schmidt, M., Le Roux, N., and Bach, F. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162(1):83–112, Mar 2017.

Shang, F., Jiao, L., Zhou, K., Cheng, J., Ren, Y., and Jin, Y. Asvrg: Accelerated proximal svrg. In *Proc. Asian Conference on Machine Learning*, volume 95, pp. 815–830, 2018.

Tseng, P. Approximation accuracy, gradient methods, and error bound for structured convex optimization. *Mathematical Programming*, 125(2):263–295, Oct 2010.

Wang, Z., Ji, K., Zhou, Y., Liang, Y., and Tarokh, V. SpiderBoost: A class of faster variance-reduced algorithms for nonconvex optimization. *ArXiv:1810.10690*, October 2018a.

Wang, Z., Zhou, Y., Liang, Y., and Lan, G. Cubic regularization with momentum for nonconvex optimization. *ArXiv:1810.03763*, October 2018b.

Xiao, L. and Zhang, T. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.

Zhang, J., Zhang, H., and Sra, S. R-SPIDER: A Fast Riemannian Stochastic Optimization Algorithm with Curvature Independent Rate. *ArXiv:811.04194*, 2018.

Zhou, D., Xu, P., and Gu, Q. Stochastic nested variance reduced gradient descent for nonconvex optimization. In *Proc. Advances in Neural Information Processing Systems (NeurIps)*, pp. 3925–3936. 2018.

Zhou, P., Yuan, X., and Feng, J. Faster First-Order Methods for Stochastic Non-Convex Optimization on Riemannian Manifolds. 2019.

# Supplementary Materials

## A. Auxiliary Lemmas for Analysis of Algorithm 1

In this section, we collect some auxiliary results that facilitate the analysis of Algorithm 1. Throughout, for any $k \in \mathbb{N}$, denote $\tau(k) \in \mathbb{N}$ the unique integer such that $(\tau(k)-1)q \leq k \leq \tau(k)q - 1$. We also define $\Gamma_0 = 0, \Gamma_1 = 1$ and $\Gamma_k = (1-\alpha_k)\Gamma_{k-1}$ for $k = 2, 3, ....$ Since we set $\alpha_k = \frac{2}{k+1}$, it is easy to check that $\Gamma_k = \frac{2}{k(k+1)}$.

We first present an auxiliary lemma from (Fang et al., 2018).

**Lemma 1.** *(Fang et al., 2018) Under Assumption 1, the estimation $v_k$ of gradient constructed by SPIDER satisfies that for all $(\tau(k)-1)q + 1 \leq k \leq \tau(k)q - 1$,*

$$\mathbb{E}\|v_k - \nabla f(z_k)\|^2 \leq \frac{L^2}{|\xi_k|}\mathbb{E}\|z_k - z_{k-1}\|^2 + \mathbb{E}\|v_{k-1} - \nabla f(z_{k-1})\|^2.$$

Telescoping Lemma 1 and noting that $v_k = \nabla f(z_k)$ for all $k$ such that $\mathrm{mod}(k, q) = 0$, we obtain the following bound.

**Lemma 2.** *Under Assumption 1, the estimation $v_k$ of gradient constructed by SPIDER satisfies that for all $k \in \mathbb{N}$,*

$$\mathbb{E}\|v_k - \nabla f(z_k)\|^2 \leq \sum_{i=(\tau(k)-1)q}^{k-1} \frac{L^2}{|\xi_i|}\mathbb{E}\|z_{i+1} - z_i\|^2. \tag{8}$$

Next, recall the following definition of the gradient mapping for some $\eta > 0$ and $x, u \in \mathbb{R}^d$:

$$G_\eta(x, u) := \frac{1}{\eta}\big(x - \mathrm{prox}_{\eta g}(x - \eta u)\big).$$

Based on this definition, we can rewrite the updates of Algorithm 1 as follows:

$$\begin{aligned} z_k &= (1-\alpha_{k+1})y_k + \alpha_{k+1}x_k, \\ x_{k+1} &= x_k - \lambda_k G_{\lambda_k}(x_k, v_k), \\ y_{k+1} &= z_k - \beta_k G_{\lambda_k}(x_k, v_k). \end{aligned}$$

Next, we prove the following auxiliary lemma.

**Lemma 3.** *Let the sequences $\{x_k\}_k, \{y_k\}_k, \{z_k\}_k$ be generated by Algorithm 1. Then, the following inequalities hold*

$$y_k - x_k = \Gamma_k \sum_{t=1}^{k} \frac{\lambda_{t-1} - \beta_{t-1}}{\Gamma_t} G_{\lambda_{t-1}}(x_{t-1}, v_{t-1}), \tag{9}$$

$$\|y_k - x_k\|^2 \leq \Gamma_k \sum_{t=1}^{k} \frac{\lambda_{t-1} - \beta_{t-1}}{\alpha_t \Gamma_t} \|G_{\lambda_{t-1}}(x_{t-1}, v_{t-1})\|^2, \tag{10}$$

$$\|z_{k+1} - z_k\|^2 \leq 2\beta_k^2 \|G_{\lambda_k}(x_k, v_k)\|^2 + 2\alpha_{k+2}^2 \Gamma_{k+1} \sum_{t=1}^{k+1} \frac{(\lambda_{t-1} - \beta_{t-1})^2}{\alpha_t \Gamma_t} \|G_{\lambda_{t-1}}(x_{t-1}, v_{t-1})\|^2. \tag{11}$$

*Proof.* We prove the first equality. By the update rule of the momentum scheme, we obtain that

$$\begin{aligned} y_k - x_k &= z_{k-1} - \beta_{k-1} G_{\lambda_{k-1}}(x_{k-1}, v_{k-1}) - (x_{k-1} - \lambda_{k-1} G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})) \\ &= (1-\alpha_k)(y_{k-1} - x_{k-1}) + (\lambda_{k-1} - \beta_{k-1}) G_{\lambda_{k-1}}(x_{k-1}, v_{k-1}). \end{aligned} \tag{12}$$

Dividing both sides by $\Gamma_k$ and noting that $\frac{1-\alpha_k}{\Gamma_k} = \Gamma_{k-1}$, we further obtain that

$$\frac{y_k - x_k}{\Gamma_k} = \frac{y_{k-1} - x_{k-1}}{\Gamma_{k-1}} + \frac{\lambda_{k-1} - \beta_{k-1}}{\Gamma_k} G_{\lambda_{k-1}}(x_{k-1}, v_{k-1}). \tag{13}$$

Telescoping the above equality over $k$ yields the first desired equality.

Next, we prove the second inequality. Based on the first equality, we obtain that

$$\begin{aligned}\|y_k - x_k\|^2 &= \|\Gamma_k \sum_{t=1}^{k} \frac{\lambda_{t-1} - \beta_{t-1}}{\Gamma_t} G_{\lambda_{t-1}}(x_{t-1}, v_{t-1})\|^2 \\ &= \|\Gamma_k \sum_{t=1}^{k} \frac{\alpha_t}{\Gamma_t} \frac{\lambda_{t-1} - \beta_{t-1}}{\alpha_t} G_{\lambda_{t-1}}(x_{t-1}, v_{t-1})\|^2 \\ &\overset{(i)}{\leq} \Gamma_k \sum_{t=1}^{k} \frac{\alpha_t}{\Gamma_t} \frac{(\lambda_{t-1} - \beta_{t-1})^2}{\alpha_t^2} \|G_{\lambda_{t-1}}(x_{t-1}, v_{t-1})\|^2 \\ &= \Gamma_k \sum_{t=1}^{k} \frac{(\lambda_{t-1} - \beta_{t-1})^2}{\Gamma_t \alpha_t} \|G_{\lambda_{t-1}}(x_{t-1}, v_{t-1})\|^2, \end{aligned} \tag{14}$$

where (i) uses the facts that $\{\Gamma_k\}_k$ is a decreasing sequence, $\sum_{t=1}^{k} \frac{\alpha_t}{\Gamma_t} = \frac{1}{\Gamma_k}$ and Jensen's inequality.

Finally, we prove the third inequality. By the update rule of the momentum scheme, we obtain that $z_{k+1} - z_k = y_{k+1} - z_k + \alpha_{k+2}(x_{k+1} - y_{k+1})$. Then, we further obtain that

$$\begin{aligned}\|z_{k+1} - z_k\| &\leq \|y_{k+1} - z_k\| + \alpha_{k+2}\|x_{k+1} - y_{k+1}\| \\ &\leq \beta_k \|G_{\lambda_k}(x_k, v_k)\| + \alpha_{k+2}\sqrt{\|x_{k+1} - y_{k+1}\|^2} \\ &\leq \beta_k \|G_{\lambda_k}(x_k, v_k)\| + \alpha_{k+2}\sqrt{\Gamma_{k+1} \sum_{t=1}^{k+1} \frac{(\lambda_{t-1} - \beta_{t-1})^2}{\Gamma_t \alpha_t} \|G_{\lambda_{t-1}}(x_{t-1}, v_{t-1})\|^2}. \end{aligned}$$

The desired result follows by taking the square on both sides of the above inequality and using the fact that $(a+b)^2 \leq 2a^2 + 2b^2$. □

We also need the following lemma, which was established as Lemma 1 and Proposition 1 in (Ghadimi et al., 2016).

**Lemma 4** (Lemma 1 and Proposition 1, (Ghadimi et al., 2016))**.** *Let $g$ be a proper and closed convex function. Then, for all $u, v, x \in \mathbb{R}^d$ and $\eta > 0$, the following statements hold:*

$$\begin{aligned}&\langle u, G_\eta(x, u)\rangle \geq \|G_\eta(x, u)\|^2 + \frac{1}{\eta}\big(g(\mathrm{prox}_{\eta g}(x - \eta u)) - g(x)\big), \\ &\|G_\eta(x, u) - G_\eta(x, v)\| \leq \|u - v\|. \end{aligned}$$

## B. Proof of Theorem 1

Consider any iteration $k$ of the algorithm. By smoothness of $f$, we obtain that

$$\begin{aligned}f(x_k) &\leq f(x_{k-1}) + \langle \nabla f(x_{k-1}), x_k - x_{k-1}\rangle + \frac{L}{2}\|x_k - x_{k-1}\|^2 \\ &= f(x_{k-1}) + \left\langle \nabla f(x_{k-1}), -\lambda_{k-1} G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\right\rangle + \frac{L\lambda_{k-1}^2}{2}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 \\ &= f(x_{k-1}) - \lambda_{k-1}\left\langle \nabla f(x_{k-1}) - v_{k-1}, G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\right\rangle - \lambda_{k-1}\left\langle v_{k-1}, G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\right\rangle \\ &\quad + \frac{L\lambda_{k-1}^2}{2}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 \\ &\overset{(i)}{\leq} f(x_{k-1}) - \lambda_{k-1}\left\langle \nabla f(x_{k-1}) - v_{k-1}, G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\right\rangle - \lambda_{k-1}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 \\ &\quad - \big(g(\mathrm{prox}_{\lambda_{k-1} g}(x_{k-1} - \lambda_{k-1} v_{k-1})) - g(x_{k-1})\big) + \frac{L\lambda_{k-1}^2}{2}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 \end{aligned}$$

$$= f(x_{k-1}) - \lambda_{k-1}\left\langle \nabla f(x_{k-1}) - v_{k-1}, G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\right\rangle - \lambda_{k-1}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2$$
$$- \big(g(x_k) - g(x_{k-1})\big) + \frac{L\lambda_{k-1}^2}{2}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2,$$

where (i) follows from Lemma 4. Rearranging the above inequality and using Cauchy-Swartz inequality yields that

$$F(x_k) \le F(x_{k-1}) - \lambda_{k-1}(1 - \frac{L\lambda_{k-1}}{2})\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 + \lambda_{k-1}\|\nabla f(x_{k-1}) - v_{k-1}\|\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|. \tag{15}$$

Note that

$$\begin{aligned}\|\nabla f(x_{k-1}) - v_{k-1}\| &\le \|\nabla f(x_{k-1}) - \nabla f(z_{k-1})\| + \|\nabla f(z_{k-1}) - v_{k-1}\| \\ &\overset{(i)}{\le} L\|x_{k-1} - z_{k-1}\| + \|\nabla f(z_{k-1}) - v_{k-1}\| \\ &\overset{(ii)}{\le} L(1-\alpha_k)\|y_{k-1} - x_{k-1}\| + \|\nabla f(z_{k-1}) - v_{k-1}\|,\end{aligned}$$

where (i) uses the Lipschitz continuity of $\nabla f$ and (ii) follows from the update rule of the momentum scheme. Substituting the above inequality into eq. (15) yields that

$$\begin{aligned}F(x_k) &\le F(x_{k-1}) - \lambda_{k-1}(1 - \frac{L\lambda_{k-1}}{2})\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 + L\lambda_{k-1}(1-\alpha_k)\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|\|y_{k-1} - x_{k-1}\| \\ &\quad + \lambda_{k-1}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|\|\nabla f(z_{k-1}) - v_{k-1}\| \\ &\le F(x_{k-1}) - \lambda_{k-1}(1 - \frac{L\lambda_{k-1}}{2})\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 + \frac{L\lambda_{k-1}^2}{2}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 + \frac{L(1-\alpha_k)^2}{2}\|y_{k-1} - x_{k-1}\|^2 \\ &\quad + \frac{\lambda_{k-1}}{2}\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 + \frac{\lambda_{k-1}}{2}\|\nabla f(z_{k-1}) - v_{k-1}\|^2 \\ &= F(x_{k-1}) - \lambda_{k-1}(\frac{1}{2} - L\lambda_{k-1})\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 + \frac{L(1-\alpha_k)^2}{2}\|y_{k-1} - x_{k-1}\|^2 + \frac{\lambda_{k-1}}{2}\|\nabla f(z_{k-1}) - v_{k-1}\|^2 \\ &\le F(x_{k-1}) - \lambda_{k-1}(\frac{1}{2} - L\lambda_{k-1})\|G_{\lambda_{k-1}}(x_{k-1}, v_{k-1})\|^2 + \frac{L\Gamma_{k-1}}{2}\sum_{t=1}^{k-1}\frac{\lambda_{t-1} - \beta_{t-1}}{\alpha_t\Gamma_t}\|G_{\lambda_{t-1}}(x_{t-1}, v_{t-1})\|^2 \\ &\quad + \frac{\lambda_{k-1}}{2}\|\nabla f(z_{k-1}) - v_{k-1}\|^2,\end{aligned}$$

where the last inequality uses item 2 of Lemma 3 and the fact that $0 < \alpha_k < 1$. Telescoping the above inequality over $k$ from 1 to $K$ yields that

$$\begin{aligned}F(x_K) &\le F(x_0) - \sum_{k=0}^{K-1}\lambda_k(\frac{1}{2} - L\lambda_k)\|G_{\lambda_k}(x_k, v_k)\|^2 + \sum_{k=0}^{K-1}\frac{L\Gamma_k}{2}\sum_{t=0}^{k-1}\frac{(\lambda_t - \beta_t)^2}{\Gamma_{t+1}\alpha_{t+1}}\|G_{\lambda_t}(x_t, v_t)\|^2 + \sum_{k=0}^{K-1}\frac{\lambda_k}{2}\|\nabla f(z_k) - v_k\|^2 \\ &= F(x_0) - \sum_{k=0}^{K-1}\lambda_k(\frac{1}{2} - L\lambda_k)\|G_{\lambda_k}(x_k, v_k)\|^2 + \frac{L}{2}\sum_{k=0}^{K-1}\frac{(\lambda_k - \beta_k)^2}{\Gamma_{k+1}\alpha_{k+1}}\|G_{\lambda_k}(x_k, v_k)\|^2(\sum_{t=k}^{K-1}\Gamma_t) \\ &\quad + \sum_{k=0}^{K-1}\frac{\lambda_k}{2}\|\nabla f(z_k) - v_k\|^2,\end{aligned} \tag{16}$$

where we have exchanged the order of summation in the second equality. Furthermore, note that $\sum_{t=k}^{K-1}\Gamma_t = 2\sum_{t=k}^{K-1}\frac{1}{t} - \frac{1}{t+1} \le \frac{2}{k}$. Then, substituting this bound into the above inequality and taking expectation on both sides yield that

$$\begin{aligned}\mathbb{E}[F(x_K)] &\le F(x_0) - \sum_{k=0}^{K-1}\lambda_k(\frac{1}{2} - L\lambda_k)\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + \frac{L}{2}\sum_{k=0}^{K-1}\frac{2(\lambda_k - \beta_k)^2}{k\Gamma_{k+1}\alpha_{k+1}}\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 \\ &\quad + \sum_{k=0}^{K-1}\frac{\lambda_k}{2}\mathbb{E}\|\nabla f(z_k) - v_k\|^2.\end{aligned} \tag{17}$$

Next, we bound the term $\mathbb{E}\|\nabla f(z_k) - v_k\|^2$ in the above inequality. By Lemma 2 we obtain that

$$\begin{aligned}\mathbb{E}\|\nabla f(z_k) - v_k\|^2 &\leq \sum_{i=(\tau(k)-1)q}^{k-1} \frac{L^2}{|\xi_i|}\mathbb{E}\|z_{i+1} - z_i\|^2 \\ &\leq \sum_{i=(\tau(k)-1)q}^{k-1} \frac{L^2}{|\xi_i|}\big[2\beta_i^2\|G_{\lambda_i}(x_i, v_i)\|^2 + 2\alpha_{i+2}^2\Gamma_{i+1}\sum_{t=0}^{i}\frac{(\lambda_t - \beta_t)^2}{\alpha_{t+1}\Gamma_{t+1}}\|G_{\lambda_t}(x_t, v_t)\|^2\big], \end{aligned} \tag{18}$$

where the last inequality uses item 3 of Lemma 3. Substituting eq. (18) into eq. (17) and simplifying yield that

$$\begin{aligned}\mathbb{E}[F(x_K)] \leq F(x_0) &- \sum_{k=0}^{K-1}\Big[\lambda_k(\frac{1}{2} - L\lambda_k) - \frac{L(\lambda_k - \beta_k)^2}{k\Gamma_{k+1}\alpha_{k+1}}\Big]\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 \\ &+ \underbrace{\sum_{k=0}^{K-1}\frac{\lambda_k}{2}\mathbb{E}\Big[\sum_{i=(\tau(k)-1)q}^{k-1}\frac{L^2}{|\xi_i|}\Big[2\beta_i^2\|G_{\lambda_i}(x_i, v_i)\|^2 + 2\alpha_{i+2}^2\Gamma_{i+1}\sum_{t=0}^{i}\frac{(\lambda_t - \beta_t)^2}{\alpha_{t+1}\Gamma_{t+1}}\|G_{\lambda_t}(x_t, v_t)\|^2\Big]\Big]}_{T}. \end{aligned} \tag{19}$$

Before we proceed the proof, we first specify the choices of all the parameters. Specifically, we choose a constant mini-batch size $|\xi_k| \equiv |\xi|$, a constant $q = |\xi|$, a constant $\beta_k \equiv \beta > 0$, $\lambda_k \in [\beta, (1 + \alpha_{k+1})\beta]$. Based on these parameter settings, the term $T$ in the above inequality can be bounded as follows.

$$\begin{aligned} T &\overset{(i)}{\leq} \sum_{k=0}^{K-1}\frac{\lambda_k}{2}\mathbb{E}\Big[\sum_{i=(\tau(k)-1)q}^{\tau(k)q-1}\frac{L^2}{|\xi_i|}\Big[2\beta_i^2\|G_{\lambda_i}(x_i, v_i)\|^2 + 2\alpha_{i+2}^2\Gamma_{i+1}\sum_{t=0}^{k-1}\frac{(\lambda_t - \beta_t)^2}{\alpha_{t+1}\Gamma_{t+1}}\|G_{\lambda_t}(x_t, v_t)\|^2\Big]\Big] \\ &\overset{(ii)}{\leq} \sum_{k=0}^{K-1}\frac{\lambda_k L^2 q\beta^2}{|\xi|}\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + \sum_{k=0}^{K-1}\frac{2\lambda_k L^2}{|\xi|[(\tau(k)-1)q+1]^3}\sum_{t=0}^{k-1}\frac{(\lambda_t - \beta_t)^2}{\alpha_{t+1}\Gamma_{t+1}}\mathbb{E}\|G_{\lambda_t}(x_t, v_t)\|^2 \\ &\overset{(iii)}{\leq} \sum_{k=0}^{K-1}\lambda_k L^2\beta^2\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + \frac{2L^2\beta^2}{|\xi|}\sum_{k=0}^{K-1}\frac{\alpha_{k+1}}{\Gamma_{k+1}}\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2\Big(\sum_{t=k}^{K-1}\frac{\lambda_k}{[(\tau(t)-1)q+1]^3}\Big) \\ &\overset{(iv)}{\leq} \sum_{k=0}^{K-1}\lambda_k L^2\beta^2\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + \frac{4L^2\beta^3}{|\xi|}\sum_{k=0}^{K-1}(k+1)\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2\Big(\sum_{t=(\tau(k)-1)q}^{\tau(K)q}\frac{1}{[(\tau(t)-1)q+1]^3}\Big) \\ &= \sum_{k=0}^{K-1}\lambda_k L^2\beta^2\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + \frac{4L^2\beta^3}{|\xi|}\sum_{k=0}^{K-1}(k+1)\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2\Big(\sum_{t=\tau(k)-1}^{\tau(K)}\frac{q}{(tq+1)^3}\Big) \\ &\leq \sum_{k=0}^{K-1}\lambda_k L^2\beta^2\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + \frac{2L^2\beta^3}{q}\sum_{k=0}^{K-1}(k+1)\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2\frac{1}{[(\tau(k)-1)q+1]^2} \\ &\overset{(v)}{\leq} \sum_{k=0}^{K-1}\lambda_k L^2\beta^2\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + 2L^2\beta^3\sum_{k=0}^{K-1}\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2\frac{\tau(k)}{[(\tau(k)-1)q+1]^2} \\ &\leq \sum_{k=0}^{K-1}\lambda_k L^2\beta^2\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + 2L^2\beta^3\sum_{k=0}^{K-1}\mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2, \end{aligned} \tag{20}$$

where (i) follows from the facts that $i \leq k-1$ and $k-1 \leq \tau(k)q - 1$, (ii) uses the fact that $\sum_{i=(\tau(k)-1)q}^{\tau(k)q-1}\alpha_{i+2}^2\Gamma_{i+1} \leq \frac{2}{(\tau(k)-1)q+1)^3}$, (iii) uses the parameter settings $q = |\xi|$ and $\lambda_t - \beta_t \leq \alpha_t\beta$, (iv) uses the facts that $\lambda_k \leq 2\beta$ and $(\tau(k)-1)q \leq k \leq \tau(k)q$ and (v) uses the fact that $k \leq \tau(k)q - 1$. Substituting the above inequality into eq. (19) and

simplifying, we obtain that

$$\mathbb{E}[F(x_K)] \leq F(x_0) - \sum_{k=0}^{K-1} \Big[\lambda_k(\frac{1}{2} - L\lambda_k - L^2\beta^2) - \frac{L(\lambda_k-\beta_k)^2}{k\Gamma_{k+1}\alpha_{k+1}} - 2L^2\beta^3\Big]\mathbb{E}\|G_{\lambda_k}(x_k,v_k)\|^2 \tag{21}$$

$$\leq F(x_0) - \sum_{k=0}^{K-1} \Big[\beta(\frac{1}{2} - 2L\beta - L^2\beta^2) - L\beta^2 - 2L^2\beta^3\Big]\mathbb{E}\|G_{\lambda_k}(x_k,v_k)\|^2. \tag{22}$$

Choosing $\beta \leq \frac{1}{8L}$, the above inequality further implies that

$$\mathbb{E}[F(x_K)] \leq F(x_0) - \sum_{k=0}^{K-1} \frac{\beta}{16}\mathbb{E}\|G_{\lambda_k}(x_k,v_k)\|^2. \tag{23}$$

Then, it follows that $\frac{1}{K}\sum_{k=0}^{K-1}\mathbb{E}\|G_{\lambda_k}(x_k,v_k)\|^2 \leq 16(F(x_0)-F^*)/(K\beta)$. Next, we bound the term $\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,\nabla f(z_\zeta))\|^2$, where $\zeta$ is selected uniformly at random from $\{0,\ldots,K-1\}$. Observe that

$$\begin{aligned}
\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,\nabla f(z_\zeta))\|^2 &= \mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,\nabla f(z_\zeta)) - G_{\lambda_\zeta}(z_\zeta,v_\zeta) + G_{\lambda_\zeta}(z_\zeta,v_\zeta)\|^2 \\
&\leq 2\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,\nabla f(z_\zeta)) - G_{\lambda_\zeta}(z_\zeta,v_\zeta)\|^2 + 2\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,v_\zeta)\|^2 \\
&\overset{(i)}{\leq} 2\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2 + 2\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,v_\zeta) - G_{\lambda_\zeta}(x_\zeta,v_\zeta) + G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2 \\
&\leq 2\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2 + 4\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,v_\zeta) - G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2 + 4\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2 \\
&\leq 2\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2 + 4\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2 \\
&\quad + \frac{4}{\lambda_\zeta^2}\mathbb{E}\|z_\zeta - x_\zeta + \mathrm{prox}_{\lambda_\zeta g}(x_\zeta - \lambda_\zeta v_\zeta) - \mathrm{prox}_{\lambda_\zeta g}(z_\zeta - \lambda_\zeta v_\zeta)\|^2 \\
&\leq 2\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2 + 4\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2 \\
&\quad + \frac{8}{\lambda_\zeta^2}\mathbb{E}\|z_\zeta - x_\zeta\|^2 + \frac{8}{\lambda_\zeta^2}\mathbb{E}\|\mathrm{prox}_{\lambda_\zeta g}(x_\zeta - \lambda_\zeta v_\zeta) - \mathrm{prox}_{\lambda_\zeta g}(z_\zeta - \lambda_\zeta v_\zeta)\|^2 \\
&\overset{(ii)}{\leq} 2\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2 + 4\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2 + \frac{8}{\lambda_\zeta^2}\mathbb{E}\|z_\zeta - x_\zeta\|^2 + \frac{8}{\lambda_\zeta^2}\mathbb{E}\|z_\zeta - x_\zeta\|^2 \\
&\overset{(iii)}{\leq} 2\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2 + 4\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2 + \frac{16}{\lambda_\zeta^2}\mathbb{E}\|y_\zeta - x_\zeta\|^2
\end{aligned} \tag{24}$$

where (i) uses the non-expansiveness property of the operator $G$ in Lemma 4, (ii) follows from the non-expansiveness of the proximal operator, and (iii) uses the update rule and the fact that $0 < \alpha_k < 1$.

Next, we bound the three terms on the right hand side of the above inequality separately. First, note that

$$\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2 = \frac{1}{K}\sum_{k=0}^{K-1}\mathbb{E}\|G_{\lambda_k}(x_k,v_k)\|^2 \leq \frac{16(F(x_0)-F^*)}{K\beta}.$$

Second, note that eq. (18) implies that

$$\begin{aligned}
\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2 &\leq \mathbb{E}\sum_{i=(\tau(\zeta)-1)q}^{\zeta-1}\frac{L^2}{|\xi_i|}\big[2\beta_i^2\|G_{\lambda_i}(x_i,v_i)\|^2 + 2\alpha_{i+2}^2\Gamma_{i+1}\sum_{t=0}^{i}\frac{(\lambda_t-\beta_t)^2}{\alpha_{t+1}\Gamma_{t+1}}\|G_{\lambda_t}(x_t,v_t)\|^2\big] \\
&\leq \frac{2L^2\beta^2}{|\xi|}\mathbb{E}\Big(\sum_{i=(\tau(\zeta)-1)q}^{\tau(\zeta)q-1}\|G_{\lambda_i}(x_i,v_i)\|^2\Big) + \frac{L^2}{|\xi|}\mathbb{E}\Big(\sum_{i=(\tau(\zeta)-1)q}^{\zeta-1}2\alpha_{i+2}^2\Gamma_{i+1}\sum_{t=0}^{i}\frac{(\lambda_t-\beta_t)^2}{\alpha_{t+1}\Gamma_{t+1}}\|G_{\lambda_t}(x_t,v_t)\|^2\Big) \\
&\leq \frac{2L^2\beta^2}{|\xi|}\frac{1}{K}\sum_{\zeta=0}^{K-1}\Big(\sum_{i=(\tau(\zeta)-1)q}^{\tau(\zeta)q-1}\mathbb{E}\|G_{\lambda_i}(x_i,v_i)\|^2\Big)
\end{aligned}$$

$$
\begin{aligned}
&+\frac{L^2\beta^2}{|\xi|}\frac{1}{K}\sum_{\zeta=0}^{K-1}\Bigg(\sum_{i=(\tau(\zeta)-1)q}^{\tau(\zeta)q-1}2\alpha_{i+2}^2\Gamma_{i+1}\sum_{t=0}^{\zeta-1}(t+1)\mathbb{E}\|G_{\lambda_t}(x_t,v_t)\|^2\Bigg)\\
&\le\frac{2L^2\beta^2 q}{|\xi|}\frac{1}{K}\sum_{\zeta=0}^{K-1}\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2+\frac{L^2\beta^2}{|\xi|}\frac{1}{K}\sum_{\zeta=0}^{K-1}\Bigg(\frac{4}{[(\tau(\zeta)-1)q+1]^3}\sum_{t=0}^{\zeta-1}(t+1)\mathbb{E}\|G_{\lambda_t}(x_t,v_t)\|^2\Bigg)\\
&\le 2L^2\beta^2\Bigg(\frac{1}{K}\sum_{\zeta=0}^{K-1}\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2\Bigg)+\frac{L^2\beta^2}{|\xi|}\frac{1}{K}\sum_{\zeta=0}^{K-1}(\zeta+1)\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2\sum_{t=\zeta}^{K-1}\frac{4}{[\tau(t)-1)q+1]^3}\\
&\le 2L^2\beta^2\Bigg(\frac{1}{K}\sum_{\zeta=0}^{K-1}\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2\Bigg)+L^2\beta^2\frac{1}{K}\sum_{\zeta=0}^{K-1}\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2\frac{2\tau(\zeta)}{[\tau(\zeta)-1)q+1]^2}\\
&\le 3L^2\beta\frac{16(F(x_0)-F^*)}{K},
\end{aligned}
$$

where we have used the fact that $\zeta$ is sampled uniformly from $0,...,K-1$ at random.

Third, note that by item 2 of Lemma 3, we know that

$$
\begin{aligned}
\mathbb{E}\|y_\zeta-x_\zeta\|^2&\le\mathbb{E}\Bigg(\Gamma_\zeta\sum_{t=0}^{\zeta-1}\frac{\lambda_t-\beta_t}{\alpha_{t+1}\Gamma_{t+1}}\|G_{\lambda_t}(x_t,v_t)\|^2\Bigg)\\
&\le\frac{1}{K}\sum_{\zeta=0}^{K-1}\Gamma_\zeta\sum_{t=0}^{\zeta-1}\frac{\lambda_t-\beta_t}{\alpha_{t+1}\Gamma_{t+1}}\mathbb{E}\|G_{\lambda_t}(x_t,v_t)\|^2\\
&\le\frac{\beta^2}{K}\sum_{\zeta=0}^{K-1}\Gamma_\zeta\sum_{t=0}^{\zeta-1}(t+1)\mathbb{E}\|G_{\lambda_t}(x_t,v_t)\|^2\\
&=\frac{\beta^2}{K}\sum_{\zeta=0}^{K-1}(\zeta+1)\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2\Bigg(\sum_{t=\zeta+1}^{K-1}\Gamma_t\Bigg)\\
&\le\frac{\beta^2}{K}\sum_{\zeta=0}^{K-1}\mathbb{E}\|G_{\lambda_\zeta}(x_\zeta,v_\zeta)\|^2.
\end{aligned}
\tag{25}
$$

Combining the above three inequalities and note that $L\beta=\Theta(1)$ and $\frac{\beta}{\lambda_k}\le 1$, we finally obtain that

$$
\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,\nabla f(z_\zeta))\|^2\le\Theta\Bigg(\frac{L(F(x_0)-F^*)}{K}\Bigg).
\tag{26}
$$

This further implies that

$$
\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,\nabla f(z_\zeta))\|\le\sqrt{\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta,\nabla f(z_\zeta))\|^2}\le\Theta\Bigg(\sqrt{\frac{L(F(x_0)-F^*)}{K}}\Bigg).
$$

Setting the right hand side of the above inequality to be bounded by $\epsilon$, we obtain that $K\ge\Theta\Big(\frac{L(F(x_0)-F^*)}{\epsilon^2}\Big)$. Then, the total number of stochastic gradient calls is bounded by $(K+q)\frac{n}{q}+K|\xi|\le\Theta(n+\sqrt{n}\epsilon^{-2})$.

## C. Proof of Theorem 2

The convergence proof of Proximal SPIDER-M with both epochwise-diminishing momentum and epochwise restart momentum follow from that of Proximal SPIDER-M, and therefore we only describe the key steps to adapt the proof.

We first prove the result of Proximal SPIDER-M with epochwise-diminishing momentum. Under the epochwise-diminishing momentum scheme, the momentum coefficient is set to be $\alpha_k=\frac{2}{\lceil k/q\rceil+1}$. Consequently, we have $\Gamma_k=\frac{2}{\lceil k/q\rceil(\lceil k/q\rceil+1)}$.

First, one can check that eq. (16) still holds, and now we have $\sum_{t=k}^{K-1}\Gamma_t \le \frac{2}{\lceil k/q\rceil}$. Then, we follow the steps that bound the accumulation error term $T$ in eq. (19). In the derivation of (ii), we now have that $\sum_{i=(\tau(k)-1)q}^{\tau(k)q-1}\alpha_{i+2}^2\Gamma_{i+1} \le \frac{2}{\tau(k)^3}$. Substituting this new bound into (ii) and noting that in (iii) we now have $\frac{\alpha_{k+1}}{\Gamma_{k+1}} = (\lceil k/q\rceil + 1)$, one can follow the subsequent steps and show that the upper bound for T in eq. (20) still holds. Moreover, in eq. (21) we should replace $\frac{L(\lambda_k-\beta_k)^2}{k\Gamma_{k+1}\alpha_{k+1}}$ with $\frac{L(\lambda_k-\beta_k)^2}{\lceil k/q\rceil\Gamma_{k+1}\alpha_{k+1}}$, and consequently eq. (22) is still valid. Then, one can follow the same analysis and show that eq. (23) is still valid. In summary, the convergence rate and the corresponding oracle complexity remain in the same order.

The convergence proof of Proximal SPIDER-M with periodic restart follows from that of Proximal SPIDER-M. The core idea is to apply the result of Proximal SPIDER-M to each restart period. Specifically, consider the iterations $k = 0, 1, ..., q-2$. As no restart is performed within these iterations, we can apply the result in eq. (26) (note that $F^*$ is the relaxation of $F(x_K)$) obtained from the analysis of Proximal SPIDER-M and conclude that

$$\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\|^2 \le \Theta\left(\frac{L(F(x_0) - \mathbb{E}[F(x_{q-1})])}{q-1}\right), \text{ where } \zeta \overset{\text{Unif}}{\sim} \{0, ..., q-2\}. \tag{27}$$

Due to the periodic restart, the above bound also holds similarly for the iterations $k = tq, tq+1, ..., (t+1)q-2$ for any $t \in \mathbb{N}$, which yields that

$$\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\|^2 \le \Theta\left(\frac{L(F(x_{tq}) - \mathbb{E}[F(x_{(t+1)q-1})])}{q-1}\right), \text{ where } \zeta \overset{\text{Unif}}{\sim} \{tq, ..., (t+1)q-2\}. \tag{28}$$

Next, consider running the algorithm with restart for iterations $k = 0, ..., K-1$, and the output index $\zeta$ is selected from $\{k : 0 \le k \le K-1, \text{mod}(k, q-1) \ne 0\}$ uniformly at random. Let $T = \left\lceil \frac{K}{q-1} \right\rceil$. Then, we can obtain the following estimate

$$\begin{aligned}
\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\|^2 &\le \frac{1}{K-T}\sum_{t=0}^{T}\sum_{k=tq}^{(t+1)q-2}\mathbb{E}\|G_{\lambda_k}(z_k, \nabla f(z_k))\|^2 \\
&\overset{(i)}{\le} \Theta\left(\frac{1}{K-T}\sum_{t=0}^{T} L\mathbb{E}(F(x_{tq}) - F(x_{(t+1)q-1}))\right) \\
&\overset{(ii)}{\le} \Theta\left(\frac{L(F(x_0) - F^*)}{K}\right),
\end{aligned}$$

where (i) uses the results inductively derived from eq. (28) and (ii) uses the fact that $x_{(t+1)q-1} = x_{(t+1)q}$ due to restart.

Therefore, it follows that $\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\| \le \epsilon$ whenever $K \ge \Theta\left(\frac{L(F(x_0)-F^*)}{\epsilon^2}\right)$, and the total number of stochastic gradient calls is in the order of $\Theta(n + \sqrt{n}\epsilon^{-2})$.

## D. Proof of Theorem 3

Note that when $\text{mod}(k, q) = 0$, the Online Proximal SPIDER-M samples $\xi_1$ data points to estimate the gradient, and we obtain the following variance bound based on Assumption 2.

$$\mathbb{E}\|v_k - \nabla f(x_k)\|^2 = \mathbb{E}\left\|\frac{1}{|\xi_1|}\sum_{i=1}^{|\xi_1|}\nabla \ell_{u_i}(x_k) - \nabla f(x_k)\right\|^2 \le \frac{1}{|\xi_1|^2}\sum_{i=1}^{|\xi_1|}\mathbb{E}\left\|\nabla \ell_{u_i}(x_k) - \nabla f(x_k)\right\|^2 \le \frac{\sigma^2}{|\xi_1|}. \tag{29}$$

By telescoping Lemma 1 and using the above bound, we obtain the following lemma.

**Lemma 5.** *Under Assumptions 1 and 2, the estimation of gradient $v_k$ constructed by Online Proximal SPIDER-M satisfies that for all $k \in \mathbb{N}$,*

$$\mathbb{E}\|v_k - \nabla f(z_k)\|^2 \le \sum_{i=(\tau(k)-1)q}^{k-1}\frac{L^2}{|\xi_i|}\mathbb{E}\|z_{i+1} - z_i\|^2 + \frac{\sigma^2}{|\xi_1|}. \tag{30}$$

Then, one can check that eq. (17) still holds in the online case. Therefore, we can apply the bound in Lemma 5 to eq. (17) and follow the proof of eq. (23). One can check that there is an additional term $\sum_{k=0}^{K-1} \frac{\lambda_k \sigma^2}{2|\xi_1|}$ in the online case, and we obtain the following bound.

$$\begin{aligned}\mathbb{E}[F(x_K)] &\leq F(x_0) - \sum_{k=0}^{K-1} \frac{\beta}{16} \mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + \sum_{k=0}^{K-1} \frac{\lambda_k \sigma^2}{2|\xi_1|} \\ &\leq F(x_0) - \sum_{k=0}^{K-1} \frac{\beta}{16} \mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 + \frac{K\beta\sigma^2}{|\xi_1|}. \end{aligned} \tag{31}$$

Then, it follows that $\frac{1}{K}\sum_{k=0}^{K-1} \mathbb{E}\|G_{\lambda_k}(x_k, v_k)\|^2 \leq 16(F(x_0) - F^*)/(K\beta) + \frac{16\sigma^2}{|\xi_1|}$. One can check that eq. (24) still holds, and we only need to update the bound for the term $\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2$ as follows

$$\mathbb{E}\|\nabla f(z_\zeta) - v_\zeta\|^2 \leq 3L^2\beta \frac{16(F(x_0) - F^*)}{K} + \frac{\sigma^2}{|\xi_1|}. \tag{32}$$

Then, we finally obtain that

$$\mathbb{E}\|G_{\lambda_\zeta}(z_\zeta, \nabla f(z_\zeta))\|^2 \leq \Theta\left(\frac{L(F(x_0) - F^*)}{K} + \frac{\sigma^2}{|\xi_1|}\right). \tag{33}$$

To make the right hand side be smaller than $\epsilon^2$, we can set $K \geq \frac{2L(F(x_0)-F^*)}{\epsilon^2}$ and $|\xi_1| \geq \frac{2\sigma^2}{\epsilon^2}$. This proves the desired iteration complexity. On the other hand, the total number of stochastic gradient oracle calls is at most $(K+q)\frac{|\xi_1|}{q} + K|\xi_2|$. By setting $q = |\xi_2| = \sqrt{|\xi_1|}$, we obtain the total oracle complexity as $\Theta(\epsilon^{-3})$.

## E. Proof of Theorem 4

The proof follows exactly from that of Theorem 2 (the same treatment of the momentum schemes apply).