
A Simple Stochastic Variance Reduced Algorithm with Fast Convergence Rates

Kaiwen Zhou¹ Fanhua Shang² James Cheng¹

Abstract

Recent years have witnessed exciting progress in the study of stochastic variance reduced gradient methods (e.g., SVRG, SAGA), their accelerated variants (e.g., Katyusha) and their extensions in many different settings (e.g., online, sparse, asynchronous, distributed). Among them, accelerated methods enjoy improved convergence rates but have complex coupling structures, which makes them hard to be extended to more settings (e.g., sparse and asynchronous) due to the existence of perturbation. In this paper, we introduce a simple stochastic variance reduced algorithm (MiG), which enjoys the best-known convergence rates for both strongly convex and non-strongly convex problems. Moreover, we also present its efficient sparse and asynchronous variants, and theoretically analyze its convergence rates in these settings. Finally, extensive experiments for various machine learning problems such as logistic regression are given to illustrate the practical improvement in both serial and asynchronous settings.

1. Introduction

In this paper, we consider the following convex optimization problem with a finite-sum structure, which is prevalent in machine learning and statistics such as regularized empirical risk minimization (ERM):

$$\min_{x \in \mathbb{R}^d} \left\{ F(x) \triangleq f(x) + g(x) \right\}, \quad (1)$$

where $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$ is a finite average of n smooth convex function $f_i(x)$, and $g(x)$ is a relatively simple (but possibly non-differentiable) convex function.

For the strongly convex problem (1), traditional gradient descent (GD) yields a linear convergence rate but with a

high per-iteration cost. As an alternative, SGD (Robbins & Monro, 1951) enjoys significantly lower per-iteration complexity than GD, i.e., $O(d)$ vs. $O(nd)$. However, due to the variance of random sampling, standard SGD usually obtains slow convergence and poor performance (Johnson & Zhang, 2013). Recently, many stochastic variance reduced methods (e.g., SAG (Roux et al., 2012), SDCA (Shalev-Shwartz & Zhang, 2013), SVRG (Johnson & Zhang, 2013), SAGA (Defazio et al., 2014), and their proximal variants, such as (Schmidt et al., 2017), (Shalev-Shwartz & Zhang, 2016), (Xiao & Zhang, 2014) and (Konečný et al., 2016)) have been proposed to solve Problem (1). All these methods enjoy low per-iteration complexities comparable with SGD, but with the help of certain variance reduction techniques, they obtain a linear convergence rate as GD. More accurately, these methods achieve an improved oracle complexity $\mathcal{O}((n+\kappa) \log(1/\epsilon))^1$, compared with $\mathcal{O}(n\sqrt{\kappa} \log(1/\epsilon))$ for accelerated deterministic methods (e.g., Nesterov’s accelerated gradient descent (Nesterov, 2004)). In summary, these methods dramatically reduce the overall computational cost compared with deterministic methods in theory.

More recently, researchers have proposed accelerated stochastic variance reduced methods for Problem (1), which include Acc-Prox-SVRG (Nitanda, 2014), APCG (Lin et al., 2014), Catalyst (Lin et al., 2015), SPDC (Zhang & Xiao, 2015), point-SAGA (Defazio, 2016), and Katyusha (Allen-Zhu, 2017). For strongly convex problems, both Acc-Prox-SVRG (Nitanda, 2014) and Catalyst (Lin et al., 2015) make good use of the “Nesterov’s momentum” in (Nesterov, 2004) and attain the corresponding oracle complexities $\mathcal{O}((n+b\sqrt{\kappa}) \log(1/\epsilon))$ (with a sufficiently large mini-batch size b) and $\mathcal{O}((n+\sqrt{\kappa n}) \log(\kappa) \log(1/\epsilon))$. APCG, SPDC, point-SAGA and Katyusha essentially achieve the best-known oracle complexity $\mathcal{O}((n+\sqrt{\kappa n}) \log(1/\epsilon))$.

Inspired by emerging multi-core computer architectures, asynchronous variants of the above stochastic gradient methods have been proposed in recent years, e.g., Hogwild! (Recht et al., 2011), Lock-Free SVRG (Reddi et al., 2015), KroMagnon (Mania et al., 2017) and ASAGA (Leblond et al., 2017). Among them, KroMagnon and ASAGA (as the sparse and asynchronous variants of SVRG and SAGA)

¹Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong ²School of Artificial Intelligence, Xidian University, China. Correspondence to: Fanhua Shang <fhshang@xidian.edu.cn>.

¹We denote $\kappa \triangleq \frac{L}{\sigma}$ throughout the paper, known as the condition number of an L -smooth and σ -strongly convex function.

enjoy a fast linear convergence rate for strongly convex objectives. However, there still lacks a variant of accelerated algorithms in these settings.

The main issue for those accelerated algorithms is that most of their algorithm designs (e.g., (Allen-Zhu, 2017) and (Hien et al., 2017)) involve tracking at least two highly correlated coupling vectors² (in the inner loop). This kind of algorithm structure prevents us from deriving efficient (lock-free) asynchronous sparse variants for those algorithms. More critically, when the number of concurrent threads is large (e.g., 20 threads), the high perturbation (i.e., updates on shared variables from concurrent threads) may even destroy their convergence guarantees. This leads us to the key question we study in this paper:

Can we design an accelerated algorithm that keeps track of only one variable vector?

We answer this question by a simple stochastic variance reduced algorithm (MiG), which has the following features:

- **Simple.** The algorithm construction of MiG requires tracking only one variable vector in the inner loop, which means its computational overhead and memory overhead are exactly the same as SVRG (or Prox-SVRG (Xiao & Zhang, 2014)). This feature allows MiG to be extended to more strict settings such as the sparse and asynchronous settings. We theoretically analyze its variants in Section 4.
- **Theoretically Fast.** MiG achieves the best-known oracle complexity of $\mathcal{O}((n + \sqrt{\kappa n}) \log(1/\epsilon))$ for strongly convex problems. For non-strongly convex problems, MiG also achieves an optimal convergence rate $\mathcal{O}(1/T^2)$, where T is the total number of stochastic iterations. These rates keep up with those of Katyusha and are consistently faster than non-accelerated algorithms, e.g., SVRG and SAGA.
- **Practically Fast.** Due to its light-weighted algorithm structure, our experiments verify that the running time of MiG is shorter than its counterparts in the serial dense setting. In the sparse and asynchronous settings, MiG achieves significantly better performance than KroMagnon and ASAGA in terms of both gradient evaluations and running time.
- **Implementable.** Unlike many incremental gradient methods (e.g., SAGA), MiG does not require an additional gradient table which is not practical for large-scale problems. Our algorithm layout is similar to SVRG, which means that most existing techniques designed for SVRG (such as a distributed variant) can be modified for MiG without much effort.

²Here we refer to the number of variable vectors involved in one update.

Table 1. Comparison of different stochastic variance reduced algorithms. (“Complexity” is for strongly-convex problems. “Memory” is those used to store variables. “S&A” refers to efficient (lock-free) Sparse & Asynchronous variant.)

Algorithm	Complexity	Memory	S&A
SVRG	$\mathcal{O}((n + \kappa) \log \frac{1}{\epsilon})$	1 Vec.	✓
SAGA	$\mathcal{O}((n + \kappa) \log \frac{1}{\epsilon})$	1 Vec. 1 ∇ Table.	✓
Katyusha	$\mathcal{O}((n + \sqrt{\kappa n}) \log \frac{1}{\epsilon})$	2 Vec.	×
MiG	$\mathcal{O}((n + \sqrt{\kappa n}) \log \frac{1}{\epsilon})$	1 Vec.	✓

We summarize some properties of the existing methods and MiG in Table 1.

2. Notations

We mainly consider Problem (1) in standard Euclidean space with the Euclidean norm denoted by $\|\cdot\|$. We use \mathbb{E} to denote that the expectation is taken with respect to all randomness in one epoch. To further categorize the objective functions, we define that a convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is said to be L -smooth if for all $x, y \in \mathbb{R}^d$, it holds that

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

and σ -strongly convex if for all $x, y \in \mathbb{R}^d$,

$$f(x) \geq f(y) + \langle \mathcal{G}, x - y \rangle + \frac{\sigma}{2} \|x - y\|^2, \quad (3)$$

where $\mathcal{G} \in \partial f(y)$, the set of sub-gradient of f at y . If f is differentiable, we replace $\mathcal{G} \in \partial f(y)$ with $\mathcal{G} = \nabla f(y)$. Then we make the following assumptions to categorize Problem (1):

Assumption 1 (Strongly Convex). *In Problem (1), each $f_i(\cdot)$ ³ is L -smooth and convex, $g(\cdot)$ is σ -strongly convex.*

Assumption 2 (Non-strongly Convex). *In Problem (1), each $f_i(\cdot)$ is L -smooth and convex, and $g(\cdot)$ is convex.*

3. A Simple Accelerated Algorithm

In this section, we introduce a simple accelerated stochastic algorithm (MiG) for both strongly convex and non-strongly convex problems.

3.1. MiG for Strongly Convex Objectives

We first consider Problem (1) that satisfies Assumption 1.

Now we formally introduce MiG in Algorithm 1. In order to

³In fact, if each $f_i(\cdot)$ is L -smooth, the averaged function $f(\cdot)$ is itself L -smooth — but probably with a smaller L . We keep using L as the smoothness constant for a consistent analysis.

Algorithm 1 MiG

1: **Input:** Initial vector x_0 , epoch length m , learning rate η , parameter θ .

2: $\tilde{x}_0 = x_0^1 = x_0, \omega = 1 + \eta\sigma$;

3: **for** $s = 1 \dots \mathcal{S}$ **do**

4: $\mu_s = \nabla f(\tilde{x}_{s-1})$;

5: **for** $j = 1 \dots m$ **do**

6: Sample i_j uniformly in $\{1 \dots n\}$;

7: $y_{j-1} = \theta x_{j-1}^s + (1-\theta)\tilde{x}_{s-1}$; //temp variable y

8: $\tilde{\nabla} = \nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1}) + \mu_s$;

9: $x_j^s = \arg \min_x \left\{ \frac{1}{2\eta} \|x - x_{j-1}^s\|^2 + \langle \tilde{\nabla}, x \rangle + g(x) \right\}$;

10: **end for**

11: $\tilde{x}_s = \theta \left(\sum_{j=0}^{m-1} \omega^j \right)^{-1} \sum_{j=0}^{m-1} \omega^j x_{j+1}^s + (1-\theta)\tilde{x}_{s-1}$;

12: $x_0^{s+1} = x_m^s$;

13: **end for**

14: **return** $\tilde{x}_{\mathcal{S}}$.

further illustrate some ideas behind the algorithm structure, we make the following remarks:

- *Temp variable y .* As we can see in Algorithm 1, y is a convex combination of x and \tilde{x} with the parameter θ . So for implementation, we do not need to keep track of y in the whole inner loop. For the purpose of giving a clean proof, we mark y with iteration number j .
- *Fancy update for \tilde{x}_s .* One can easily verify that this update for \tilde{x}_s is equivalent to using ω^j weighted averaged y_{j+1} to update \tilde{x}_s , which is written as: $\tilde{x}_s = \left(\sum_{j=0}^{m-1} \omega^j \right)^{-1} \sum_{j=0}^{m-1} \omega^j y_{j+1}$. Since we only keep track of x , we adopt this expended fancy update for \tilde{x}_s — but it is still quite simple in implementation.
- *Choice of x_0^{s+1} .* In recent years, some existing stochastic algorithms such as (Zhang et al., 2013; Xiao & Zhang, 2014) choose to use \tilde{x}_s as the initial vector for new epoch. For MiG, when using \tilde{x}_s , the overall oracle complexity will degenerate to a non-accelerated one for some ill-conditioned problems, which is $\mathcal{O}((n+\kappa) \log(1/\epsilon))$. It is reported that even in practice, using the last iterate yields a better performance as discussed in (Allen-Zhu & Hazan, 2016b).

Next we give the convergence rate of MiG in terms of oracle complexity as follows (the proofs to all theorems in this paper are given in the Supplementary Material):

Theorem 1 (Strongly Convex). *Let x^* be the optimal solution of Problem (1). If Assumption 1 holds, then by choosing $m = \Theta(n)$, MiG achieves an ϵ -additive error with the following oracle complexities in expectation:*

$$\begin{cases} \mathcal{O}(\sqrt{\kappa n} \log \frac{F(x_0) - F(x^*)}{\epsilon}), & \text{if } \frac{m}{\kappa} \leq \frac{3}{4}, \\ \mathcal{O}(n \log \frac{F(x_0) - F(x^*)}{\epsilon}), & \text{if } \frac{m}{\kappa} > \frac{3}{4}. \end{cases}$$

 Table 2. The theoretical settings of the parameters η and θ .

CONDITION	LEARNING RATE η	PARAMETER θ
$\frac{m}{\kappa} \leq \frac{3}{4}$	$\sqrt{\frac{1}{3\sigma m L}}$	$\sqrt{\frac{m}{3\kappa}}$
$\frac{m}{\kappa} > \frac{3}{4}$	$\frac{2}{3L}$	$\frac{1}{2}$

In other words, the overall oracle complexity of MiG is $\mathcal{O}((n + \sqrt{\kappa n}) \log \frac{F(x_0) - F(x^*)}{\epsilon})$.

This result implies that in the strongly convex setting, MiG enjoys the best-known oracle complexity for stochastic first-order algorithms, e.g., APCG, SPDC, and Katyusha. The theoretical suggestions⁴ of the learning rate η and the parameter θ are shown in Table 2.

3.1.1. COMPARISON BETWEEN MiG AND RELATED METHODS

We carefully compare the algorithm structure of MiG with Katyusha, and find that MiG corresponds to a case of Katyusha, when $1 - \tau_1 - \tau_2 = 0$, and Option II in Katyusha is used. However, this setting is neither suggested nor analyzed in (Allen-Zhu, 2017), and thus without a convergence guarantee. In some sense, MiG can be regarded as a “simplified Katyusha”, while this simplification does not hurt its oracle complexity. Since this simplification discards all the proximal gradient steps in Katyusha, MiG enjoys a lower memory overhead in practice and a cleaner proof in theory. Detailed comparison with Katyusha can be found in the Supplementary Material B.1.1.

MiG does not use any kind of “Nesterov’s Momentum”, which is used in some accelerated algorithms, e.g., Acc-Prox-SVRG (Nitanda, 2014) and Catalyst (Lin et al., 2015).

3.2. MiG for Non-strongly Convex Objectives

In this part, we consider Problem (1) when Assumption 2 holds. Since non-strongly convex optimization problems (e.g., LASSO) are becoming prevalent these days, making a direct variant of MiG for these problems is of interest.

In this setting, we summarize MiG^{NSC} with the optimal $\mathcal{O}(1/T^2)$ convergence rate in Algorithm 2.

Theorem 2 (Non-strongly Convex). *If Assumption 2 holds, then by choosing $m = \Theta(n)$, MiG^{NSC} achieves the following oracle complexity in expectation:*

$$\mathcal{O}\left(n \sqrt{\frac{F(x_0) - F(x^*)}{\epsilon}} + \sqrt{\frac{nL \|x_0 - x^*\|^2}{\epsilon}}\right).$$

This result implies that MiG^{NSC} attains the optimal con-

⁴We recommend users to tune these two parameters for better performance in practice, or to use the tuning criteria mentioned in Table 3 with only tuning θ .

Algorithm 2 MiG^{NSC}

```

1: Input: Initial vector  $x_0$ , epoch length  $m$ , learning rate
    $\eta$ , parameter  $\theta$ .
2:  $\tilde{x}_0 = x_0^1 = x_0$ ;
3: for  $s = 1 \dots \mathcal{S}$  do
4:    $\mu_s = \nabla f(\tilde{x}_{s-1})$ ,  $\theta = \frac{2}{s+4}$ ,  $\eta = \frac{1}{4L\theta}$ ;
5:   for  $j = 1 \dots m$  do
6:     Sample  $i_j$  uniformly in  $\{1 \dots n\}$ ;
7:      $y_{j-1} = \theta x_{j-1}^s + (1-\theta)\tilde{x}_{s-1}$ ; //temp variable  $y$ 
8:      $\tilde{\nabla} = \nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1}) + \mu_s$ ;
9:      $x_j^s = \arg \min_x \left\{ \frac{1}{2\eta} \|x - x_{j-1}^s\|^2 + \langle \tilde{\nabla}, x \rangle + g(x) \right\}$ ;
10:   end for
11:    $\tilde{x}_s = \frac{\theta}{m} \sum_{j=1}^m x_j^s + (1-\theta)\tilde{x}_{s-1}$ ;
12:    $x_0^{s+1} = x_m^s$ ;
13: end for
14: return  $\tilde{x}_S$ .

```

vergence rate $\mathcal{O}(1/T^2)$, where $T = \mathcal{S}(m+n)$ is the total number of stochastic iterations.

The result in Theorem 2 shows that MiG^{NSC} enjoys the same oracle complexity as Katyusha^{ns} (Allen-Zhu, 2017), which is close to the best-known complexity in this case⁵. If the reduction techniques in (Allen-Zhu & Hazan, 2016a; Xu et al., 2016) are used in our algorithm, our algorithm can obtain the best-known oracle complexity.

3.3. Extensions

It is common to apply reductions to extend the algorithms designed for L -smooth and σ -strongly convex objectives to other cases (e.g., non-strongly convex or non-smooth). For example, Allen-Zhu & Hazan (2016a) proposed several reductions for the algorithms that satisfy *homogeneous objective decrease* (HOOD), which is defined as follows.

Definition 1 (Allen-Zhu & Hazan (2016a)). *An algorithm \mathcal{A} that solves Problem (1) with Assumption 1 satisfies HOOD if, for every starting vector x_0 , \mathcal{A} produces an output x' satisfying⁶ $F(x') - F(x^*) \leq \frac{F(x_0) - F(x^*)}{4}$ in $\text{Time}(L, \sigma)$.*

Based on Theorem 1, it is a direct corollary that MiG (refers to Algorithm 1) satisfies HOOD.

Corollary 1. *MiG satisfies the HOOD property in $\text{Time}(n + \sqrt{\kappa n})$.*

The reductions in (Allen-Zhu & Hazan, 2016a) use either decaying regularization (AdaptReg) or certain smoothing

⁵Note that the best-known oracle complexity for non-strongly convex problems is $\mathcal{O}(n \log(1/\epsilon) + \sqrt{nL/\epsilon})$.

⁶This definition can be extended to the probabilistic guarantee, which is $\mathbb{E}[F(x')] - F(x^*) \leq \frac{F(x_0) - F(x^*)}{4}$ (Allen-Zhu & Hazan, 2016a).

trick (AdaptSmooth) to achieve optimal reductions that shave off a non-optimal log factor comparing to other reduction techniques. Thus, we can apply AdaptReg to MiG and get an improved $\mathcal{O}(n \log(1/\epsilon) + \sqrt{nL/\epsilon})$ rate for non-strongly convex problems. Moreover, we can also apply AdaptSmooth to MiG to tackle non-smooth optimization problems, e.g., SVM.

4. Sparse and Asynchronous Variants

In order to further elaborate the importance of keeping track of only one variable vector (in the inner loop), in this section we propose the variants of MiG for both the serial sparse and asynchronous sparse settings.

Adopting the sparse update technique in (Mania et al., 2017) for sparse datasets is a very practical choice to reduce collisions between threads. However, due to additional sparse approximating variance and asynchronous perturbation, we need to compensate it with a slower theoretical speed. On the other hand, asynchrony (in the lock-free style) may even destroy convergence guarantees if the algorithm requires tracking many highly correlated vectors. In practice, it is reported that maintaining more atomic⁷ variables also degrades the performance. Thus, in this section, we mainly focus on practical issues and experimental performance.

As we can see, MiG has only one variable vector. This feature gives us convenience in both theoretical analysis and practical implementation. In order to give a clean proof, we first make a simpler assumption on the objective function, which is identical to those in (Recht et al., 2011; Mania et al., 2017; Leblond et al., 2017):

$$\min_{x \in \mathbb{R}^d} F(x) \triangleq \frac{1}{n} \sum_{i=1}^n f_i(x). \quad (4)$$

Assumption 3 (Sparse and Asynchronous Settings). *In Problem (4), each $f_i(\cdot)$ is L -smooth, and the averaged function $F(\cdot)$ is σ -strongly convex.*

Next we start with analyzing MiG in the serial sparse setting and then extend it to a sparse and asynchronous one.

4.1. Serial Sparse MiG

The sparse variant (as shown in Algorithm 3) of MiG is slightly different from MiG in the dense case. We explain these differences by making the following remarks:

- *Sparse approximate gradient* $\tilde{\nabla}_S$. In order to perform fully sparse updates, following (Mania et al., 2017), we use a diagonal matrix D to re-weigh the dense vector

⁷Atomic write of some necessary variables is a requirement to achieve high precision in practice (Leblond et al., 2017).

Algorithm 3 Serial Sparse MiG

1: **Input:** Initial vector x_0 , epoch length m , learning rate η , parameter θ .
2: $\tilde{x}_0 = x_0^1 = x_0$;
3: **for** $s = 1 \dots \mathcal{S}$ **do**
4: $\mu_s = \nabla F(\tilde{x}_{s-1})$;
5: **for** $j = 1 \dots m$ **do**
6: $T_{i_j} := \text{support of sample } i_j$;
7: $[y_{j-1}]_{T_{i_j}} = \theta \cdot [x_{j-1}^s]_{T_{i_j}} + (1 - \theta) \cdot [\tilde{x}_{s-1}]_{T_{i_j}}$;
8: $\tilde{\nabla}_S = \nabla f_{i_j}([y_{j-1}]_{T_{i_j}}) - \nabla f_{i_j}([\tilde{x}_{s-1}]_{T_{i_j}}) + D_{i_j} \mu_s$;
9: $[x_j^s]_{T_{i_j}} = [x_{j-1}^s]_{T_{i_j}} - \eta \cdot \tilde{\nabla}_S$;
10: **end for**
11: Option I: $x_0^{s+1} = \tilde{x}_s = \frac{\theta}{m} \sum_{j=0}^{m-1} x_j^s + (1 - \theta) \tilde{x}_{s-1}$;
12: Option II: $x_0^{s+1} = x_m^s, \tilde{x}_s = \frac{\theta}{m} \sum_{j=1}^m x_j^s + (1 - \theta) \tilde{x}_{s-1}$;
13: **end for**
14: **return** $\tilde{x}_\mathcal{S}$.

μ_s , whose entries are the inverse probabilities $\{p_k^{-1}\}$ of the corresponding coordinates $\{k \mid k = 1, \dots, d\}$ belonging to a uniformly sampled support T_{i_j} of sample i_j . P_{i_j} is the projection matrix for the support T_{i_j} . We define $D_{i_j} = P_{i_j} D$, which ensures the unbiasedness $\mathbb{E}_{i_j}[D_{i_j} \mu_s] = \mu_s$. Here we also define $D_m = \max_{k=1 \dots d} p_k^{-1}$ for future usage. Note that we only need to compute y on the support of sample i_j , and hence the entire inner loop updates sparsely.

- *Update \tilde{x} with uniform average.* In the sparse and asynchronous setting, a weighted average in Algorithm 1 is not effective due to the perturbation both in theory and in practice. Thus, we choose a simple uniform average scheme for a better practical performance.

We now consider the convergence property of Algorithm 3.

Theorem 3 (Option I). *Let x^* be the optimal solution of Problem (4). If Assumption 3 holds, then by choosing $\eta = 1/L$, $\theta = 1/10$, $m = 25\kappa$, Algorithm 3 with Option I satisfies the following inequality in one epoch s :*

$$\mathbb{E}[(F(\tilde{x}_s) - F(x^*))] \leq 0.75 \cdot (F(\tilde{x}_{s-1}) - F(x^*)),$$

which means that the total oracle complexity of the serial sparse MiG is $\mathcal{O}((n + \kappa) \log \frac{F(x_0) - F(x^*)}{\epsilon})$.

Since it is natural to ask whether we can get an improved bound for the Serial Sparse MiG, we analyze Algorithm 3 with Option II and a somewhat intriguing restart scheme. The convergence result is given as follows:

Theorem 4 (Option II). *If Assumption 3 holds, then by executing Algorithm 3 with Option II and restarting⁸ the algorithm every $\mathcal{S} = \left\lceil 2 \cdot \frac{(1-\theta) \cdot (1+\zeta) + \frac{\theta}{\eta m \sigma}}{\theta + \zeta \theta - \zeta} \right\rceil$ epochs, where*

⁸We set $x_0 = \frac{1}{\mathcal{S}} \sum_{s=1}^{\mathcal{S}} \tilde{x}_s$ as the initial vector after each restart.

$\zeta = D_m^2 - D_m$, the oracle complexity of the entire procedure is divided into the cases,

$$\begin{cases} \mathcal{O}(\sqrt{\kappa n} \log \frac{F(x_0) - F(x^*)}{\epsilon}), & \text{if } \frac{m}{\kappa} \leq \frac{3}{4} \text{ with } \zeta \leq \sqrt{\frac{m}{4\kappa}}, \\ \mathcal{O}(n \log \frac{F(x_0) - F(x^*)}{\epsilon}), & \text{if } \frac{m}{\kappa} > \frac{3}{4} \text{ with } \zeta \leq C_\zeta, \end{cases}$$

where C_ζ is a constant for the sparse estimator variance. Detailed parameter settings are given in the Supplementary Material C.2.1.

Remark: This result indeed imposes some strong assumptions on D_m , which may not be true for real world datasets, because the variance bound used for Option II highly correlates to D_m , and D_m can be as large as n for extreme datasets. Detailed discussion is given in the Supplementary Material C.2.2.

The result of Theorem 4 shows that under some constraints on sparse variance, Serial Sparse MiG attains a faster convergence rate than Sparse SVRG (Mania et al., 2017) and Sparse SAGA (Leblond et al., 2017). Although these constraints are strong and the restart scheme is not quite practical, we keep the result here as a reference for both the sparse ($\zeta > 0$) and dense ($\zeta = 0$) cases.

4.2. Asynchronous Sparse MiG

In this part, we extend the Serial Sparse MiG to the Asynchronous Sparse MiG.

Our algorithm is given in Algorithm 4. Notice that Option I and II correspond to the update options mentioned in Algorithm 3. The difference is that Option II corresponds to averaging “fake” iterates defined at (5), while Option I is the average of inconsistent read⁹ of x . Since the averaging scheme in Option II is not proposed before, we refer to it as “fake average”. Just like the analysis in the serial sparse case, Option I leads to a direct and clean proof while Option II may require restart and leads to troublesome theoretical analysis. So we only analyze Option I in this setting.

However, Option II is shown to be highly practical since the “fake average” scheme only requires updates on the support of samples and enjoys strong robustness when the actual number of inner loops does not equal to m ¹⁰. Thus, Option II leads to a very practical implementation.

Following (Mania et al., 2017), our analysis is based on the “fake” iterates x and y , which are defined as:

$$x_j = x_0 - \eta \sum_{i=0}^{j-1} \tilde{\nabla}(\hat{y}_i), \quad y_j = \theta x_j + (1 - \theta) \tilde{x}_{s-1}, \quad (5)$$

where the “perturbed” iterates \hat{y}, \hat{x} with perturbation ξ are

⁹We could use “fake average” in Option I, but it leads to a complex proof and a worse convergence rate with factor ($\propto \kappa^{-2}$).

¹⁰This phenomenon is prevalent in the asynchronous setting.

Algorithm 4 Asynchronous Sparse MiG

- 1: **Input:** Initial vector x_0 , epoch length m , learning rate η , parameter θ .
- 2: $x :=$ shared variable, $\bar{x} :=$ average of x ;
- 3: $\tilde{x}_0 = x = x_0$;
- 4: **for** $s = 1 \dots \mathcal{S}$ **do**
- 5: Compute $\mu_s = \nabla F(\tilde{x}_{s-1})$ in parallel;
- 6: Option I: $\bar{x} = \mathbf{0}$;
- 7: Option II: $\bar{x} = x$;
- 8: $j = 0$; {inner loop counter}
- 9: **while** $j < m$ **do** {in parallel}
- 10: $j = j + 1$; // atomic increase counter j
- 11: Sample i_j uniformly in $\{1 \dots n\}$;
- 12: $T_{i_j} :=$ support of sample i_j ;
- 13: $[\hat{x}]_{T_{i_j}} :=$ inconsistent read of $[x]_{T_{i_j}}$;
- 14: $[\hat{y}]_{T_{i_j}} = \theta \cdot [\hat{x}]_{T_{i_j}} + (1 - \theta) \cdot [\tilde{x}_{s-1}]_{T_{i_j}}$;
- 15: $\tilde{\nabla}(\hat{y}) = \nabla f_{i_j}([\hat{y}]_{T_{i_j}}) - \nabla f_{i_j}([\tilde{x}_{s-1}]_{T_{i_j}}) + D_{i_j} \mu_s$;
- 16: $[u]_{T_{i_j}} = -\eta \cdot \tilde{\nabla}(\hat{y})$;
- 17: // atomic write x, \bar{x} for each coordinate
- 18: $[x]_{T_{i_j}} = [x]_{T_{i_j}} + [u]_{T_{i_j}}$;
- 19: Option I: $\bar{x} = \bar{x} + \frac{1}{m} \cdot \hat{x}$;
- 20: Option II: $[\bar{x}]_{T_{i_j}} = [\bar{x}]_{T_{i_j}} + [u]_{T_{i_j}} \cdot \frac{(m+1-j)_+}{m}$;
- 21: **end while**
- 22: $\tilde{x}_s = \theta \bar{x} + (1 - \theta) \tilde{x}_{s-1}$;
- 23: Option I: $x = \tilde{x}_s$;
- 24: Option II: keep x unchanged;
- 25: **end for**
- 26: **return** $\tilde{x}_{\mathcal{S}}$.

defined as

$$\hat{y}_j = \theta \hat{x}_j + (1 - \theta) \tilde{x}_{s-1}, \hat{x}_j = x_j + \xi_j. \quad (6)$$

The labeling order and detailed analysis framework are given in the Supplementary Material C.3.


Note that y is a temp variable, so the only source of perturbation comes from x . This is the benefit of keeping track of only one variable vector since it controls the perturbation and allows us to give a smooth analysis in asynchrony.

Next we give our convergence result as follows:


Theorem 5. *If Assumption 3 holds, then by choosing $m = 60\kappa$, $\eta = 1/(5L)$, $\theta = 1/6$, suppose τ satisfies $\tau \leq \min\{\frac{5}{4\sqrt{\Delta}}, 2\kappa, \sqrt{\frac{2\kappa}{\sqrt{\Delta}}}\}$ (the linear speed-up condition), Algorithm 4 with Option I has the following oracle complexity:*

$$\mathcal{O}\left((n + \kappa) \log \frac{F(x_0) - F(x^*)}{\epsilon}\right),$$

where τ represents the maximum number of overlaps between concurrent threads (Mania et al., 2017) and $\Delta = \max_{k=1 \dots d} p_k$, which is a measure of sparsity (Leblond et al., 2017).



(a) Relatively small regularization parameter $\lambda = 10^{-8}$



(b) Relatively large regularization parameter $\lambda = 10^{-4}$

Figure 1. Comparison of the effect of the momentum techniques used in Acc-Prox-SVRG (Nitanda, 2014), Katyusha (Allen-Zhu, 2017) and MiG for ℓ_2 -logistic regression on covtype.

This result is better than that of KroMagnon, which correlates to κ^2 (Mania et al., 2017), and keeps up with ASAGA (Leblond et al., 2017). Although without significant improvement on theoretical bounds due to the existence of perturbation, the coupling step of MiG can still be regarded as a simple add-on boosting and stabilizing the performance of SVRG variants. We show this improvement by empirical evaluations in Section 5.3.

5. Experiments

In this section, we evaluate the performance of MiG on real-world datasets for both serial dense and asynchronous sparse¹¹ cases. All the algorithms were implemented in C++ and executed through MATLAB interface for a fair comparison. Detailed experimental setup is given in the Supplementary Material D.

We first give a detailed comparison between MiG and other algorithms in the sequential dense setting.

5.1. Comparison of Momentums

The parameter θ in MiG, similar to the parameter τ_2 in Katyusha, is referred as the parameter for “Katyusha Momentum” in (Allen-Zhu, 2017). So intuitively, MiG can be regarded as adding “Katyusha Momentum” on top of SVRG. Katyusha equipped with the linear coupling framework in (Allen-Zhu & Orecchia, 2017) and thus can be re-

¹¹Experiments for the serial sparse variant are omitted since it corresponds to the asynchronous sparse variant with 1 thread.

Table 3. Compared algorithms and parameter tuning criteria.

	Momentum	Parameter Tuning
SVRG	None	learning rate η
Acc-Prox-SVRG	Nestrv.	same η , tune momentum β
Katyusha	Nestrv.&Katyu.	$\tau_2 = \frac{1}{2}$, $\alpha = \frac{1}{3\tau_1 L}$, tune τ_1
MiG	Katyu.	$\eta = \frac{1}{3\theta L}$, tune θ

garded as the combination of “*Nesterov’s Momentum*” with “*Katyusha Momentum*”.

We empirically evaluate the effect of the two kinds of momentums. Moreover, we also examine the performance of Acc-Prox-SVRG (Nitanda, 2014), which can be regarded as SVRG with pure “*Nesterov’s Momentum*”. Note that the mini-batch size used in Acc-Prox-SVRG is set to 1. The algorithms and parameter settings are listed in Table 3 (we use the same notations as in their original papers).

The results in Figure 1 correspond to the two typical conditions with relatively large $\lambda = 10^{-4}$ and relatively small $\lambda = 10^{-8}$. One can verify that with the epoch length $m = 2n$ for all algorithms, these two conditions fall into the two regions correspondingly in Table 2.

For the case of $\frac{m}{\kappa} \leq \frac{3}{4}$ (see the first row in Figure 1), we set the parameters for MiG and Katyusha¹² with their theoretical suggestions (e.g., $\theta = \tau_1 = \sqrt{\frac{m}{3\kappa}}$). For fair comparison, we set the learning rate $\eta = \frac{1}{4L}$ for SVRG and Acc-Prox-SVRG, which is theoretically reasonable. The results imply that MiG and Katyusha have close convergence results and outperform SVRG and Acc-Prox-SVRG. This justifies the improvement of $\sqrt{\kappa n}$ convergence rate in theory.

We notice that Katyusha is slightly faster than MiG in terms of the number of oracle calls, which is reasonable since Katyusha has one more “*Nesterov’s Momentum*”. From the result of Acc-Prox-SVRG, we see that “*Nesterov’s Momentum*” is effective in this case, but without significant improvement. As analyzed in (Nitanda, 2014), using large enough mini-batch is a requirement to make Acc-Prox-SVRG improve its convergence rate in theory (see Table 1 in (Nitanda, 2014)), which also explains the limited difference between MiG and Katyusha.

When comparing running time (millisecond, ms), MiG outperforms other algorithms. Using “*Nesterov’s Momentum*” requires tracking of at least two variable vectors, which increases both memory consumption and computational overhead. More severely, it prevents the algorithms with this trick to have an efficient sparse and asynchronous variant.

¹²We choose to implement Katyusha with Option I, which is analyzed theoretically in (Allen-Zhu, 2017).

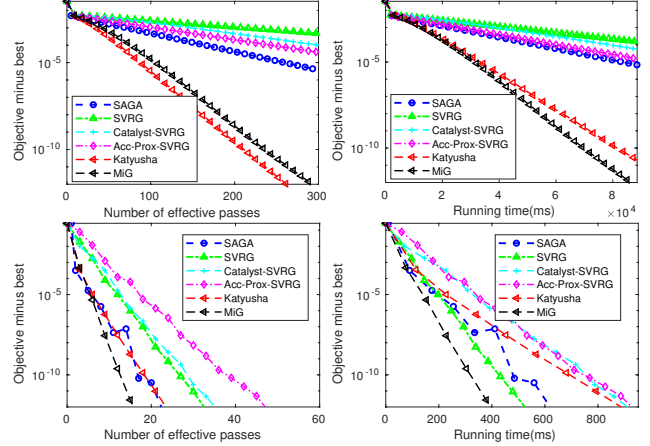


Figure 2. First row: Theoretical evaluation of MiG and state-of-the-art algorithms for ℓ_2 -logistic regression ($\lambda = 10^{-8}$) on covtype. Second row: Practical evaluation of MiG and the state-of-the-art algorithms for ridge regression ($\lambda = 10^{-4}$) on a9a.

For the case of $\frac{m}{\kappa} > \frac{3}{4}$ (see the second row in Figure 1), we tuned all the parameters in Table 3. From the parameter tuning of Acc-Prox-SVRG, we found that using a smaller momentum parameter β yields a better performance, but still worse than the original SVRG. This result seems to indicate that the “*Nesterov’s Momentum*” is not effective in this case. Katyusha yields a poor performance in this case because the parameter suggestion limits $\tau_1 \leq \frac{1}{2}$. When tuning both τ_2 and τ_1 , Katyusha performs much better, but with increasing difficulty of parameter tuning. MiG performs quite well with tuning only one parameter θ , which further verifies its simplicity and efficiency.

5.2. Comparison with state-of-the-art Algorithms

We compare MiG with many state-of-the-art accelerated algorithms (e.g., Acc-Prox-SVRG (Nitanda, 2014), Catalyst (based on SVRG) (Lin et al., 2015), and Katyusha (Allen-Zhu, 2017)) and non-accelerated algorithms (e.g., SVRG (Johnson & Zhang, 2013) or Prox-SVRG (Xiao & Zhang, 2014) for non-smooth regularizer, and SAGA (Defazio et al., 2014)), as shown in Figure 2.

In order to give clear comparisons, we designed two different types of experiments. One is called “theoretical evaluation” with a relatively small λ ¹³, where most of the parameter settings follow the corresponding theoretical recommendations¹⁴ to justify the improvement of $\sqrt{\kappa n}$ convergence rate. Another is “practical evaluation” for a relatively large

¹³Note that we normalize data vectors to ensure a uniform L .

¹⁴Except for Acc-Prox-SVRG and Catalyst, we carefully tuned the parameters for them, and the detailed parameter settings are given in the Supplementary Material D.1.




Figure 3. Comparison of KroMagnon (Mania et al., 2017), ASAGA (Leblond et al., 2017), and MiG with 16 threads. First row: RCV1, ℓ_2 -logistic regression with $\lambda = 10^{-9}$. Second row: KDD2010, ℓ_2 -logistic regression with $\lambda = 10^{-10}$.

λ , where we carefully tuned the parameters for all the algorithms since in this condition, all the algorithms have similar convergence rates.

For a relatively large λ , the results (see the Supplementary Material D.1 for more results) show that MiG performs consistently better than Katyusha in terms of both oracle calls and running time. In other words, we see that MiG achieves satisfactory performance in both conditions. Moreover, experimental results for non-strongly convex objectives are also given in the Supplementary Material D.1.

5.3. In Sparse and Asynchronous Settings

To further stress the simplicity and implementability of MiG, we make some experiments to assess the performance of its asynchronous variant. We also compare MiG (i.e., Algorithm 4 with Option II¹⁵) with KroMagnon (Mania et al., 2017) and ASAGA (Leblond et al., 2017).

Unlike in the serial dense case where we have strong theoretical guarantees, in these settings, we mainly focus on practical performance and stability. So we carefully tuned the parameter(s) for each algorithm to achieve a best-tuned performance (detailed setting and parameter tuning criteria is given in the Supplementary Material D.2). We measure the performance on the two sparse datasets listed in Table 4.

When comparing performance in terms of oracle calls, MiG significantly outperforms other algorithms, as shown in Figure 3. When considering running time, the difference is narrowed due to the high simplicity of KroMagnon (which

¹⁵We omit the tricky restart scheme required in theory to examine the most practical variant.

Table 4. Summary of the two sparse data sets.

Dataset	# Data	# Features	Density
RCV1	697,641	47,236	1.5×10^{-3}
KDD2010	19,264,097	1,163,024	10^{-6}




Figure 4. Speed-up evaluation on RCV1. Left: Evaluation of sub-optimality in terms of running time for asynchronous versions (20 threads) and SS (Serial Sparse) versions. Right: Speed-up of achieving 10^{-5} sub-optimality in terms of the number of threads.

only uses one atomic vector) compared with ASAGA (which uses atomic gradient table and atomic gradient average vector) and MiG (which only uses atomic “fake average”).

We then examine the speed-up gained from more parallel threads on RCV1. We evaluate the improvement of using asynchronous variants (20 threads) and the speed-up ratio as a function of the number of threads as shown in Figure 4. For the latter evaluation, the running time is recorded when the algorithms achieve 10^{-5} sub-optimality. The speed-up ratio is calculated based on the running time of a single core.

Since we used MiG with Option II for the above experiments which only has a theoretical analysis (with restart) in the serial case, we further designed an experiment to evaluate the effectiveness of θ . The results in the Supplementary Material D.2.2 indicate the effectiveness of our acceleration trick.

6. Conclusion

We proposed a simple stochastic variance reduction algorithm (MiG) with the best-known oracle complexities for stochastic first-order algorithms. These elegant results further reveal the mystery of acceleration tricks in stochastic first-order optimization. Moreover, the high simplicity of MiG allows us to derive the variants for the asynchronous and sparse settings, which shows its potential to be applied to other cases (e.g., online (Borsos et al., 2018), distributed (Lee et al., 2017; Lian et al., 2017)) as well as to tackle more complex problems (e.g., structured prediction). In general, our approach can be implemented to boost the speed of large-scale real-world optimization problems.

Acknowledgements

We thank the reviewers for their valuable comments. This work was supported in part by Grants (CUHK 14206715 & 14222816) from the Hong Kong RGC, the Major Research Plan of the National Natural Science Foundation of China (Nos. 91438201 and 91438103), Project supported the Foundation for Innovative Research Groups of the National Natural Science Foundation of China (No. 61621005), the National Natural Science Foundation of China (Nos. U1701267, 61573267, 61502369 and 61473215), the Program for Cheung Kong Scholars and Innovative Research Team in University (No. IRT_15R53) and the Fund for Foreign Scholars in University Research and Teaching Programs (the 111 Project) (No. B07048).

References

- Allen-Zhu, Z. Katyusha: The first direct acceleration of stochastic gradient methods. In *STOC*, 2017.
- Allen-Zhu, Z. and Hazan, E. Optimal black-box reductions between optimization objectives. In *NIPS*, pp. 1606–1614, 2016a.
- Allen-Zhu, Z. and Hazan, E. Variance reduction for faster non-convex optimization. In *ICML*, pp. 699–707, 2016b.
- Allen-Zhu, Z. and Orecchia, L. Linear coupling: An ultimate unification of gradient and mirror descent. In *ITCS*, 2017.
- Boros, Z., Krause, A., and Levy, K. Y. Online variance reduction for stochastic optimization. *arXiv:1802.04715v2*, 2018.
- Defazio, A. A simple practical accelerated method for finite sums. In *NIPS*, pp. 676–684, 2016.
- Defazio, A., Bach, F., and Lacoste-Julien, S. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *NIPS*, pp. 1646–1654, 2014.
- Hien, L., Lu, C., Xu, H., and Feng, J. Accelerated stochastic mirror descent algorithms for composite non-strongly convex optimization. *arXiv:1605.06892v4*, 2017.
- Johnson, R. and Zhang, T. Accelerating stochastic gradient descent using predictive variance reduction. In *NIPS*, pp. 315–323, 2013.
- Konečný, J., Liu, J., Richtárik, P., and Takáč, M. Mini-batch semi-stochastic gradient descent in the proximal setting. *IEEE J. Sel. Top. Sign. Proces.*, 10(2):242–255, 2016.
- Leblond, R., Pedregosa, F., and Lacoste-Julien, S. ASAGA: Asynchronous parallel SAGA. In *AISTATS*, pp. 46–54, 2017.
- Lee, J. D., Lin, Q., Ma, T., and Yang, T. Distributed stochastic variance reduced gradient methods by sampling extra data with replacement. *J. Mach. Learn. Res.*, 18:1–43, 2017.
- Lian, X., Zhang, C., Zhang, H., Hsieh, C.-J., Zhang, W., and Liu, J. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *NIPS*, pp. 5336–5346, 2017.
- Lin, H., Mairal, J., and Harchaoui, Z. A universal catalyst for first-order optimization. In *NIPS*, pp. 3366–3374, 2015.
- Lin, Q., Lu, Z., and Xiao, L. An accelerated proximal coordinate gradient method. In *NIPS*, pp. 3059–3067, 2014.
- Mania, H., Pan, X., Papailiopoulos, D., Recht, B., Ramchandran, K., and Jordan, M. I. Perturbed iterate analysis for asynchronous stochastic optimization. *SIAM J. Optim.*, 27(4):2202–2229, 2017.
- Nesterov, Y. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer Academic Publ., Boston, 2004.
- Nitanda, A. Stochastic proximal gradient descent with acceleration techniques. In *NIPS*, pp. 1574–1582, 2014.
- Pedregosa, F., Leblond, R., and Lacoste-Julien, S. Breaking the nonsmooth barrier: A scalable parallel method for composite optimization. In *NIPS*, pp. 55–64, 2017.
- Recht, B., Re, C., Wright, S., and Niu, F. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *NIPS*, pp. 693–701, 2011.
- Reddi, S., Hefny, A., Sra, S., Póczos, B., and Smola, A. On variance reduction in stochastic gradient descent and its asynchronous variants. In *NIPS*, pp. 2629–2637, 2015.
- Robbins, H. and Monro, S. A stochastic approximation method. *Ann. Math. Statist.*, 22(3):400–407, 1951.
- Roux, N. L., Schmidt, M., and Bach, F. A stochastic gradient method with an exponential convergence rate for finite training sets. In *NIPS*, pp. 2672–2680, 2012.
- Schmidt, M., Roux, N. L., and Bach, F. Minimizing finite sums with the stochastic average gradient. *Math. Program.*, 162:83–112, 2017.
- Shalev-Shwartz, S. and Zhang, T. Stochastic dual coordinate ascent methods for regularized loss minimization. *J. Mach. Learn. Res.*, 14:567–599, 2013.

Shalev-Shwartz, S. and Zhang, T. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Math. Program.*, 155:105–145, 2016.

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

Xu, Y., Yan, Y., Lin, Q., and Yang, T. Homotopy smoothing for non-smooth problems with lower complexity than $O(1/\epsilon)$. In *NIPS*, pp. 1208–1216, 2016.

Zhang, L., Mahdavi, M., and Jin, R. Linear convergence with condition number independent access of full gradients. In *NIPS*, pp. 980–988, 2013.

Zhang, Y. and Xiao, L. Stochastic primal-dual coordinate method for regularized empirical risk minimization. In *ICML*, pp. 353–361, 2015.

Supplementary Materials for “A Simple Stochastic Variance Reduced Algorithm with Fast Convergence Rates”

Notations. We use \mathbb{E}_{i_j} to denote that the expectation is taken with respect to the j th sample in one epoch, while \mathbb{E} means taking expectation with respect to all randomness in one epoch. We use $O(\cdot)$ to denote computational complexity and $\mathcal{O}(\cdot)$ to denote oracle complexity. x^* refers to the solution to Problem (1) (the proofs for Section 3) or Problem (4) (the proofs for Section 4). S is the total number of epochs to be executed. Boldface number like $\mathbf{0}$ refers to a vector with all 0.

Note: In order to give a clean proof, we omit the superscripts for iterates in the same epoch s as x_j instead of x_j^s unless otherwise specified.

A. Useful Lemmas

Lemma 1. (Variance Bound) *Suppose each component function f_i is L -smooth, let $\tilde{\nabla} = \nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1}) + \nabla f(\tilde{x}_{s-1})$, which is the approximate gradient used in MiG. Then the following inequality holds:*

$$\mathbb{E}_{i_j} [\|\nabla f(y_{j-1}) - \tilde{\nabla}\|^2] \leq 2L(f(\tilde{x}_{s-1}) - f(y_{j-1}) - \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle).$$

Proof. This lemma is identical to Lemma 3.4 in (Allen-Zhu, 2017), which provides a tighter upper bound on the gradient estimator variance than those in (Johnson & Zhang, 2013; Xiao & Zhang, 2014). \square

Lemma 2. (3-points property) *Assume that z^* is an optimal solution to the following problem,*

$$\min_x \frac{\tau}{2} \|x - z_0\|^2 + \psi(x),$$

where $\tau > 0$, and $\psi(\cdot)$ is a convex function (but possibly non-differentiable). Then for all $z \in \mathbb{R}^d$, there exists a vector $\mathcal{G} \in \partial\psi(z^*)$ with

$$\langle \mathcal{G}, z - z^* \rangle = \frac{\tau}{2} \|z^* - z_0\|^2 - \frac{\tau}{2} \|z - z_0\|^2 + \frac{\tau}{2} \|z - z^*\|^2,$$

where $\partial\psi(z^*)$ denotes the sub-differential of $\psi(\cdot)$ at z^* . If $\psi(\cdot)$ is differentiable, we can simply replace $\mathcal{G} \in \partial\psi(\cdot)$ with $\mathcal{G} = \nabla\psi(\cdot)$.

Proof. By the optimality of z^* , there exists a vector $\mathcal{G} \in \partial\psi(z^*)$ (or $\mathcal{G} = \nabla\psi(z^*)$ for differentiable $\psi(\cdot)$) satisfying

$$\tau(z^* - z_0) + \mathcal{G} = \mathbf{0}.$$

Thus for all $z \in \mathbb{R}^d$,

$$\begin{aligned} 0 &= \langle \tau(z^* - z_0) + \mathcal{G}, z^* - z \rangle \\ &= \tau \langle z^* - z_0, z^* - z \rangle + \langle \mathcal{G}, z^* - z \rangle \\ &\stackrel{(*)}{=} \frac{\tau}{2} \|z^* - z_0\|^2 - \frac{\tau}{2} \|z - z_0\|^2 + \frac{\tau}{2} \|z - z^*\|^2 + \langle \mathcal{G}, z^* - z \rangle, \end{aligned}$$

where $(*)$ uses the fact that $\langle a - b, a - c \rangle = \frac{1}{2} \|a - b\|^2 - \frac{1}{2} \|b - c\|^2 + \frac{1}{2} \|a - c\|^2$. \square

Lemma 3. *If two vector $x_j, x_{j-1} \in \mathbb{R}^d$ satisfy $x_j = \arg \min_x \{ \frac{1}{2\eta} \|x - x_{j-1}\|^2 + \langle \tilde{\nabla}, x \rangle + g(x) \}$ with a constant vector $\tilde{\nabla}$ and a general convex function $g(\cdot)$, then for all $u \in \mathbb{R}^d$, we have*

$$\langle \tilde{\nabla}, x_j - u \rangle \leq -\frac{1}{2\eta} \|x_{j-1} - x_j\|^2 + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1}{2\eta} \|x_j - u\|^2 + g(u) - g(x_j).$$

Moreover, if $g(\cdot)$ is σ -strongly convex, the above inequality becomes

$$\langle \tilde{\nabla}, x_j - u \rangle \leq -\frac{1}{2\eta} \|x_{j-1} - x_j\|^2 + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1 + \eta\sigma}{2\eta} \|x_j - u\|^2 + g(u) - g(x_j).$$

Proof. Applying Lemma 2 with $z = u, z_0 = x_{j-1}, z^* = x_j, \tau = \frac{1}{\eta}, \psi(x) = \langle \tilde{\nabla}, x \rangle + g(x)$, there exists a vector $\mathcal{G} \in \partial g(x_j)$ (or $\mathcal{G} = \nabla g(x_j)$ for differentiable $g(\cdot)$) satisfying

$$\langle \tilde{\nabla}, u - x_j \rangle + \langle \mathcal{G}, u - x_j \rangle = \frac{1}{2\eta} \|x_{j-1} - x_j\|^2 - \frac{1}{2\eta} \|x_{j-1} - u\|^2 + \frac{1}{2\eta} \|x_j - u\|^2.$$

Using the convexity of $g(\cdot)$, we get $g(u) - g(x_j) \geq \langle \mathcal{G}, u - x_j \rangle$ by definition. After rearranging, we conclude that

$$\langle \tilde{\nabla}, x_j - u \rangle \leq -\frac{1}{2\eta} \|x_{j-1} - x_j\|^2 + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1}{2\eta} \|x_j - u\|^2 + g(u) - g(x_j).$$

If $g(\cdot)$ is further σ -strongly convex, we have $g(u) - g(x_j) \geq \langle \mathcal{G}, u - x_j \rangle + \frac{\sigma}{2} \|x_j - u\|^2$ by (3). Similarly, we can write

$$\langle \tilde{\nabla}, x_j - u \rangle \leq -\frac{1}{2\eta} \|x_{j-1} - x_j\|^2 + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1 + \eta\sigma}{2\eta} \|x_j - u\|^2 + g(u) - g(x_j).$$

□

B. Proofs for Section 3

B.1. Proof of Theorem 1

First, we add the following constraint on the parameters η and θ , which is crucial in the proof of Theorem 1:

$$L\theta + \frac{L\theta}{1 - \theta} \leq \frac{1}{\eta}, \text{ or equivalently } \eta \leq \frac{1 - \theta}{L\theta(2 - \theta)}. \quad (7)$$

We start with convexity of $f(\cdot)$ at y_{j-1} . By definition, for any vector $u \in \mathbb{R}^d$, we have

$$\begin{aligned} f(y_{j-1}) - f(u) &\leq \langle \nabla f(y_{j-1}), y_{j-1} - u \rangle \\ &= \langle \nabla f(y_{j-1}), y_{j-1} - x_{j-1} \rangle + \langle \nabla f(y_{j-1}), x_{j-1} - u \rangle \\ &\stackrel{(\star)}{=} \frac{1 - \theta}{\theta} \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \langle \nabla f(y_{j-1}), x_{j-1} - u \rangle, \end{aligned} \quad (8)$$

where (\star) follows from the fact that $y_{j-1} = \theta x_{j-1} + (1 - \theta)\tilde{x}_{s-1}$.

Then we further expand $\langle \nabla f(y_{j-1}), x_{j-1} - u \rangle$ as

$$\langle \nabla f(y_{j-1}), x_{j-1} - u \rangle = \langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_{j-1} - u \rangle + \langle \tilde{\nabla}, x_{j-1} - x_j \rangle + \langle \tilde{\nabla}, x_j - u \rangle. \quad (9)$$

Using L -smooth (2) of $f(\cdot)$ at (y_j, y_{j-1}) , we get

$$\begin{aligned} f(y_j) - f(y_{j-1}) &\leq \langle \nabla f(y_{j-1}), y_j - y_{j-1} \rangle + \frac{L}{2} \|y_j - y_{j-1}\|^2 \\ &\stackrel{(\star)}{=} \theta \langle \nabla f(y_{j-1}), x_j - x_{j-1} \rangle + \frac{L\theta^2}{2} \|x_j - x_{j-1}\|^2 \\ &= \theta [\langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - x_{j-1} \rangle + \langle \tilde{\nabla}, x_j - x_{j-1} \rangle] + \frac{L\theta^2}{2} \|x_j - x_{j-1}\|^2, \\ \langle \tilde{\nabla}, x_{j-1} - x_j \rangle &\leq \frac{1}{\theta} (f(y_{j-1}) - f(y_j)) + \langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - x_{j-1} \rangle + \frac{L\theta}{2} \|x_j - x_{j-1}\|^2, \end{aligned}$$

where (\star) uses the definition of y_{j-1} .

After plugging in the constraint (7), we have

$$\langle \tilde{\nabla}, x_{j-1} - x_j \rangle \leq \frac{1}{\theta} (f(y_{j-1}) - f(y_j)) + \langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - x_{j-1} \rangle + \frac{1}{2\eta} \|x_j - x_{j-1}\|^2 - \frac{L\theta}{2(1-\theta)} \|x_j - x_{j-1}\|^2. \quad (10)$$

Then we are ready to combine (8), (9), (10), as well as Lemma 3 (here $g(x)$ is σ -strongly convex by Assumption 1), which gives

$$\begin{aligned} f(y_{j-1}) - f(u) &\leq \frac{1-\theta}{\theta} \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - u \rangle + \frac{1}{\theta} (f(y_{j-1}) - f(y_j)) \\ &\quad - \frac{L\theta}{2(1-\theta)} \|x_j - x_{j-1}\|^2 + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1+\eta\sigma}{2\eta} \|x_j - u\|^2 + g(u) - g(x_j). \end{aligned}$$

After taking expectation with respect to the sample i_j , we obtain

$$\begin{aligned} f(y_{j-1}) - f(u) &\stackrel{(a)}{\leq} \frac{1-\theta}{\theta} \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \mathbb{E}_{i_j} [\langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - x_{j-1} \rangle] + \frac{1}{\theta} (f(y_{j-1}) - \mathbb{E}_{i_j} [f(y_j)]) \\ &\quad - \frac{L\theta}{2(1-\theta)} \mathbb{E}_{i_j} [\|x_j - x_{j-1}\|^2] + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1+\eta\sigma}{2\eta} \mathbb{E}_{i_j} [\|x_j - u\|^2] + g(u) - \mathbb{E}_{i_j} [g(x_j)] \\ &\stackrel{(b)}{\leq} \frac{1-\theta}{\theta} \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \frac{1}{2\beta} \mathbb{E}_{i_j} [\|\nabla f(y_{j-1}) - \tilde{\nabla}\|^2] + \frac{\beta}{2} \mathbb{E}_{i_j} [\|x_j - x_{j-1}\|^2] \\ &\quad + \frac{1}{\theta} (f(y_{j-1}) - \mathbb{E}_{i_j} [f(y_j)]) - \frac{L\theta}{2(1-\theta)} \mathbb{E}_{i_j} [\|x_j - x_{j-1}\|^2] + \frac{1}{2\eta} \|x_{j-1} - u\|^2 \\ &\quad - \frac{1+\eta\sigma}{2\eta} \mathbb{E}_{i_j} [\|x_j - u\|^2] + g(u) - \mathbb{E}_{i_j} [g(x_j)], \end{aligned}$$

where (a) holds due to the unbiasedness of the gradient estimator $\mathbb{E}_{i_j} [\nabla f(y_{j-1}) - \tilde{\nabla}] = \mathbf{0}$, and (b) uses the Young's inequality to expand $\mathbb{E}_{i_j} [\langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - x_{j-1} \rangle]$ with the parameter $\beta > 0$.

Applying Lemma 1 to bound the variance term $\mathbb{E}_{i_j} [\|\nabla f(y_{j-1}) - \tilde{\nabla}\|^2]$, we get

$$\begin{aligned} f(y_{j-1}) - f(u) &\leq \frac{1-\theta}{\theta} \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \frac{L}{\beta} (f(\tilde{x}_{s-1}) - f(y_{j-1}) - \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle) \\ &\quad + \frac{\beta}{2} \mathbb{E}_{i_j} [\|x_j - x_{j-1}\|^2] + \frac{1}{\theta} (f(y_{j-1}) - \mathbb{E}_{i_j} [f(y_j)]) - \frac{L\theta}{2(1-\theta)} \mathbb{E}_{i_j} [\|x_j - x_{j-1}\|^2] + \frac{1}{2\eta} \|x_{j-1} - u\|^2 \\ &\quad - \frac{1+\eta\sigma}{2\eta} \mathbb{E}_{i_j} [\|x_j - u\|^2] + g(u) - \mathbb{E}_{i_j} [g(x_j)]. \end{aligned}$$

Let $\beta = \frac{L\theta}{1-\theta} > 0$, by rearranging the above inequality, we obtain

$$\begin{aligned} 0 &\leq \frac{1-\theta}{\theta} f(\tilde{x}_{s-1}) - \frac{1}{\theta} \mathbb{E}_{i_j} [f(y_j)] + F(u) - \mathbb{E}_{i_j} [g(x_j)] + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1+\eta\sigma}{2\eta} \mathbb{E}_{i_j} [\|x_j - u\|^2] \\ &\stackrel{(\star)}{\leq} \frac{1-\theta}{\theta} F(\tilde{x}_{s-1}) - \frac{1}{\theta} \mathbb{E}_{i_j} [F(y_j)] + F(u) + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1+\eta\sigma}{2\eta} \mathbb{E}_{i_j} [\|x_j - u\|^2], \\ \frac{1}{\theta} (\mathbb{E}_{i_j} [F(y_j)] - F(u)) &\leq \frac{1-\theta}{\theta} (F(\tilde{x}_{s-1}) - F(u)) + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1+\eta\sigma}{2\eta} \mathbb{E}_{i_j} [\|x_j - u\|^2], \quad (11) \end{aligned}$$

where (\star) follows from the Jensen's inequality and the definition of y_{j-1} , which leads to $-g(x_j) \leq \frac{1-\theta}{\theta} g(\tilde{x}_{s-1}) - \frac{1}{\theta} g(y_j)$.

Let $u = x^*$, using our choice of $\omega = 1 + \eta\sigma$ to sum (11) over $j = 1 \dots m$ with increasing weight ω^{j-1} . After taking expectation with respect to all randomness in this epoch, we have

$$\frac{1}{\theta} \sum_{j=0}^{m-1} \omega^j (\mathbb{E} [F(y_{j+1})] - F(x^*)) + \frac{\omega^m}{2\eta} \mathbb{E} [\|x_m - x^*\|^2] \leq \frac{1-\theta}{\theta} \sum_{j=0}^{m-1} \omega^j (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{1}{2\eta} \|x_0 - x^*\|^2.$$

Using the Jensen's inequality and $\tilde{x}_s = \theta \left(\sum_{j=0}^{m-1} \omega^j \right)^{-1} \sum_{j=0}^{m-1} \omega^j x_{j+1} + (1-\theta)\tilde{x}_{s-1} = \left(\sum_{j=0}^{m-1} \omega^j \right)^{-1} \sum_{j=0}^{m-1} \omega^j y_{j+1}$, we have

$$\frac{1}{\theta} \sum_{j=0}^{m-1} \omega^j (\mathbb{E}[F(\tilde{x}_s)] - F(x^*)) + \frac{\omega^m}{2\eta} \mathbb{E}[\|x_m - x^*\|^2] \leq \frac{1-\theta}{\theta} \sum_{j=0}^{m-1} \omega^j (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{1}{2\eta} \|x_0 - x^*\|^2. \quad (12)$$

(I) Consider the first case in Theorem 1 with $\frac{m}{\kappa} \leq \frac{3}{4}$, we set $\eta = \sqrt{\frac{1}{3\sigma m L}}$, $\theta = \sqrt{\frac{m}{3\kappa}} \leq \frac{1}{2}$, and $m = \Theta(n)$.

First, we evaluate the crucial constraint (7). By substituting in our parameter settings, the constraint becomes

$$L\theta + \frac{L\theta}{1-\theta} \leq \frac{1}{\eta} \rightarrow \sqrt{\frac{m}{\kappa}} \leq \frac{\sqrt{3}}{2}.$$

Thus the constraint is satisfied by meeting the case assumption.

Then we focus on $(1-\theta)\omega^m$, observed that

$$(1-\theta)\omega^m = \left(1 - \sqrt{\frac{m}{3\kappa}}\right) \cdot \left(1 + \sqrt{\frac{1}{3m\kappa}}\right)^m.$$

Let $\zeta = \sqrt{\frac{m}{\kappa}}$, $\zeta \in (0, \frac{\sqrt{3}}{2}]$, we can denote

$$\phi(\zeta) = \left(1 - \frac{\sqrt{3}}{3}\zeta\right) \cdot \left(1 + \frac{\sqrt{3}}{3} \cdot \frac{\zeta}{m}\right)^m$$

as a function of ζ .

By taking derivative with respect to ζ , we find that $\phi(\zeta)$ is monotonically decreasing on $[0, \frac{\sqrt{3}}{2}]$ for any $m > 0$, which means

$$(1-\theta)\omega^m \leq \max_{\zeta \in (0, \frac{\sqrt{3}}{2}]} \phi(\zeta) \leq \phi(0) = 1.$$

Thus we have $\frac{1}{\theta} \geq \frac{1-\theta}{\theta} \omega^m$. By using this inequality in (12), we get

$$\begin{aligned} & \frac{1-\theta}{\theta} \sum_{j=0}^{m-1} \omega^j (\mathbb{E}[F(\tilde{x}_s)] - F(x^*)) + \frac{1}{2\eta} \mathbb{E}[\|x_m - x^*\|^2] \\ & \leq \omega^{-m} \cdot \left(\frac{1-\theta}{\theta} \sum_{j=0}^{m-1} \omega^j (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{1}{2\eta} \|x_0 - x^*\|^2 \right). \end{aligned}$$

Dividing both sides of the above inequality by $\frac{1-\theta}{\theta} \sum_{j=0}^{m-1} \omega^j$, we get

$$\begin{aligned} & (\mathbb{E}[F(\tilde{x}_s)] - F(x^*)) + \frac{\theta}{2\eta(1-\theta) \sum_{j=0}^{m-1} \omega^j} \mathbb{E}[\|x_m - x^*\|^2] \\ & \leq \omega^{-m} \cdot \left((F(\tilde{x}_{s-1}) - F(x^*)) + \frac{\theta}{2\eta(1-\theta) \sum_{j=0}^{m-1} \omega^j} \|x_0 - x^*\|^2 \right). \end{aligned}$$

Summing the above inequality over $s = 1 \dots \mathcal{S}$, we get

$$\begin{aligned} & (\mathbb{E}[F(\tilde{x}_{\mathcal{S}})] - F(x^*)) + \frac{\theta}{2\eta(1-\theta) \sum_{j=0}^{m-1} \omega^j} \mathbb{E}[\|x_{\mathcal{S}}^m - x^*\|^2] \\ & \leq \omega^{-\mathcal{S}m} \cdot \left((F(\tilde{x}_0) - F(x^*)) + \frac{\theta}{2\eta(1-\theta) \sum_{j=0}^{m-1} \omega^j} \|x_0^1 - x^*\|^2 \right). \end{aligned}$$

Notice that in order to prevent confusion, we mark iterates with epoch number, such as $x_m^{\mathcal{S}}$ represent the last iterate in epoch \mathcal{S} .

Using the fact that $\sum_{j=0}^{m-1} \omega^j \geq m$, we have

$$\left(\mathbb{E}[F(\tilde{x}_S)] - F(x^*)\right) \leq \omega^{-Sm} \cdot \left((F(\tilde{x}_0) - F(x^*)) + \frac{\theta}{2\eta(1-\theta)m} \|x_0^1 - x^*\|^2 \right).$$

Using the σ -strongly convexity of $F(\cdot)$ to bound $\|x_0^1 - x^*\|^2$, which is $\|x_0^1 - x^*\|^2 \leq \frac{2}{\sigma} (F(x_0^1) - F(x^*))$, we obtain

$$\mathbb{E}[F(\tilde{x}_S) - F(x^*)] \leq (1 + \eta\sigma)^{-Sm} \cdot \left(1 + \frac{\theta}{\eta(1-\theta)m\sigma} \right) \cdot (F(\tilde{x}_0) - F(x^*)).$$

Note that $\tilde{x}_0 = x_0^1 = x_0$.

By substituting with our parameters setting, we get

$$\begin{aligned} \mathbb{E}[F(\tilde{x}_S) - F(x^*)] &\stackrel{(*)}{\leq} \left(O\left(1 + \sqrt{\frac{1}{3n\kappa}}\right) \right)^{-Sm} \cdot O\left(1 + 2\theta\sqrt{\frac{\kappa}{n}}\right) \cdot (F(\tilde{x}_0) - F(x^*)) \\ &\leq \left(O\left(1 + \sqrt{\frac{1}{3n\kappa}}\right) \right)^{-Sm} \cdot O(F(\tilde{x}_0) - F(x^*)), \end{aligned}$$

where $(*)$ holds due to the fact that $\theta \leq \frac{1}{2}$.

The above result implies that the oracle complexity in the case $\frac{m}{\kappa} \leq \frac{3}{4}$ to achieve an ϵ -additive error is $\mathcal{O}\left(\sqrt{\kappa n} \log \frac{F(\tilde{x}_0) - F(x^*)}{\epsilon}\right)$.

(II) For another case with $\frac{m}{\kappa} > \frac{3}{4}$, we set $\eta = \frac{2}{3L}$, $\theta = \frac{1}{2}$, and $m = \Theta(n)$.

Again, we evaluate the constraint (7) first. By substituting the parameter setting, the constraint becomes

$$L\theta + \frac{L\theta}{1-\theta} \leq \frac{1}{\eta} \rightarrow \eta \leq \frac{2}{3L}.$$

Thus the constraint is satisfied by our parameter choice.

Substituting the parameter setting into (12), we get

$$2 \sum_{j=0}^{m-1} \omega^j (\mathbb{E}[F(\tilde{x}_s)] - F(x^*)) + \frac{3L\omega^m}{4} \mathbb{E}[\|x_m - x^*\|^2] \leq \sum_{j=0}^{m-1} \omega^j (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{3L}{4} \|x_0 - x^*\|^2.$$

Notice that based on the Bernoulli's inequality, $\omega^m = (1 + \frac{2}{3\kappa})^m \geq 1 + \frac{2m}{3\kappa} \geq \frac{3}{2}$, which leads to

$$\begin{aligned} \frac{3}{2} \sum_{j=0}^{m-1} \omega^j (\mathbb{E}[F(\tilde{x}_s)] - F(x^*)) + \frac{9L}{8} \mathbb{E}[\|x_m - x^*\|^2] &\leq \sum_{j=0}^{m-1} \omega^j (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{3L}{4} \|x_0 - x^*\|^2, \\ \sum_{j=0}^{m-1} \omega^j (\mathbb{E}[F(\tilde{x}_s)] - F(x^*)) + \frac{3L}{4} \mathbb{E}[\|x_m - x^*\|^2] &\leq \frac{2}{3} \cdot \left(\sum_{j=0}^{m-1} \omega^j (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{3L}{4} \|x_0 - x^*\|^2 \right). \end{aligned}$$

Again, by telescoping the above inequality from $s = 1 \dots S$, we get

$$\sum_{j=0}^{m-1} \omega^j (\mathbb{E}[F(\tilde{x}_S)] - F(x^*)) + \frac{3L}{4} \mathbb{E}[\|x_m^S - x^*\|^2] \leq \left(\frac{2}{3}\right)^S \cdot \left(\sum_{j=0}^{m-1} \omega^j (F(\tilde{x}_0) - F(x^*)) + \frac{3L}{4} \|x_0^1 - x^*\|^2 \right).$$

Since $\sum_{j=0}^{m-1} \omega^j \geq m$, the above inequality can be rewritten as follows:

$$\begin{aligned} (\mathbb{E}[F(\tilde{x}_S)] - F(x^*)) &\stackrel{(*)}{\leq} \left(\frac{2}{3}\right)^S \cdot \left(1 + \frac{3\kappa}{2m}\right) \cdot (F(\tilde{x}_0) - F(x^*)) \\ &\leq \left(\frac{2}{3}\right)^S \cdot O(F(\tilde{x}_0) - F(x^*)), \end{aligned}$$

where $(*)$ uses the σ -strongly convexity of $F(\cdot)$, that is, $\|x_0^1 - x^*\|^2 \leq \frac{2}{\sigma} (F(x_0^1) - F(x^*))$.

This result implies that the oracle complexity in this case is $\mathcal{O}\left(n \log \frac{F(\tilde{x}_0) - F(x^*)}{\epsilon}\right)$.

B.1.1. DETAILED COMPARISON BETWEEN MiG AND KATYUSHA

As mentioned in Section 3.1.1, MiG is a special case of Katyusha. However, it is a non-trivial special case and there are some significant differences from the theoretical perspective between them for the ill-conditioned problems:

The intuition of Katyusha is that Katyusha Momentum is used to further reduce the variance of the iterates so as to make Nesterov’s Momentum effective, which can be clearly seen from the parameter choice: the author fixed Katyusha Momentum (for \tilde{x} , in the notation of original work) as $\tau_2 = 1/2$, and chose Nesterov’s Momentum (for y and z) with $(1/2 - \tau_1, \tau_1)$. This idea is similar to Acc-Prox-SVRG (Nitanda, 2014), which uses sufficiently large mini-batch to make the Nesterov’s Momentum effective. Thus, the acceleration comes from Nesterov’s Momentum, which is common in previous work (e.g. Catalyst).

In contrast, MiG uses only Katyusha Momentum (negative momentum) to achieve acceleration (i.e., choosing τ_2 in a dynamic way), which is somehow counter-intuitive. To the best of our knowledge, this is the first work to yield an acceleration without using Nesterov’s Momentum (which is at the heart of all existing accelerated first-order methods, as stated in (Allen-Zhu, 2017)).

B.2. Proof of Theorem 2

Again, we first impose the following constraint on η and θ :

$$L\theta + \frac{L\theta}{1-\theta} \leq \frac{1}{\eta}, \text{ or equivalently } \eta \leq \frac{1-\theta}{L\theta(2-\theta)}. \quad (13)$$

We start with convexity of $f(\cdot)$ at y_{j-1} . By definition, for any vector $u \in \mathbb{R}^d$, we have

$$\begin{aligned} f(y_{j-1}) - f(u) &\leq \langle \nabla f(y_{j-1}), y_{j-1} - u \rangle \\ &\stackrel{(*)}{=} \frac{1-\theta}{\theta} \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \langle \nabla f(y_{j-1}), x_{j-1} - u \rangle, \end{aligned} \quad (14)$$

where $(*)$ follows from the fact that $y_{j-1} = \theta x_{j-1} + (1-\theta)\tilde{x}_{s-1}$.

Then we further expand $\langle \nabla f(y_{j-1}), x_{j-1} - u \rangle$ as

$$\langle \nabla f(y_{j-1}), x_{j-1} - u \rangle = \langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_{j-1} - u \rangle + \langle \tilde{\nabla}, x_{j-1} - x_j \rangle + \langle \tilde{\nabla}, x_j - u \rangle. \quad (15)$$

Using L -smooth (2) of $f(\cdot)$ at (y_j, y_{j-1}) , we get

$$\begin{aligned} f(y_j) - f(y_{j-1}) &\leq \langle \nabla f(y_{j-1}), y_j - y_{j-1} \rangle + \frac{L}{2} \|y_j - y_{j-1}\|^2 \\ &= \theta [\langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - x_{j-1} \rangle + \langle \tilde{\nabla}, x_j - x_{j-1} \rangle] + \frac{L\theta^2}{2} \|x_j - x_{j-1}\|^2, \\ \langle \tilde{\nabla}, x_{j-1} - x_j \rangle &\leq \frac{1}{\theta} (f(y_{j-1}) - f(y_j)) + \langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - x_{j-1} \rangle + \frac{L\theta}{2} \|x_j - x_{j-1}\|^2. \end{aligned}$$

Using the constraint (13), we have

$$\langle \tilde{\nabla}, x_{j-1} - x_j \rangle \leq \frac{1}{\theta} (f(y_{j-1}) - f(y_j)) + \langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - x_{j-1} \rangle + \frac{1}{2\eta} \|x_j - x_{j-1}\|^2 - \frac{L\theta}{2(1-\theta)} \|x_j - x_{j-1}\|^2. \quad (16)$$

Then we can combine (14), (15), (16), as well as Lemma 3 (with $g(x)$ convex), which leads to

$$\begin{aligned} f(y_{j-1}) - f(u) &\leq \frac{1-\theta}{\theta} \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \langle \nabla f(y_{j-1}) - \tilde{\nabla}, x_j - u \rangle + \frac{1}{\theta} (f(y_{j-1}) - f(y_j)) \\ &\quad - \frac{L\theta}{2(1-\theta)} \|x_j - x_{j-1}\|^2 + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1}{2\eta} \|x_j - u\|^2 + g(u) - g(x_j). \end{aligned}$$

After taking expectation with respect to the sample i_j , we get

$$\begin{aligned} f(y_{j-1}) - f(u) &\leq \frac{1-\theta}{\theta} \langle \nabla f(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \frac{1}{2\beta} \mathbb{E}_{i_j} [\|\nabla f(y_{j-1}) - \tilde{\nabla}\|^2] + \frac{\beta}{2} \mathbb{E}_{i_j} [\|x_j - x_{j-1}\|^2] \\ &\quad + \frac{1}{\theta} (f(y_{j-1}) - \mathbb{E}_{i_j} [f(y_j)]) - \frac{L\theta}{2(1-\theta)} \mathbb{E}_{i_j} [\|x_j - x_{j-1}\|^2] + \frac{1}{2\eta} \|x_{j-1} - u\|^2 \\ &\quad - \frac{1}{2\eta} \mathbb{E}_{i_j} [\|x_j - u\|^2] + g(u) - \mathbb{E}_{i_j} [g(x_j)], \end{aligned}$$

which the inequality uses the unbiasedness of the gradient estimator and the Young’s inequality with the parameter $\beta > 0$.

Using Lemma 1 to bound the variance term and choosing $\beta = \frac{L\theta}{1-\theta}$, the above inequality becomes

$$\frac{1}{\theta} \mathbb{E}_{i_j} [f(y_j)] - F(u) \leq \frac{1-\theta}{\theta} f(\tilde{x}_{s-1}) + \frac{1}{2\eta} \|x_{j-1} - u\|^2 - \frac{1}{2\eta} \mathbb{E}_{i_j} [\|x_j - u\|^2] - \mathbb{E}_{i_j} [g(x_j)].$$

Let $u = x^*$ and using the fact that $-g(x_j) \leq \frac{1-\theta}{\theta} g(\tilde{x}_{s-1}) - \frac{1}{\theta} g(y_j)$, we get

$$\frac{1}{\theta} \mathbb{E}_{i_j} [F(y_j) - F(x^*)] \leq \frac{1-\theta}{\theta} (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{1}{2\eta} \|x_{j-1} - x^*\|^2 - \frac{1}{2\eta} \mathbb{E}_{i_j} [\|x_j - x^*\|^2].$$

Summing the above inequality over $j = 1 \dots m$ and taking expectation with respect to all randomness in this epoch, we obtain

$$\begin{aligned} \frac{1}{\theta} \mathbb{E} \left[\frac{1}{m} \sum_{j=1}^m F(y_j) - F(x^*) \right] &\stackrel{(a)}{\leq} \frac{1-\theta}{\theta} (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{2\theta L}{m} \|x_0 - x^*\|^2 - \frac{2\theta L}{m} \mathbb{E} [\|x_m - x^*\|^2], \\ \frac{1}{\theta^2} \mathbb{E} [F(\tilde{x}_s) - F(x^*)] &\stackrel{(b)}{\leq} \frac{1-\theta}{\theta^2} (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{2L}{m} \|x_0 - x^*\|^2 - \frac{2L}{m} \mathbb{E} [\|x_m - x^*\|^2], \end{aligned} \quad (17)$$

where (a) follows from our parameter choice of $\eta = \frac{1}{4L\theta}$, and (b) uses the Jensen’s inequality.

At this point, we first check the constraint (13),

$$\eta \leq \frac{1-\theta}{L\theta(2-\theta)} \rightarrow \theta \leq \frac{2}{3},$$

which is satisfied by the parameter setting $\theta = \frac{2}{s+4} \leq \frac{1}{2}$.

Then by evaluating (17) with $s = 1$, we have

$$\frac{1}{\theta_1^2} \mathbb{E} [F(\tilde{x}_1) - F(x^*)] + \frac{2L}{m} \mathbb{E} [\|x_m^1 - x^*\|^2] \leq \frac{1-\theta_1}{\theta_1^2} (F(\tilde{x}_0) - F(x^*)) + \frac{2L}{m} \|x_0^1 - x^*\|^2.$$

Notice that here we mark iterates and θ with epoch number to prevent confusion.

Since $\theta_s = \frac{2}{s+4}$, one can easily verify that $\frac{1-\theta_{s+1}}{\theta_{s+1}^2} \leq \frac{1}{\theta_s^2}$. Thus we can telescope the above inequality from $s = 1 \dots \mathcal{S}$,

$$\begin{aligned} \frac{1}{\theta_{\mathcal{S}}^2} \mathbb{E} [F(\tilde{x}_{\mathcal{S}}) - F(x^*)] + \frac{2L}{m} \mathbb{E} [\|x_m^{\mathcal{S}} - x^*\|^2] &\leq \frac{1-\theta_1}{\theta_1^2} (F(\tilde{x}_0) - F(x^*)) + \frac{2L}{m} \|x_0^1 - x^*\|^2, \\ \mathbb{E} [F(\tilde{x}_{\mathcal{S}}) - F(x^*)] &\leq \frac{4(1-\theta_1)}{(\mathcal{S}+4)^2 \theta_1^2} (F(\tilde{x}_0) - F(x^*)) + \frac{8L}{(\mathcal{S}+4)^2 m} \|x_0^1 - x^*\|^2. \end{aligned}$$

In other words, by choosing $m = \Theta(n)$, the total oracle complexity is $\mathcal{O}(n \sqrt{\frac{F(\tilde{x}_0) - F(x^*)}{\epsilon}} + \sqrt{\frac{nL \|x_0^1 - x^*\|^2}{\epsilon}})$.

C. Proofs for Section 4

Lemma 4 (Sparse Variance Bound 1). *If Assumption 3 holds, for any $y, \tilde{x} \in \mathbb{R}^d$ and the sample i_j , denote $\tilde{\nabla} = \nabla f_{i_j}(y) - \nabla f_{i_j}(\tilde{x}) + D_{i_j} \nabla F(\tilde{x})$, where D_{i_j} is defined in Section 4.1, we can bound the variance $\mathbb{E}_{i_j} [\|\tilde{\nabla}\|^2]$ as*

$$\mathbb{E}_{i_j} [\|\tilde{\nabla}\|^2] \leq 4L(F(y) - F(x^*)) + 4L(F(\tilde{x}) - F(x^*)).$$

Proof. From Lemma 10 in (Mania et al., 2017), we have

$$\mathbb{E}_{i_j} [\|\tilde{\nabla}\|^2] \leq 2\mathbb{E}_{i_j} [\|\nabla f_{i_j}(y) - \nabla f_{i_j}(x^*)\|^2] + 2\mathbb{E}_{i_j} [\|\nabla f_{i_j}(\tilde{x}) - \nabla f_{i_j}(x^*)\|^2], \quad (18)$$

which provides an upper bound for the variance of the sparse stochastic variance reduced gradient estimator.

From Theorem 2.1.5 in (Nesterov, 2004), we have

$$\|\nabla f_{i_j}(y) - \nabla f_{i_j}(x^*)\|^2 \leq 2L(f_{i_j}(y) - f_{i_j}(x^*) - \langle \nabla f_{i_j}(x^*), y - x^* \rangle).$$

Taking expectation with the sample i_j , the above inequality becomes

$$\mathbb{E}_{i_j} [\|\nabla f_{i_j}(y) - \nabla f_{i_j}(x^*)\|^2] \leq 2L(F(y) - F(x^*)).$$

Similarly, we have

$$\mathbb{E}_{i_j} [\|\nabla f_{i_j}(\tilde{x}) - \nabla f_{i_j}(x^*)\|^2] \leq 2L(F(\tilde{x}) - F(x^*)).$$

Substituting the above inequalities into (18) yields the desired result. \square

Lemma 5 (Sparse Variance Bound 2). *If Assumption 3 holds, using same notations as in Algorithm 3, we can bound the variance as*

$$\begin{aligned} \mathbb{E}_{i_j} [\|\nabla F(y_{j-1}) - \tilde{\nabla}_S\|^2] &\leq 2L(F(\tilde{x}_{s-1}) - F(y_{j-1}) - \langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle) \\ &\quad + 2L(D_m^2 - D_m) \cdot (F(\tilde{x}_{s-1}) - F(x^*)), \end{aligned}$$

where $D_m = \max_{k=1\dots d} \frac{1}{p_k}$, i.e., the inverse probability of the most sparse coordinates.

Proof. First, we expand the variance term as

$$\begin{aligned} &\mathbb{E}_{i_j} [\|\nabla F(y_{j-1}) - \tilde{\nabla}_S\|^2] \\ &= \mathbb{E}_{i_j} \left[\left\| (\nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1})) - (\nabla F(y_{j-1}) - D_{i_j} \nabla F(\tilde{x}_{s-1})) \right\|^2 \right] \\ &= \mathbb{E}_{i_j} [\|\nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1})\|^2] - 2\mathbb{E}_{i_j} [\langle \nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1}), \nabla F(y_{j-1}) - D_{i_j} \nabla F(\tilde{x}_{s-1}) \rangle] \\ &\quad + \mathbb{E}_{i_j} [\|\nabla F(y_{j-1}) - D_{i_j} \nabla F(\tilde{x}_{s-1})\|^2]. \end{aligned} \quad (19)$$

Then we focus on the last two terms in (19),

$$\begin{aligned} &\mathbb{E}_{i_j} [\|\nabla F(y_{j-1}) - D_{i_j} \nabla F(\tilde{x}_{s-1})\|^2] \\ &= \|\nabla F(y_{j-1})\|^2 - 2\mathbb{E}_{i_j} [\langle \nabla F(y_{j-1}), D_{i_j} \nabla F(\tilde{x}_{s-1}) \rangle] + \mathbb{E}_{i_j} [\|D_{i_j} \nabla F(\tilde{x}_{s-1})\|^2] \\ &\stackrel{(\star)}{=} \|\nabla F(y_{j-1})\|^2 - 2\langle \nabla F(y_{j-1}), \nabla F(\tilde{x}_{s-1}) \rangle + \langle \nabla F(\tilde{x}_{s-1}), D \nabla F(\tilde{x}_{s-1}) \rangle, \end{aligned} \quad (20)$$

where (\star) uses the fact that $P_{i_j} = P_{i_j} \cdot P_{i_j}$ and the unbiased sparse estimator $\mathbb{E}_{i_j} [D_{i_j} \nabla F(x)] = \nabla F(x)$.

$$\begin{aligned} &\mathbb{E}_{i_j} [\langle \nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1}), \nabla F(y_{j-1}) - D_{i_j} \nabla F(\tilde{x}_{s-1}) \rangle] \\ &= \|\nabla F(y_{j-1})\|^2 - \langle \nabla F(y_{j-1}), \nabla F(\tilde{x}_{s-1}) \rangle - \mathbb{E}_{i_j} [\langle \nabla f_{i_j}(y_{j-1}), D_{i_j} \nabla F(\tilde{x}_{s-1}) \rangle] \\ &\quad + \mathbb{E}_{i_j} [\langle \nabla f_{i_j}(\tilde{x}_{s-1}), D_{i_j} \nabla F(\tilde{x}_{s-1}) \rangle] \\ &\stackrel{(\star)}{=} \|\nabla F(y_{j-1})\|^2 - \langle \nabla F(y_{j-1}), \nabla F(\tilde{x}_{s-1}) \rangle - \langle \nabla F(y_{j-1}), D \nabla F(\tilde{x}_{s-1}) \rangle + \langle \nabla F(\tilde{x}_{s-1}), D \nabla F(\tilde{x}_{s-1}) \rangle, \end{aligned} \quad (21)$$

where (\star) follows from $\langle \nabla f_{i_j}(y_{j-1}), D_{i_j} \nabla F(\tilde{x}_{s-1}) \rangle = \langle \nabla f_{i_j}(y_{j-1}), D \nabla F(\tilde{x}_{s-1}) \rangle$, since $\nabla f_{i_j}(y_{j-1})$ and D_{i_j} are both supported on the sample i_j .

By substituting (20) and (21) into (19), we get

$$\begin{aligned} &\mathbb{E}_{i_j} [\|\nabla F(y_{j-1}) - \tilde{\nabla}_S\|^2] \\ &= \mathbb{E}_{i_j} [\|\nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1})\|^2] - \|\nabla F(y_{j-1})\|^2 + 2\langle \nabla F(y_{j-1}), D \nabla F(\tilde{x}_{s-1}) \rangle - \|D \nabla F(\tilde{x}_{s-1})\|^2 \\ &\quad + \langle (D - I) \nabla F(\tilde{x}_{s-1}), D \nabla F(\tilde{x}_{s-1}) \rangle \\ &= \mathbb{E}_{i_j} [\|\nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1})\|^2] - \|\nabla F(y_{j-1}) - D \nabla F(\tilde{x}_{s-1})\|^2 + \langle (D - I) \nabla F(\tilde{x}_{s-1}), D \nabla F(\tilde{x}_{s-1}) \rangle \\ &\leq \mathbb{E}_{i_j} [\|\nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1})\|^2] + \langle (D - I) \nabla F(\tilde{x}_{s-1}), D \nabla F(\tilde{x}_{s-1}) \rangle. \end{aligned} \quad (22)$$

From Theorem 2.1.5 in (Nesterov, 2004), we have

$$\|\nabla f_{i_j}(y_{j-1}) - \nabla f_{i_j}(\tilde{x}_{s-1})\|^2 \leq 2L(f_{i_j}(\tilde{x}_{s-1}) - f_{i_j}(y_{j-1}) - \langle \nabla f_{i_j}(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle). \quad (23)$$

By substituting (23) into (22), we obtain

$$\begin{aligned} & \mathbb{E}_{i_j} [\|\nabla F(y_{j-1}) - \tilde{\nabla}_S\|^2] \\ & \leq 2L(F(\tilde{x}_{s-1}) - F(y_{j-1}) - \langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle) + \langle (D - I)\nabla F(\tilde{x}_{s-1}), D\nabla F(\tilde{x}_{s-1}) \rangle. \end{aligned}$$

Since $D_m = \max_{k=1\dots d} \frac{1}{p_k}$, we can bound the last term as follows:

$$\begin{aligned} & \langle (D - I)\nabla F(\tilde{x}_{s-1}), D\nabla F(\tilde{x}_{s-1}) \rangle \\ & = \sum_{k=1}^d \left(\frac{1}{p_k^2} - \frac{1}{p_k} \right) [\nabla F(\tilde{x}_{s-1})]_k^2 \\ & \leq (D_m^2 - D_m) \|\nabla F(\tilde{x}_{s-1}) - \nabla F(x^*)\|^2 \\ & \stackrel{(*)}{\leq} 2L(D_m^2 - D_m) \cdot (F(\tilde{x}_{s-1}) - F(x^*)), \end{aligned}$$

where $(*)$ follows from applying Theorem 2.1.5 in (Nesterov, 2004). \square

C.1. Proof of Theorem 3

For the purpose of analyzing in the asynchronous setting, we need to analyze the convergence of MiG starting with iterate difference.

We start with the iterate difference between y_j and x^* . By expanding iterate difference and taking expectation with respect to the sample i_j , we get

$$\begin{aligned} \mathbb{E}_{i_j} [\|y_j - x^*\|^2] & \stackrel{(a)}{=} \mathbb{E}_{i_j} [\|\theta(x_{j-1} - \eta \cdot \tilde{\nabla}_S) + (1 - \theta)\tilde{x}_{s-1} - x^*\|^2] \\ & = \mathbb{E}_{i_j} [\|y_{j-1} - \eta\theta \cdot \tilde{\nabla}_S - x^*\|^2] \\ & \stackrel{(b)}{=} \|y_{j-1} - x^*\|^2 - 2\eta\theta \langle \nabla F(y_{j-1}), y_{j-1} - x^* \rangle + \eta^2\theta^2 \mathbb{E}_{i_j} [\|\tilde{\nabla}_S\|^2], \end{aligned} \quad (24)$$

where (a) uses the definition of y , and (b) follows from the unbiasedness of the sparse gradient estimator $\mathbb{E}_{i_j} [\tilde{\nabla}_S] = \nabla F(y_{j-1})$.

Using σ -strongly convex, we get a bound for $-\langle \nabla F(y_{j-1}), y_{j-1} - x^* \rangle$ as

$$\langle \nabla F(y_{j-1}), y_{j-1} - x^* \rangle \geq F(y_{j-1}) - F(x^*) + \frac{\sigma}{2} \|y_{j-1} - x^*\|^2. \quad (25)$$

Using Lemma 4, we have the following variance bound:

$$\mathbb{E}_{i_j} [\|\tilde{\nabla}_S\|^2] \leq 4L(F(y_{j-1}) - F(x^*)) + 4L(F(\tilde{x}_{s-1}) - F(x^*)). \quad (26)$$

By combining (24), (25) and (26), we get

$$\begin{aligned} \mathbb{E}_{i_j} [\|y_j - x^*\|^2] & \leq \|y_{j-1} - x^*\|^2 - \eta\theta\sigma \|y_{j-1} - x^*\|^2 - 2\eta\theta(F(y_{j-1}) - F(x^*)) + \eta^2\theta^2 \mathbb{E}_{i_j} [\|\tilde{\nabla}_S\|^2] \\ & \leq (1 - \eta\theta\sigma) \cdot \|y_{j-1} - x^*\|^2 + (4L\eta^2\theta^2 - 2\eta\theta) \cdot (F(y_{j-1}) - F(x^*)) \\ & \quad + 4L\eta^2\theta^2(F(\tilde{x}_{s-1}) - F(x^*)), \\ & \quad (2\eta\theta - 4L\eta^2\theta^2) \cdot (F(y_{j-1}) - F(x^*)) \\ & \leq \|y_{j-1} - x^*\|^2 - \mathbb{E}_{i_j} [\|y_j - x^*\|^2] + 4L\eta^2\theta^2(F(\tilde{x}_{s-1}) - F(x^*)). \end{aligned}$$

Summing the above inequality over $j = 1 \dots m$ and taking expectation with respect to all randomness in this epoch, we get

$$\begin{aligned} & (2\eta\theta - 4L\eta^2\theta^2) \cdot \mathbb{E} \left[\sum_{j=1}^m (F(y_{j-1}) - F(x^*)) \right] \\ & \leq \|y_0 - x^*\|^2 - \mathbb{E} [\|y_m - x^*\|^2] + 4L\eta^2\theta^2 m (F(\tilde{x}_{s-1}) - F(x^*)). \end{aligned}$$

Using Jensen’s inequality and $\tilde{x}_s = \frac{1}{m} \sum_{j=0}^{m-1} y_j$, we obtain

$$\begin{aligned} (2\eta\theta - 4L\eta^2\theta^2)m \cdot \mathbb{E}[(F(\tilde{x}_s) - F(x^*))] &\leq \|y_0 - x^*\|^2 + 4L\eta^2\theta^2m(F(\tilde{x}_{s-1}) - F(x^*)), \\ \mathbb{E}[(F(\tilde{x}_s) - F(x^*))] &\stackrel{(\star)}{\leq} \frac{\frac{2}{\sigma} + 4L\eta^2\theta^2m}{(2\eta\theta - 4L\eta^2\theta^2)m} \cdot (F(\tilde{x}_{s-1}) - F(x^*)), \end{aligned}$$

where (\star) follows from the fact $x_0 = \tilde{x}_{s-1} = y_0$ and σ -strongly convexity of $F(\cdot)$.

By choosing $m = 25\kappa$, $\eta = \frac{1}{L}$, $\theta = \frac{1}{10}$, the above inequality becomes

$$\mathbb{E}[(F(\tilde{x}_s) - F(x^*))] \leq 0.75 \cdot (F(\tilde{x}_{s-1}) - F(x^*)),$$

which means that the total oracle complexity is $\mathcal{O}((n + \kappa) \log \frac{F(\tilde{x}_0) - F(x^*)}{\epsilon})$.

C.2. Proof of Theorem 4

Here we consider improving the convergence rate for the Serial Sparse MiG, referencing to Algorithm 3 with Option II. Our analysis is based on function difference similar to the proofs in Appendix B, which make the parameter θ effective.

Similarly, we first add the following constraint:

$$L\theta + \frac{L\theta}{1 - \theta} \leq \frac{1}{\eta}. \quad (27)$$

We start with the convexity of $F(\cdot)$ at y_{j-1} and x^* ,

$$\begin{aligned} F(y_{j-1}) - F(x^*) &\leq \langle \nabla F(y_{j-1}), y_{j-1} - x^* \rangle \\ &= \frac{1 - \theta}{\theta} \langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \langle \nabla F(y_{j-1}), x_{j-1} - x^* \rangle. \end{aligned} \quad (28)$$

Then we further expand $\langle \nabla F(y_{j-1}), x_{j-1} - x^* \rangle$ as

$$\langle \nabla F(y_{j-1}), x_{j-1} - x^* \rangle = \langle \nabla F(y_{j-1}) - \tilde{\nabla}_S, x_{j-1} - x^* \rangle + \langle \tilde{\nabla}_S, x_{j-1} - x_j \rangle + \langle \tilde{\nabla}_S, x_j - x^* \rangle. \quad (29)$$

Using L -smooth (2) of $F(\cdot)$ at (y_j, y_{j-1}) , we can bound $\langle \tilde{\nabla}_S, x_{j-1} - x_j \rangle$ as

$$\begin{aligned} F(y_j) - F(y_{j-1}) &\leq \langle \nabla F(y_{j-1}), y_j - y_{j-1} \rangle + \frac{L}{2} \|y_j - y_{j-1}\|^2 \\ &\stackrel{(\star)}{=} \theta [\langle \nabla F(y_{j-1}) - \tilde{\nabla}_S, x_j - x_{j-1} \rangle + \langle \tilde{\nabla}_S, x_j - x_{j-1} \rangle] + \frac{L\theta^2}{2} \|x_j - x_{j-1}\|^2, \\ \langle \tilde{\nabla}_S, x_{j-1} - x_j \rangle &\leq \frac{1}{\theta} (F(y_{j-1}) - F(y_j)) + \langle \nabla F(y_{j-1}) - \tilde{\nabla}_S, x_j - x_{j-1} \rangle + \frac{L\theta}{2} \|x_j - x_{j-1}\|^2, \end{aligned}$$

where (\star) uses the definition of y .

Using the constraint (27), we have

$$\begin{aligned} \langle \tilde{\nabla}_S, x_{j-1} - x_j \rangle &\leq \frac{1}{\theta} (F(y_{j-1}) - F(y_j)) + \langle \nabla F(y_{j-1}) - \tilde{\nabla}_S, x_j - x_{j-1} \rangle + \frac{1}{2\eta} \|x_j - x_{j-1}\|^2 \\ &\quad - \frac{L\theta}{2(1 - \theta)} \|x_j - x_{j-1}\|^2. \end{aligned} \quad (30)$$

Since we can write an equivalent update as $x_j = \arg \min_x \{\frac{1}{2} \|x - x_{j-1}\|^2 + \eta \langle \tilde{\nabla}_S, x \rangle\}$, by applying Lemma 2 with $z = x^*$, $z_0 = x_{j-1}$, $z^* = x_j$, $\tau = \frac{1}{\eta}$, $\psi(x) = \langle \tilde{\nabla}_S, x \rangle$, we have

$$\langle \tilde{\nabla}_S, x_j - x^* \rangle = -\frac{1}{2\eta} \|x_{j-1} - x_j\|^2 + \frac{1}{2\eta} \|x_{j-1} - x^*\|^2 - \frac{1}{2\eta} \|x_j - x^*\|^2. \quad (31)$$

Combining (28), (29), (30) and (31), we get

$$F(y_{j-1}) - F(x^*) \leq \frac{1-\theta}{\theta} \langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \frac{1}{\theta} (F(y_{j-1}) - F(y_j)) + \langle \nabla F(y_{j-1}) - \tilde{\nabla}_S, x_j - x^* \rangle - \frac{L\theta}{2(1-\theta)} \|x_j - x_{j-1}\|^2 + \frac{1}{2\eta} \|x_{j-1} - x^*\|^2 - \frac{1}{2\eta} \|x_j - x^*\|^2.$$

Taking expectation with respect to the sample i_j , we obtain

$$\begin{aligned} F(y_{j-1}) - F(x^*) &\stackrel{(a)}{\leq} \frac{1-\theta}{\theta} \langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \frac{1}{\theta} (F(y_{j-1}) - \mathbb{E}_{i_j}[F(y_j)]) - \frac{L\theta}{2(1-\theta)} \mathbb{E}_{i_j}[\|x_j - x_{j-1}\|^2] \\ &\quad + \mathbb{E}_{i_j}[\langle \nabla F(y_{j-1}) - \tilde{\nabla}_S, x_j - x_{j-1} \rangle] + \frac{1}{2\eta} \|x_{j-1} - x^*\|^2 - \frac{1}{2\eta} \mathbb{E}_{i_j}[\|x_j - x^*\|^2] \\ &\stackrel{(b)}{\leq} \frac{1-\theta}{\theta} \langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \frac{1}{\theta} (F(y_{j-1}) - \mathbb{E}_{i_j}[F(y_j)]) - \frac{L\theta}{2(1-\theta)} \mathbb{E}_{i_j}[\|x_j - x_{j-1}\|^2] \\ &\quad + \frac{1}{2\beta} \mathbb{E}_{i_j}[\|\nabla F(y_{j-1}) - \tilde{\nabla}_S\|^2] + \frac{\beta}{2} \mathbb{E}_{i_j}[\|x_j - x_{j-1}\|^2] + \frac{1}{2\eta} \|x_{j-1} - x^*\|^2 - \frac{1}{2\eta} \mathbb{E}_{i_j}[\|x_j - x^*\|^2], \end{aligned}$$

where (a) uses the unbiasedness of the sparse gradient estimator $\mathbb{E}_{i_j}[\tilde{\nabla}_S] = \nabla F(y_{j-1})$, and (b) uses the Young's inequality with the parameter $\beta > 0$.

Using Lemma 5 to bound the variance term, we get

$$\begin{aligned} F(y_{j-1}) - F(x^*) &\leq \frac{1-\theta}{\theta} \langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle + \frac{1}{\theta} (F(y_{j-1}) - \mathbb{E}_{i_j}[F(y_j)]) - \frac{L\theta}{2(1-\theta)} \mathbb{E}_{i_j}[\|x_j - x_{j-1}\|^2] \\ &\quad + \frac{L}{\beta} (F(\tilde{x}_{s-1}) - F(y_{j-1}) - \langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle) + \frac{L(D_m^2 - D_m)}{\beta} (F(\tilde{x}_{s-1}) - F(x^*)) \\ &\quad + \frac{\beta}{2} \mathbb{E}_{i_j}[\|x_j - x_{j-1}\|^2] + \frac{1}{2\eta} \|x_{j-1} - x^*\|^2 - \frac{1}{2\eta} \mathbb{E}_{i_j}[\|x_j - x^*\|^2]. \end{aligned}$$

By choosing $\beta = \frac{L\theta}{1-\theta} > 0$, the above inequality becomes

$$\begin{aligned} F(y_{j-1}) - F(x^*) &\leq \frac{1}{\theta} (F(y_{j-1}) - \mathbb{E}_{i_j}[F(y_j)]) + \frac{1-\theta}{\theta} (F(\tilde{x}_{s-1}) - F(y_{j-1})) \\ &\quad + \frac{1-\theta}{\theta} (D_m^2 - D_m) \cdot (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{1}{2\eta} \|x_{j-1} - x^*\|^2 - \frac{1}{2\eta} \mathbb{E}_{i_j}[\|x_j - x^*\|^2], \\ \frac{1}{\theta} (\mathbb{E}_{i_j}[F(y_j)] - F(x^*)) &\leq \frac{1-\theta}{\theta} (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{1-\theta}{\theta} (D_m^2 - D_m) \cdot (F(\tilde{x}_{s-1}) - F(x^*)) \\ &\quad + \frac{1}{2\eta} \|x_{j-1} - x^*\|^2 - \frac{1}{2\eta} \mathbb{E}_{i_j}[\|x_j - x^*\|^2]. \end{aligned}$$

Summing the above inequality over $j = 1 \dots m$ and taking expectation with respect to all randomness in this epoch, we get

$$\begin{aligned} &\frac{1}{\theta} \left(\frac{1}{m} \sum_{j=1}^m \mathbb{E}[F(y_j)] - F(x^*) \right) \\ &\leq \frac{1-\theta}{\theta} (1 + D_m^2 - D_m) \cdot (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{1}{2\eta m} \left(\|x_0 - x^*\|^2 - \mathbb{E}[\|x_m - x^*\|^2] \right). \end{aligned}$$

Using Jensen's inequality and $\tilde{x}_s = \frac{1}{m} \sum_{j=1}^m y_j$ in Option II, we have

$$\begin{aligned} &\frac{1}{\theta} \mathbb{E}[F(\tilde{x}_s) - F(x^*)] \\ &\leq \frac{1-\theta}{\theta} (1 + D_m^2 - D_m) \cdot (F(\tilde{x}_{s-1}) - F(x^*)) + \frac{1}{2\eta m} \left(\|x_0 - x^*\|^2 - \mathbb{E}[\|x_m - x^*\|^2] \right). \end{aligned}$$

In order to give a clean proof, we set $\tilde{D}_s \triangleq F(\tilde{x}_s) - F(x^*)$, $P_0^s \triangleq \|x_0^s - x^*\|^2$ and the terms that related to D_m as $\zeta = D_m^2 - D_m$. Then the inequality becomes

$$\begin{aligned} \frac{1}{\theta} \mathbb{E}[\tilde{D}_s] &\leq \frac{1-\theta}{\theta}(1+\zeta) \cdot \tilde{D}_{s-1} + \frac{1}{2\eta m} \mathbb{E}[P_0^s - P_m^s], \\ \frac{\theta + \theta\zeta - \zeta}{\theta} \cdot \mathbb{E}[\tilde{D}_s] &\leq \frac{1-\theta}{\theta}(1+\zeta) \cdot \mathbb{E}[\tilde{D}_{s-1} - \tilde{D}_s] + \frac{1}{2\eta m} \mathbb{E}[P_0^s - P_m^s]. \end{aligned}$$

Suppose we run \mathcal{S} epochs before a restart. By summing the above inequality over $s = 1 \dots \mathcal{S}$, we get

$$\frac{\theta + \theta\zeta - \zeta}{\theta} \cdot \sum_{s=1}^{\mathcal{S}} \mathbb{E}[\tilde{D}_s] \leq \frac{1-\theta}{\theta}(1+\zeta) \cdot \mathbb{E}[\tilde{D}_0 - \tilde{D}_{\mathcal{S}}] + \frac{1}{2\eta m} \mathbb{E}[P_0^1 - P_m^{\mathcal{S}}].$$

Choosing the initial vector for next \mathcal{S} epochs as $x_0^{new} = \frac{1}{\mathcal{S}} \sum_{s=1}^{\mathcal{S}} \tilde{x}_s$, we have

$$\begin{aligned} \mathcal{S} \cdot \frac{\theta + \theta\zeta - \zeta}{\theta} \cdot \mathbb{E}[\tilde{D}_0^{new}] &\leq \frac{1-\theta}{\theta}(1+\zeta) \cdot \mathbb{E}[\tilde{D}_0 - \tilde{D}_{\mathcal{S}}] + \frac{1}{2\eta m} \mathbb{E}[P_0^1 - P_m^{\mathcal{S}}] \\ &\stackrel{(\star)}{\leq} \left(\frac{(1-\theta)(1+\zeta)}{\theta} + \frac{1}{\eta m \sigma} \right) \cdot \tilde{D}_0, \\ \mathbb{E}[\tilde{D}_0^{new}] &\leq \frac{(1-\theta) \cdot (1+\zeta) + \frac{\theta}{\eta m \sigma}}{\mathcal{S} \cdot (\theta + \zeta\theta - \zeta)} \tilde{D}_0, \end{aligned}$$

where (\star) uses the σ -strongly convexity of $F(\cdot)$ and $\tilde{x}_0 = x_0^1$, that is, $P_0^1 \leq \frac{2}{\sigma} \tilde{D}_0$.

Setting $\mathcal{S} = \left\lceil 2 \cdot \frac{(1-\theta) \cdot (1+\zeta) + \frac{\theta}{\eta m \sigma}}{\theta + \zeta\theta - \zeta} \right\rceil$, we have that \tilde{D}_0 decreases by a factor of at least $\frac{1}{2}$ every \mathcal{S} rounds of epochs. So in order to achieve an ϵ -additive error, we need totally $O(\log \frac{\tilde{D}_0}{\epsilon})$ restarts of the above procedure.

C.2.1. PARAMETER SETTING FOR TWO CASES

(I) Consider the first case with $\frac{m}{\kappa} \leq \frac{3}{4}$, by choosing identical parameters settings $\eta = \sqrt{\frac{1}{3\sigma m L}}$, $\theta = \sqrt{\frac{m}{3\kappa}} \leq \frac{1}{2}$ and $m = \Theta(n)$ as in Section 3 (so the constraint (27) is satisfied), we have

$$\mathcal{S} = \left\lceil 2 \cdot \frac{2 - (\theta + \zeta \cdot (\theta - 1))}{\theta + \zeta \cdot (\theta - 1)} \right\rceil.$$

Imposing an additional constraint on the sparse variance: $\zeta \leq \sqrt{\frac{m}{4\kappa}}$, we have $\mathcal{S} = O(\sqrt{\frac{\kappa}{n}})$, which means that the total oracle complexity is

$$\mathcal{O} \left(\mathcal{S} \cdot \mathcal{O} \left(\log \frac{\tilde{D}_0}{\epsilon} \right) \cdot (m+n) \right) = \mathcal{O} \left(\sqrt{\kappa n} \log \frac{F(\tilde{x}_0) - F(x^*)}{\epsilon} \right)$$

(II) Consider another case with $\frac{m}{\kappa} > \frac{3}{4}$ and $\zeta \leq C_\zeta$, let $\hat{\theta} \triangleq \frac{C_\zeta + 0.5}{C_\zeta + 1}$, by choosing $\theta = \frac{\zeta + 0.5}{\zeta + 1} \in [0.5, \hat{\theta}]$, $\eta = \frac{1-\hat{\theta}}{2m\sigma\theta} \leq \frac{1-\theta}{L\theta(2-\theta)}$ (the constraint (27) is satisfied) and $m = \Theta(n)$, we have $\mathcal{S} = O(1)$ (correlated with C_ζ), the total oracle complexity $\mathcal{O} \left(n \log \frac{F(\tilde{x}_0) - F(x^*)}{\epsilon} \right)$.

C.2.2. DISCUSSION ABOUT SPARSE VARIANCE BOUND (LEMMA 4 AND LEMMA 5)

In the dense update case (with $D_m = 1$), we see that Lemma 5 degenerates to Lemma 1, this bound is much tighter than the bound in Lemma 4 (Lemma 4 is identical for both the sparse and dense cases).

However, in the sparse update case, Lemma 5 highly correlates with dataset sparsity ($\propto D_m^2$), which could be much looser than Lemma 4 in some extreme cases (imagine a dataset with some dimensions contain only one entry among the n samples, so $D_m = n$). Unfortunately, our MiG algorithm (as well as Katyusha) relies on canceling the additional variance term to yield a tight correlation inside one iteration:

$$F(\tilde{x}_{s-1}) - F(y_{j-1}) - \boxed{\langle \nabla F(y_{j-1}), \tilde{x}_{s-1} - y_{j-1} \rangle} \rightarrow \text{canceled by coupling term } (1-\theta)\tilde{x}_s$$

If D_m is as large as n , the oracle complexity in Theorem 4 could be even worse than that in Theorem 3.

It is still an open problem whether we can have a tighter variance bound in the sparse update setting that is uncorrelated with D_m .

C.3. Proof of Theorem 5

Here we analyze Algorithm 4 based on the “perturbed iterate analysis” framework (Mania et al., 2017).

To begin with, we need to specify the iterates labeling order, which is crucial in our asynchronous analysis.

Choice of labeling order. There are “Before Read” (Mania et al., 2017) and “After Read” (Leblond et al., 2017) labeling schemes proposed in recent years which are reasonable in asynchronous analysis. Among these two schemes, “Before Read” requires considering the updates from “future”, which leads to a complex analysis. “After Read” enjoys a simpler analysis but requires changing the order of sampling step to ensure uniform distributed samples¹⁶. In order to give a clean proof, we adopt the “After Read” labeling scheme and make the following assumptions:

Assumption C.1. *The labeling order increases after the step (14) in Algorithm 4 finished, so the future perturbation is not included in the effect of asynchrony in the current step.*

Assumption C.2. *We explicitly assume uniform distributed samples and the independence of the sample i_j with \hat{x}_{j-1} .*

In other words, we are analyzing the following procedure:

1. Inconsistent read the iterate \hat{x}_{j-1} .
2. Increase iterates counter j and sample a random index i_j .
3. Compute an update $-\eta \cdot \tilde{\nabla}(\hat{y}_{j-1})$.
4. Atomic write the update to shared memory coordinately.

From (Leblond et al., 2017), we can model the effect of asynchrony as follows:

$$\hat{x}_j - x_j = \eta \sum_{k=(j-1-\tau)_+}^{j-2} \mathcal{S}_k^j \tilde{\nabla}(\hat{y}_k), \quad (32)$$

where \mathcal{S}_k^j is a diagonal matrix with entries in $\{0, +1\}$. This definition models the coordinate perturbation from the past updates. Here τ represents the maximum number of overlaps between concurrent threads (Mania et al., 2017). We further denote $\Delta = \max_{k=1 \dots d} p_k$ following (Leblond et al., 2017), which provides a measure of sparsity.

Then we start our analysis with the iterate difference between “fake” y_j and x^* . By expanding the iterate difference and taking expectation with respect to the sample i_j , we get

$$\begin{aligned} \mathbb{E}_{i_j} [\|y_j - x^*\|^2] &= \mathbb{E}_{i_j} [\|\theta(x_{j-1} - \eta \cdot \tilde{\nabla}(\hat{y}_{j-1})) + (1 - \theta)\tilde{x}_{s-1} - x^*\|^2] \\ &= \mathbb{E}_{i_j} [\|y_{j-1} - \eta\theta \cdot \tilde{\nabla}(\hat{y}_{j-1}) - x^*\|^2] \\ &\stackrel{(\star)}{=} \|y_{j-1} - x^*\|^2 - 2\eta\theta \langle \nabla F(\hat{y}_{j-1}), \hat{y}_{j-1} - x^* \rangle + \eta^2 \theta^2 \mathbb{E}_{i_j} [\|\tilde{\nabla}(\hat{y}_{j-1})\|^2] \\ &\quad + 2\eta\theta \mathbb{E}_{i_j} [\langle \tilde{\nabla}(\hat{y}_{j-1}), \hat{y}_{j-1} - y_{j-1} \rangle], \end{aligned} \quad (33)$$

where (\star) uses the unbiasedness $\mathbb{E}_{i_j} [\tilde{\nabla}(\hat{y}_{j-1})] = \nabla F(\hat{y}_{j-1})$ and the independence Assumption C.2.

Using Lemma 4, we get the variance bound

$$\mathbb{E}_{i_j} [\|\tilde{\nabla}(\hat{y}_{j-1})\|^2] \leq 4L(F(\hat{y}_{j-1}) - F(x^*)) + 4L(F(\tilde{x}_{s-1}) - F(x^*)). \quad (34)$$

¹⁶So there are always two versions (analyzed, implemented) of algorithms in the works with “After Read” scheme (Leblond et al., 2017; Pedregosa et al., 2017)

Using the σ -strongly convex of $F(\cdot)$, we can bound $-\langle \nabla F(\hat{y}_{j-1}), \hat{y}_{j-1} - x^* \rangle$ as follows:

$$\begin{aligned} \langle \nabla F(\hat{y}_{j-1}), \hat{y}_{j-1} - x^* \rangle &\geq F(\hat{y}_{j-1}) - F(x^*) + \frac{\sigma}{2} \|\hat{y}_{j-1} - x^*\|^2 \\ &\stackrel{(\star)}{\geq} F(\hat{y}_{j-1}) - F(x^*) + \frac{\sigma}{4} \|y_{j-1} - x^*\|^2 - \frac{\sigma}{2} \|\hat{y}_{j-1} - y_{j-1}\|^2, \end{aligned} \quad (35)$$

where (\star) uses the fact that $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$.

Combining (33), (34) and (35), we get

$$\begin{aligned} \mathbb{E}_{i_j} [\|y_j - x^*\|^2] &\leq (1 - \frac{\eta\theta\sigma}{2}) \|y_{j-1} - x^*\|^2 + \eta\theta\sigma \|\hat{y}_{j-1} - y_{j-1}\|^2 + 2\eta\theta \mathbb{E}_{i_j} [\langle \tilde{\nabla}(\hat{y}_{j-1}), \hat{y}_{j-1} - y_{j-1} \rangle] \\ &\quad + (4L\eta^2\theta^2 - 2\eta\theta)(F(\hat{y}_{j-1}) - F(x^*)) + 4L\eta^2\theta^2(F(\tilde{x}_{s-1}) - F(x^*)). \end{aligned} \quad (36)$$

From Lemma 1 in (Leblond et al., 2017), we borrow the notations $C_1 = 1 + \sqrt{\Delta}\tau$, $C_2 = \sqrt{\Delta} + \eta\theta\sigma C_1$ and bound the asynchronous variance terms $\|\hat{y}_{j-1} - y_{j-1}\|^2$, $\mathbb{E}_{i_j} [\langle \tilde{\nabla}(\hat{y}_{j-1}), \hat{y}_{j-1} - y_{j-1} \rangle]$ using (32) as

$$\mathbb{E}_{i_j} [\langle \tilde{\nabla}(\hat{y}_{j-1}), \hat{y}_{j-1} - y_{j-1} \rangle] \leq \frac{\eta\theta\sqrt{\Delta}}{2} \sum_{k=(j-1-\tau)_+}^{j-2} \|\tilde{\nabla}(\hat{y}_k)\|^2 + \frac{\eta\theta\sqrt{\Delta}\tau}{2} \mathbb{E}_{i_j} [\|\tilde{\nabla}(\hat{y}_{j-1})\|^2], \quad (37)$$

$$\|\hat{y}_{j-1} - y_{j-1}\|^2 \leq \eta^2\theta^2(1 + \sqrt{\Delta}\tau) \sum_{k=(j-1-\tau)_+}^{j-2} \|\tilde{\nabla}(\hat{y}_k)\|^2. \quad (38)$$

Upper bounding the asynchronous terms in (36) using (37) and (38), we get

$$\begin{aligned} \mathbb{E}_{i_j} [\|y_j - x^*\|^2] &\leq (1 - \frac{\eta\theta\sigma}{2}) \|y_{j-1} - x^*\|^2 + \eta^2\theta^2(\sqrt{\Delta} + \eta\theta\sigma(1 + \sqrt{\Delta}\tau)) \sum_{k=(j-1-\tau)_+}^{j-2} \|\tilde{\nabla}(\hat{y}_k)\|^2 \\ &\quad + (4L\eta^2\theta^2(1 + \sqrt{\Delta}\tau) - 2\eta\theta)(F(\hat{y}_{j-1}) - F(x^*)) + 4L\eta^2\theta^2(1 + \sqrt{\Delta}\tau)(F(\tilde{x}_{s-1}) - F(x^*)). \end{aligned}$$

Defining $a_j \triangleq \|y_j - x^*\|^2$, $\hat{D}_{j-1} = F(\hat{y}_{j-1}) - F(x^*)$, $\tilde{D}_{s-1} = F(\tilde{x}_{s-1}) - F(x^*)$ for a clean proof and rearranging, we obtain

$$\begin{aligned} \mathbb{E}_{i_j} [a_j] &\leq (1 - \frac{\eta\theta\sigma}{2}) a_{j-1} + \eta^2\theta^2 C_2 \sum_{k=(j-1-\tau)_+}^{j-2} \|\tilde{\nabla}(\hat{y}_k)\|^2 + (4L\eta^2\theta^2 C_1 - 2\eta\theta) \hat{D}_{j-1} + 4L\eta^2\theta^2 C_1 \tilde{D}_{s-1}, \\ (2\eta\theta - 4L\eta^2\theta^2 C_1) \hat{D}_{j-1} &\stackrel{(\star)}{\leq} (a_{j-1} - \mathbb{E}_{i_j} [a_j]) + \eta^2\theta^2 C_2 \sum_{k=(j-1-\tau)_+}^{j-2} \|\tilde{\nabla}(\hat{y}_k)\|^2 + 4L\eta^2\theta^2 C_1 \tilde{D}_{s-1}, \end{aligned} \quad (39)$$

where (\star) uses the fact that $1 - \frac{\eta\theta\sigma}{2} \leq 1$.

Summing (39) over $j = 1 \dots m$ and taking expectation with all randomness in this epoch, we get

$$(2\eta\theta - 4L\eta^2\theta^2 C_1) \sum_{j=1}^m \mathbb{E} [\hat{D}_{j-1}] \leq (a_0 - \mathbb{E} [a_m]) + \eta^2\theta^2 C_2 \sum_{j=1}^m \sum_{k=(j-1-\tau)_+}^{j-2} \mathbb{E} [\|\tilde{\nabla}(\hat{y}_k)\|^2] + 4L\eta^2\theta^2 C_1 m \tilde{D}_{s-1}. \quad (40)$$

Then we focus on upper bounding the second term on the right side of (40),

$$\sum_{j=1}^m \sum_{k=(j-1-\tau)_+}^{j-2} \mathbb{E} [\|\tilde{\nabla}(\hat{y}_k)\|^2] \leq \tau \sum_{j=1}^{m-1} \mathbb{E} [\|\tilde{\nabla}(\hat{y}_{j-1})\|^2] \leq \tau \sum_{j=1}^m \mathbb{E} [\|\tilde{\nabla}(\hat{y}_{j-1})\|^2] \stackrel{(\star)}{\leq} 4L\tau (\sum_{j=1}^m \mathbb{E} [\hat{D}_{j-1}] + m \tilde{D}_{s-1}),$$

where (\star) uses the variance bound (34).

Substituting the above inequality into (40), we get

$$(2\eta\theta - 4L\eta^2\theta^2C_1 - 4L\eta^2\theta^2C_2\tau) \sum_{j=1}^m \mathbb{E}[\hat{D}_{j-1}] \leq a_0 + (4L\eta^2\theta^2C_1m + 4L\eta^2\theta^2C_2\tau m)\tilde{D}_{s-1},$$

$$\tilde{D}_s \stackrel{(\star)}{\leq} \frac{\frac{2}{\sigma} + 4L\eta^2\theta^2C_1m + 4L\eta^2\theta^2C_2\tau m}{(2\eta\theta - 4L\eta^2\theta^2C_1 - 4L\eta^2\theta^2C_2\tau)m} \cdot \tilde{D}_{s-1},$$

where (\star) uses the σ -strongly convex of $F(\cdot)$ and $\tilde{x}_0 = x_0 = y_0$, $\tilde{x}_s = \frac{1}{m} \sum_{j=0}^{m-1} \hat{y}_j$.

By choosing $m = 60\kappa$, $\eta = \frac{1}{5L}$, $\theta = \frac{1}{6}$, we get

$$\tilde{D}_s \leq \frac{2 + \frac{4}{15}(C_1 + C_2\tau)}{4 - \frac{4}{15}(C_1 + C_2\tau)} \cdot \tilde{D}_{s-1}.$$

In order to ensure linear speed up, τ needs to satisfy the following constraint:

$$\rho \triangleq \frac{2 + \frac{4}{15}(C_1 + C_2\tau)}{4 - \frac{4}{15}(C_1 + C_2\tau)} \leq 1.$$

By simply setting $\tau \leq \min\{\frac{5}{4\sqrt{\Delta}}, 2\kappa, \sqrt{\frac{2\kappa}{\sqrt{\Delta}}}\}$, the above constraint is satisfied with $\rho \leq 0.979$, which implies that the total oracle complexity is $\mathcal{O}((n + \kappa) \log \frac{F(\tilde{x}_0) - F(x^*)}{\epsilon})$.

D. Experimental Setup

All our algorithms were implemented in C++ and parameters were passed through MATLAB interface for fair comparison¹⁷. Detailed settings are divided into the following cases.

D.1. In Serial Dense Case

In this case, we ran experiments on the HP Z440 machine with single Intel Xeon E5-1630v4 with 3.70GHz cores, 16GB RAM, Ubuntu 16.04 LTS with GCC 4.9.0, MATLAB R2017b. We are optimizing the following binary problems with $a_i \in \mathbb{R}^d$, $b_i \in \{-1, +1\}$, $i = 1 \dots m$:

$$\begin{aligned} \text{Logistic Regression: } & \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-b_i a_i^T x)) + \frac{\lambda}{2} \|x\|^2, \\ \text{Ridge Regression: } & \frac{1}{n} \sum_{i=1}^n (a_i^T x + b_i)^2 + \frac{\lambda}{2} \|x\|^2, \\ \text{LASSO: } & \frac{1}{n} \sum_{i=1}^n (a_i^T x + b_i)^2 + \lambda \|x\|_1, \end{aligned} \tag{41}$$

where λ is the regularization parameter.

We used datasets from LibSVM website¹⁸, including **a9a** (32,561 samples, 123 features), **covtype.binary** (581,012 samples, 54 features), **w8a** (49,749 samples, 300 features), **ijcnn1** (49,990 samples, 22 features). We added one dimension as bias to all the datasets and then normalized all data vectors to 1 for the ease of experimental setup.

We mainly compared MiG with the following state-of-the-art algorithms:

- SVRG. For theoretical evaluation, we set the learning rate as $\frac{1}{4L}$, which is a reasonable learning rate for SVRG in theory. Otherwise in the strongly convex case, we tuned the learning rate. For non-smooth regularizers (e.g., LASSO), we ran Prox-SVRG (Xiao & Zhang, 2014) instead.

¹⁷The code of our method can be downloaded from the anonymous link:

https://www.dropbox.com/s/1a5v1gvioqbjtv3/Async_Sparse_MiG.zip?dl=0.

¹⁸<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

- SAGA. For theoretical evaluation, we set the learning rate as $\frac{1}{2(\sigma n+L)}$ following (Defazio et al., 2014). Otherwise (include the non-strongly convex case), we tuned the learning rate.
- Acc-Prox-SVRG. This algorithm is quite unstable if mini-batch size is set to 1, we set same learning rate as in SVRG and tuned the momentum parameter β (Nitanda, 2014).
- Catalyst on SVRG. We set the same learning rate as in SVRG and carefully tuned the parameters α_0 and κ (Lin et al., 2015).
- Katyusha. As suggested by the author, we fixed $\tau_2 = \frac{1}{2}$ (sometimes we tuned τ_2 for a better performance), set $\eta = \frac{1}{3\tau_1 L}$ and tuned only τ_1 (Allen-Zhu, 2017). For theoretical evaluation, we chose $\tau_1 = \sqrt{\frac{m}{3\kappa}}$. For Katyusha^{NS}, we used $\tau_1 = \frac{2}{s+4}$, $\alpha = \frac{1}{a\tau_1 L}$ and tune a .
- MiG. Similarly, we set $\eta = \frac{1}{3\theta L}$ and tuned only θ . For theoretical evaluation, we chose $\theta = \sqrt{\frac{m}{3\kappa}}$. For MiG^{NSC}, we chose $\theta = \frac{2}{s+4}$, $\eta = \frac{1}{aL\theta}$ and tuned a .

More theoretical evaluation results for ℓ_2 -logistic regression and ridge regression problems are shown in Figures 1 and 2, respectively, where the regularization parameter was set to some relatively small values, e.g., 10^{-6} , 10^{-7} , and 10^{-8} .

More practical evaluation results (with parameter tuning for all the algorithms and for relatively large λ) are shown in Figures 3 and 4.

We also compared the performance of MiG^{NSC} with that of all the algorithms mentioned above, as shown in Figures 5 and 6, where the results are for the non-strongly convex logistic regression (i.e., $\lambda=0$) and LASSO, respectively.




Figure 1. Theoretical evaluations of MiG and other state-of-the-art algorithms for solving ℓ_2 -logistic regression on covtype ($\lambda = 10^{-7}$, the first two figures) and a9a ($\lambda = 10^{-7}$, the last two figures).




Figure 2. Theoretical evaluations of MiG and other state-of-the-art algorithms for solving ridge regression on covtype ($\lambda = 10^{-8}$, the first two figures) and w8a ($\lambda = 10^{-6}$, the last two figures).




Figure 3. Practical evaluations of MiG and other state-of-the-art algorithms for solving ℓ_2 -logistic regression on covtype (a), ijcn1 (b), w8a (c), and a9a (d).




Figure 4. Practical evaluations of MiG and other state-of-the-art algorithms for solving ridge regression on covtype (a), ijcn1 (b), covtype (c), and w8a (d).




Figure 5. Comparison of MiG^{NSC} and other state-of-the-art algorithms for solving non-strongly convex logistic regression without regularizer (i.e., $\lambda = 0$) on *a9a* (the first two figures) and *ijcnn1* (the last two figures).




Figure 6. Comparison of MiG^{NSC} and other state-of-the-art algorithms for solving the non-strongly convex problem (LASSO) on *covtype* ($\lambda = 10^{-6}$, the first two figures) and *w8a* ($\lambda = 10^{-7}$, the last two figures).

D.2. In Asynchronous Sparse Case

Experiments in this setting were running on a server with 4 Intel Xeon E7-4820 v3 with 10 cores each 1.90GHz, 512 GB RAM, CentOS 7.4.1708 with GCC 4.8.5, MATLAB R2017a. Multi-threads experiments were running on certain CPU cores using `numactl`. We are optimizing the following smooth and strongly convex problem:

$$\text{Logistic Regression: } \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-b_i a_i^T x)) + \frac{\lambda}{2} \|x\|^2,$$

where λ is the regularization parameter. Notice that the ℓ_2 regularizer is dense with respect to the sample i , so we used the following sparse and unbiased stochastic gradient as in (Leblond et al., 2017):

$$\nabla_i(x) = \nabla f_i(x) + \lambda D_i x.$$

We tuned learning rates for ASAGA and KroMagnon, learning rate and θ for MiG.

D.2.1. PARAMETER TUNING CRITERIA

For the relatively small RCV1 dataset, we carefully tuned the learning rate for ASAGA and KroMagnon to achieve best performance. For MiG, following theoretical intuition in the dense case, we first chose a large learning rate and then carefully tuned θ to achieve best performance (thus the effect of θ can be regard as stabilize iterates and allow for a larger learning rate).

For the KDD2010 dataset, we first made 3 subsets with 19,000 samples, 190,000 samples and 1,900,000 samples (corresponds to $\lambda = 10^{-7}$, 10^{-8} and 10^{-9}). Then we tuned parameters for them following the above criteria and finally used relatively “safe” parameter settings for the entire dataset.

D.2.2. ON THE EFFECTIVENESS OF OUR ACCELERATION TRICK

In this part, we evaluated the effect of our acceleration trick in MiG. One can easily verify that when θ is set to 1, the coupling term is neglected and the algorithm is quite similar to SVRG (different on the choice of the snapshot point). Based

on the theoretical analysis in Section 3 and Section 4.1, we claim that the effect of θ is to stabilize the iterates to adopt a larger learning rate. Thus, we designed an experiment to justify this claim.

We compared the best-tuned performance of MiG, MiG with $\theta = 1$, and KroMagnon on RCV1 on a server with the 16 threads, as shown in Figure 7.




Figure 7. Evaluation of the effectiveness of our acceleration trick of our algorithm for solving ℓ_2 -logistic regression with $\lambda = 10^{-9}$.

Based on the parameter choice in this experiment, MiG with $\theta = 1$ cannot use the same large learning rate as MiG does. When the suitable parameter θ is chosen, MiG outperforms both algorithms as desired, which further indicates the effectiveness of our acceleration trick.