
A Universally Optimal Multistage Accelerated Stochastic Gradient Method

Necdet Serhat Aybat*
Pennsylvania State University
University Park, PA, USA
nsa10@psu.edu

Alireza Fallah*
Massachusetts Institute of Technology
Cambridge, MA, USA
afallah@mit.edu

Mert Gürbüzbalaban*
Rutgers University
Piscataway, NJ, USA
mg1366@rutgers.edu

Asuman Ozdaglar*
Massachusetts Institute of Technology
Cambridge, MA, USA
asuman@mit.edu

Abstract

We study the problem of minimizing a strongly convex, smooth function when we have noisy estimates of its gradient. We propose a novel multistage accelerated algorithm that is universally optimal in the sense that it achieves the optimal rate both in the deterministic and stochastic case and operates without knowledge of noise characteristics. The algorithm consists of stages that use a stochastic version of Nesterov’s method with a specific restart and parameters selected to achieve the fastest reduction in the bias-variance terms in the convergence rate bounds.

1 Introduction

First order optimization methods play a key role in solving large scale machine learning problems due to their low iteration complexity and scalability with large data sets. In several cases, these methods operate with noisy first order information either because the gradient is estimated from draws or subset of components of the underlying objective function [3, 8, 13, 16, 17, 21, 36, 9, 11] or noise is injected intentionally due to privacy or algorithmic considerations [4, 25, 30, 14, 15]. A fundamental question in this setting is to design fast algorithms with *optimal convergence rate*, matching the lower bounds on the oracle complexity in terms of target accuracy and other important parameters both for the deterministic and stochastic case (i.e., with or without gradient errors).

In this paper, we design an *optimal* first order method to solve the problem

$$f^* \triangleq \min_{x \in \mathbb{R}^d} f(x) \quad \text{such that } f \in S_{\mu,L}(\mathbb{R}^d), \quad (1)$$

where, for scalars $0 < \mu \leq L$, $S_{\mu,L}(\mathbb{R}^d)$ is the set of continuously differentiable functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ that are strongly convex with modulus μ and have Lipschitz-continuous gradients with constant L , which imply that for every $x, y \in \mathbb{R}^d$, f satisfies (see e.g. [27])

$$\frac{\mu}{2} \|x - y\|^2 \leq f(x) - f(y) - \nabla f(y)^\top (x - y) \leq \frac{L}{2} \|x - y\|^2. \quad (2)$$

For $f \in S_{\mu,L}(\mathbb{R}^d)$, the ratio $\kappa \triangleq \frac{L}{\mu}$ is called the *condition number* of f . Throughout the paper, we denote the solution of problem (1) by f^* which is achieved at the *unique* optimal point x^* .

*The authors are in alphabetical order.

We assume that the gradient information is available through a stochastic oracle, which at each iteration n , given the current iterate $x_n \in \mathbb{R}^d$, provides the noisy gradient $\tilde{\nabla}f(x_n, w_n)$ where $\{w_n\}_n$ is a sequence of independent random variables such that for all $n \geq 0$,

$$\mathbb{E}[\tilde{\nabla}f(x_n, w_n)|x_n] = \nabla f(x_n), \quad \mathbb{E}[\|\tilde{\nabla}f(x_n, w_n) - \nabla f(x_n)\|^2|x_n] \leq \sigma^2. \quad (3)$$

This oracle model is commonly considered in the literature (see e.g. [16, 17, 6]). In Appendix K, we show how our analysis can be extended to the following more general noise setting, same as the one studied in [3], where the variance of the noise is allowed to grow linearly with the squared distance to the optimal solution:

$$\mathbb{E}[\tilde{\nabla}f(x_n, w_n)|x_n] = \nabla f(x_n), \quad \mathbb{E}[\|\tilde{\nabla}f(x_n, w_n) - \nabla f(x_n)\|^2|x_n] \leq \sigma^2 + \eta^2\|x_n - x^*\|^2. \quad (4)$$

for some constant $\eta \geq 0$.

Under noise setting in (3), the performance of many algorithms is characterized by the expected error of the iterates (in terms of the suboptimality in function values) which admits a bound as a sum of two terms: a *bias term* that shows the decay of initialization error $f(x_0) - f^*$ and is independent of the noise parameter σ^2 , and a *variance term* that depends on σ^2 and is independent of the initial point x_0 . A lower bound on the bias term follows from the seminal work of Nemirovsky and Yudin [26], which showed that without noise ($\sigma = 0$) and after n iterations, $\mathbb{E}[f(x_n)] - f^*$ cannot be smaller than²

$$L\|x_0 - x^*\|_2^2 \exp(-\mathcal{O}(1)\frac{n}{\sqrt{\kappa}}). \quad (5)$$

With noise, Raginsky and Rakhlin [31] provided the following (much larger) lower bound³ on function suboptimality which also provides a lower bound on the variance term:

$$\Omega\left(\frac{\sigma^2}{\mu n}\right) \quad \text{for } n \text{ sufficiently large.} \quad (6)$$

Several algorithms have been proposed in the recent literature attempting to achieve these lower bounds.⁴ Xiao [38] obtains $\mathcal{O}(\log(n)/n)$ performance guarantees in expected suboptimality for an accelerated version of the dual averaging method. Dieuleveut et al. [12] consider quadratic objective function and develop an algorithm with averaging to achieve the error bound $\mathcal{O}(\frac{\sigma^2}{n} + \frac{\|x_0 - x^*\|^2}{n^2})$. Hu et al. [20] consider general strongly convex and smooth functions and achieve an error bound with similar dependence under the assumption of bounded noise. Ghadimi and Lan [16] and Chen et al. [7] extend this result to the noise model in (3) by introducing the accelerated stochastic approximation algorithm (AC-SA) and optimal regularized dual averaging algorithm (ORDA), respectively. Both AC-SA and ORDA have multistage versions presented in [17] and [7] where authors improve the bias term of their single stage methods to the optimal $\exp(-\mathcal{O}(1)n/\sqrt{\kappa})$ by exploiting knowledge of σ and the optimality gap Δ , i.e., an upper bound for $f(x_0) - f^*$, in the operation of the algorithm. Another closely related paper is [8] which proposed μ AGD+ and showed under *additive* noise model that it admits the error bound $\mathcal{O}(\frac{\sigma^2}{n} + \frac{\|x_0 - x^*\|^2}{n^p})$ for any $p \geq 1$ where the constants grow with p , and in particular, they achieve the bound $\mathcal{O}(\frac{\sigma^2 \log n}{n} + \frac{\|x_0 - x^*\|^2 \log n}{n^{\log n}})$ for $p = \log n$.

In this paper, we introduce the class of Multistage Accelerated Stochastic Gradient (M-ASG) methods that are universally optimal, achieving the lower bound both in the noiseless deterministic case and the noisy stochastic case up to some constants independent of μ and L . M-ASG proceeds in stages that use a stochastic version of Nesterov's accelerated method [27] with a specific restart and parameterization. Given an arbitrary length and constant stepsize for the first stage together with geometrically growing lengths and shrinking stepsizes for the following stages, we first provide a general convergence rate result for M-ASG (see Theorem 3.4). Given the computational budget n , a specific choice for the length of the first stage is shown to achieve the optimal error bound *without requiring knowledge of the noise bound σ^2 and the initial optimality gap* (See Corollary 3.8).

²This lower bound is shown with the additional assumption $n \leq d$

³The authors show this result for $\mu = 1$. Nonetheless, it can be generalized to any $\mu > 0$ by scaling the problem parameters properly.

⁴Here we review their error bounds after n iterations highlighting dependence on σ^2 , n , and initial point x_0 , suppressing μ and L dependence.

Table 1: Comparison of algorithms

Algorithm	Requires			Opt. Bias	Opt. Var.
	σ	Δ	n or ϵ		
AC-SA	\times	\times	\times	\times	\checkmark
Multi. AC-SA	\checkmark	\checkmark	\times	\checkmark	\checkmark
ORDA	\times	\times	\times	\times	\checkmark
Multi. ORDA	\checkmark	\checkmark	\times	\checkmark	\checkmark
Cohen et al.	\times	\times	\times	\times	\checkmark
M-ASG (With parameters in Corollary 3.7)	\times	\times	\times	\times	\checkmark
M-ASG (With parameters in Corollary 3.8)	\times	\times	$\checkmark[n]$	\checkmark	\checkmark
M-ASG (With parameters in Corollary 3.9)	\times	\checkmark	$\checkmark[\epsilon]$	\checkmark	\checkmark

To the best of our knowledge, this is the first algorithm that achieves such a lower bound under such informational assumptions. In Table 1, we provide a comparison of our algorithm with other algorithms in terms of required assumptions and optimality of their results in both bias and variance terms. In particular, we consider ACSA [16], Multistage AC-SA [17], ORDA and Multistage ORDA [7], and the algorithm proposed in [8].

Our paper builds on an analysis of Nesterov’s accelerated stochastic method with a specific momentum parameter presented in Section 2 which may be of independent interest. This analysis follows from a dynamical system representation and study of first order methods which has gained attention in the literature recently [24, 19, 2]. In Section 3, we present the M-ASG algorithm, and characterize its behavior under different assumptions as summarized in Table 1. In particular, we show that it achieves the optimal convergence rate with the given budget of iterations n . In Section 4, we show how additional information such as σ and Δ can be leveraged in our framework to improve practical performance. Finally, in Section 5, we provide numerical results on the comparison of our algorithm with some of the other most recent methods in the literature.

Preliminaries and notation: Let I_d and 0_d represent the $d \times d$ identity and zero matrices. For matrix $A \in \mathbb{R}^{d \times d}$, $\text{Tr}(A)$ and $\det(A)$ denote the trace and determinant of A , respectively. Also, for scalars $1 \leq i \leq j \leq d$ and $1 \leq k \leq l \leq d$, we use $A_{[i:j], [k:l]}$ to show the submatrix formed by rows i to j and columns k to l . We use the superscript $^\top$ to denote the transpose of a vector or a matrix depending on the context. Throughout this paper, all vectors are represented as column vectors. Let \mathbb{S}_+^m denote the set of all symmetric and positive semi-definite $m \times m$ matrices. For two matrices $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{p \times q}$, their Kronecker product is denoted by $A \otimes B$. For scalars $0 < \mu \leq L$, $S_{\mu, L}(\mathbb{R}^d)$ is the set of continuously differentiable functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ that are strongly convex with modulus μ and have Lipschitz-continuous gradients with constant L . All logarithms throughout the paper are in natural basis.

2 Modeling Accelerated Gradient method as a dynamical system

In this section we study Nesterov’s Accelerated Stochastic Gradient method (ASG) [27] with the stochastic first-order oracle in (3):

$$y_k = (1 + \beta)x_k - \beta x_{k-1}, \quad x_{k+1} = y_k - \alpha \tilde{\nabla} f(y_k, w_k) \quad (7)$$

where $\alpha \in (0, \frac{1}{L}]$ is the stepsize and $\beta = \frac{1 - \sqrt{\alpha\mu}}{1 + \sqrt{\alpha\mu}}$ is the momentum parameter. This choice of momentum parameter has already been studied in the literature, e.g., [28, 37, 33]. In the next lemma, we provide a new motivation for this choice by showing that for quadratic functions and in the noiseless setting, this momentum parameter achieves the *fastest* asymptotic convergence rate for a given fixed stepsize $\alpha \in (0, \frac{1}{L}]$. The proof of this lemma is provided in Appendix A.

Lemma 2.1. *Let $f \in S_{\mu, L}(\mathbb{R}^d)$ be a strongly convex quadratic function such that $f(x) = \frac{1}{2}x^\top Qx - p^\top x + r$ where Q is a d by d symmetric positive definite matrix with all its eigenvalues in the interval $[\mu, L]$. Consider the deterministic ASG iterations, i.e., $\sigma = 0$, as shown in (7), with constant stepsize $\alpha \in (0, 1/L]$. Then, the fastest asymptotic convergence rate, i.e. the smallest $\rho \in (0, 1)$ that satisfies the inequality*

$$\|x_k - x^*\|^2 \leq (\rho + \epsilon_k)^{2k} \|x_0 - x^*\|^2, \quad \forall x_0 \in \mathbb{R}^d,$$

for some non-negative sequence $\{\epsilon_k\}_k$ that goes to zero is $\rho = 1 - \sqrt{\alpha\mu}^5$ and it is achieved by $\beta = \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$. As a consequence, for this choice of β , there exists $\{\epsilon_k\}$ such that $\lim_{k \rightarrow \infty} \epsilon_k = 0$ and

$$f(x_k) - f^* \leq L(1 - \sqrt{\alpha\mu} + \epsilon_k)^{2k} \|x_0 - x^*\|^2.$$

Our analysis builds on the reformulation of a first-order optimization algorithm as a linear dynamical system. Following [24, 19], we write ASG iterations as

$$\xi_{k+1} = A\xi_k + B\tilde{\nabla}f(y_k, w_k), \quad y_k = C\xi_k, \quad (8)$$

where $\xi_k := [x_k^\top \quad x_{k-1}^\top]^\top \in \mathbb{R}^{2d}$ is the *state* vector and A, B and C are matrices with appropriate dimensions defined as the Kronecker products $A = \tilde{A} \otimes I_d$, $B = \tilde{B} \otimes I_d$ and $C = \tilde{C} \otimes I_d$ with

$$\tilde{A} = \begin{bmatrix} 1+\beta & -\beta \\ 1 & 0 \end{bmatrix}, \quad \tilde{B} = \begin{bmatrix} -\alpha \\ 0 \end{bmatrix}, \quad \tilde{C} = [1+\beta \quad -\beta]. \quad (9)$$

We can also relate the state ξ_k to the iterate x_k in a linear fashion through the identity $x_k = T\xi_k$, $T \triangleq [I_d \quad 0_d]$. We study the evolution of the ASG method through the following *Lyapunov function* which also arises in the study of deterministic accelerated gradient methods:

$$V_P(\xi) = (\xi - \xi^*)^\top P(\xi - \xi^*) + f(T\xi) - f^* \quad (10)$$

where P is a symmetric positive semi-definite matrix. We first state the following lemma which can be derived by adapting the proof of Proposition 4.6 in [2] to our setting with less restrictive noise assumption compared to the additive noise model of [2]. Its proof can be found in Appendix B.

Lemma 2.2. *Let $f \in S_{\mu,L}(\mathbb{R}^d)$. Consider the ASG iterations given by (7). Assume there exist $\rho \in (0, 1)$ and $\tilde{P} \in \mathbb{S}_+^2$, possibly depending on ρ , such that*

$$\rho^2 \tilde{X}_1 + (1 - \rho^2) \tilde{X}_2 \succeq \begin{bmatrix} \tilde{A}^\top \tilde{P} \tilde{A} - \rho^2 \tilde{P} & \tilde{A}^\top \tilde{P} \tilde{B} \\ \tilde{B}^\top \tilde{P} \tilde{A} & \tilde{B}^\top \tilde{P} \tilde{B} \end{bmatrix} \quad (11)$$

where

$$\tilde{X}_1 = \frac{1}{2} \begin{bmatrix} \beta^2 \mu & -\beta^2 \mu & -\beta \\ -\beta^2 \mu & \beta^2 \mu & \beta \\ -\beta & \beta & \alpha(2 - L\alpha) \end{bmatrix}, \quad \tilde{X}_2 = \frac{1}{2} \begin{bmatrix} (1+\beta)^2 \mu & -\beta(1+\beta)\mu & -(1+\beta) \\ -\beta(1+\beta)\mu & \beta^2 \mu & \beta \\ -(1+\beta) & \beta & \alpha(2 - L\alpha) \end{bmatrix}.$$

Let $P = \tilde{P} \otimes I_d$. Then, for every $k \geq 0$,

$$\mathbb{E}[V_P(\xi_{k+1})] \leq \rho^2 \mathbb{E}[V_P(\xi_k)] + \sigma^2 \alpha^2 (\tilde{P}_{1,1} + \frac{L}{2}). \quad (12)$$

We use this lemma and derive the following theorem which characterizes the behavior of ASG method for when $\alpha \in (0, 1/L]$ and $\beta = \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$ (see the proof in Appendix C).

Theorem 2.3. *Let $f \in S_{\mu,L}(\mathbb{R}^d)$. Consider the ASG iterations given in (7) with $\alpha \in (0, \frac{1}{L}]$ and $\beta = \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$. Then,*

$$\mathbb{E}[V_{P_\alpha}(\xi_{k+1})] \leq (1 - \sqrt{\alpha\mu}) \mathbb{E}[V_{P_\alpha}(\xi_k)] + \frac{\sigma^2 \alpha}{2} (1 + \alpha L) \quad (13)$$

for every $k \geq 0$, where $P_\alpha = \tilde{P}_\alpha \otimes I_d$ with $\tilde{P}_\alpha = \begin{bmatrix} \sqrt{\frac{1}{2\alpha}} \\ \sqrt{\frac{\mu}{2}} - \sqrt{\frac{1}{2\alpha}} \end{bmatrix} \begin{bmatrix} \sqrt{\frac{1}{2\alpha}} & \sqrt{\frac{\mu}{2}} - \sqrt{\frac{1}{2\alpha}} \end{bmatrix}$.

This result relies on the special structure of P_α which will also be key for our analysis in Section 3.

3 A class of multistage ASG algorithms

In this section, we introduce a class of *multistage* ASG algorithms, represented in Algorithm 1 which we denote by M-ASG. The main idea is to run ASG with properly chosen parameters (α_k, β_k) at

Algorithm 1: Multistage Accelerated Stochastic Gradient Algorithm (M-ASG)

```

1 Set  $n_0 = -1$ ;
2 for  $k = 1; k \leq K; k = k + 1$  do
3   Set  $x_0^k = x_1^k = x_{n_{k-1}+1}^{k-1}$ ;
4   for  $m = 1; m \leq n_k; m = m + 1$  do
5     Set  $\beta_k = \frac{1-\sqrt{\mu\alpha_k}}{1+\sqrt{\mu\alpha_k}}$ ;
6     Set  $y_m^k = (1 + \beta_k)x_m^k - \beta_k x_{m-1}^k$ ;
7     Set  $x_{m+1}^k = y_m^k - \alpha_k \tilde{\nabla} f(y_m^k, w_m^k)$ 
8   end
9 end

```

each stage $k \in 1, \dots, K$ for $K \geq 2$ stages. In addition, each new stage is dependent on the previous stage as the first two initial iterates of the new stage are set to the last iterate of the previous stage.

To analyze Algorithm 1, we first characterize the evolution of iterates in one specific stage through the Lyapunov function in (10). The details of the proof is provided in Appendix D.

Theorem 3.1. *Let $f \in S_{\mu,L}(\mathbb{R}^d)$. Consider running the ASG method given in (7) for n iterations with $\alpha = \frac{c^2}{L}$ and $\beta = \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$ for some $0 < c \leq 1$. Then, for P_α given in Theorem 2.3,*

$$\mathbb{E}[V_{P_\alpha}(\xi_{n+1})] \leq \exp(-n \frac{c}{\sqrt{\kappa}}) \mathbb{E}[V_{P_\alpha}(\xi_1)] + \frac{\sigma^2 \sqrt{\kappa} c}{L}. \quad (14)$$

Given a computational budget of n iterations, we use this result to choose a stepsize that help us achieve an approximately optimal decay in the *variance term* which yields the following corollary for M-ASG algorithm with $K = 1$ stage, and its proof can be found in Appendix E.

Corollary 3.2. *Let $f \in S_{\mu,L}(\mathbb{R}^d)$. Consider running M-ASG, i.e., Algorithm 1, for only one stage with $n_1 = n$ iterations and stepsize $\alpha_1 = \left(\frac{p\sqrt{\kappa} \log n}{n}\right)^2 \frac{1}{L}$ for some scalar $p \geq 1$. Then,*

$$\mathbb{E}[f(x_{n+1}^1)] - f^* \leq \frac{2}{n^p} (f(x_0^1) - f^*) + \frac{p\sigma^2 \log n}{n\mu} \quad (15)$$

provided that $n \geq p\sqrt{\kappa} \max\{2 \log(p\sqrt{\kappa}), e\}$.

For subsequent analysis, given $K \geq 1$, for all $1 \leq k \leq K$, we define the state vector $\xi_i^k = \begin{bmatrix} x_i^k, x_{i-1}^k \end{bmatrix}^\top$ for $1 \leq i \leq n_k + 1$ —recall that $x_0^k = x_1^k = x_{n_{k-1}+1}^{k-1}$, where K is the number of stages. We analyze the performance of each stage with respect to a stage-dependent Lyapunov function $V_{P_{\alpha_k}}$. The following lemma relates the performance bounds with respect to consecutive choice of Lyapunov functions, building on our specific restarting mechanism (The proof can be found in Appendix F).

Lemma 3.3. *Let $f \in S_{\mu,L}(\mathbb{R}^d)$. Consider M-ASG, i.e., Algorithm 1. Then, for every $1 \leq k \leq K-1$,*

$$\mathbb{E}[V_{P_{\alpha_{k+1}}}(\xi_1^{k+1})] \leq 2\mathbb{E}[V_{P_{\alpha_k}}(\xi_{n_k+1}^k)]. \quad (16)$$

Now, we are ready to state and prove the main result of the paper (see proof in Appendix G):

Theorem 3.4. *Let $f \in S_{\mu,L}(\mathbb{R}^d)$. Consider running M-ASG, i.e., Algorithm 1, with some $n_1 \geq 1$ and $\alpha_1 = \frac{1}{L}$ and fixing $n_k = 2^k \lceil \sqrt{\kappa} \log(2^{p+2}) \rceil$ and $\alpha_k = \frac{1}{2^{2k}L}$ for any $k \geq 2$ and $p \geq 1$. The last iterate of each stage, i.e., $x_{n_k+1}^k$, satisfies the following bound for all $k \geq 1$:*

$$\mathbb{E}[f(x_{n_k+1}^k)] - f^* \leq \frac{2}{2^{(p+1)(k-1)}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^1) - f^*)) + \frac{\sigma^2 \sqrt{\kappa}}{L 2^{k-1}}. \quad (17)$$

⁵Note that although this rate is asymptotic, its smaller than the non-asymptotic rate that we provide for general strongly convex functions in Theorem 2.3, as there $\rho = \sqrt{1 - \sqrt{\alpha\mu}}$.

We next define $N_K(p, n_1)$ as the number of iterations needed to run M-ASG for $K \geq 1$ stages, i.e., $N_K(p, n_1) \triangleq \sum_{k=1}^K n_k$. Note for $K \geq 2$ and with parameters given in Theorem 3.4,

$$N_K(p, n_1) = n_1 + (2^{K+1} - 4) \lceil \sqrt{\kappa} \log(2^{p+2}) \rceil. \quad (18)$$

We define $\{x_n\}_{n \in \mathbb{Z}_+}$ sequence such that x_n is the iterate generated by M-ASG algorithm at the end of n gradient steps for $n \geq 0$, i.e., $x_0 = x_0^0$, $x_n = x_{n+1}^1$ for $1 \leq n \leq n_1$, and for $n > n_1$ we set $x_n = x_m^k$ where $k = \lceil \log_2 \left(\frac{n - n_1}{\lceil \sqrt{\kappa} \log(2^{p+2}) \rceil} + 4 \right) - 1 \rceil$ and $m = n - N_{k-1}(p, n_1)$.

Remark 3.5. *In the absence of noise, i.e., $\sigma = 0$, the result of Theorem 3.4 recovers the linear convergence rate of deterministic gradient methods as its special case. Indeed, running M-ASG only for one stage with n iterations, i.e., $K = 1$ and $n_1 = n$ guarantees that $\mathbb{E}[f(x_n)] - f^* \leq 2 \exp(-n/\sqrt{\kappa})(f(x_0^0) - f^*)$ for all $n \geq 1$.*

The next theorem remarks the behavior of M-ASG after running it for n iterations with the parameters in the preceding theorem, and its proof is provided in Appendix H.

Theorem 3.6. *Let $f \in S_{\mu, L}(\mathbb{R}^d)$. Consider running Algorithm 1 for n iterations and with parameters given in Theorem 3.4 and $n_1 < n$. Then the error is bounded by*

$$\begin{aligned} & \mathbb{E}[f(x_n)] - f^* \\ & \leq \mathcal{O}(1) \left(\frac{(8(p+1)\sqrt{\kappa} \log(2))^{p+1}}{(n - n_1)^{p+1}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^0) - f^*)) + \frac{(p+1)\sigma^2}{(n - n_1)\mu} \right). \end{aligned} \quad (19)$$

Corollary 3.7. *Under the premise of Theorem 3.6, choosing $n_1 = \lceil (p+1)\sqrt{\kappa} \log(12(p+1)\kappa) \rceil$, the suboptimality error of M-ASG after $n \geq 2n_1$ admits*

$$\mathbb{E}[f(x_n)] - f^* \leq \mathcal{O}(1) \left(\frac{1}{n^{p+1}} (f(x_0^0) - f^*) + \frac{(p+1)\sigma^2}{n\mu} \right).$$

Theorem 3.6 immediately yields the result in Corollary 3.7, (suboptimal with respect to dependence on initial optimality gap); see Appendix I for the proof. Similar rate results have also been obtained by AC-SA [16] and ORDA [7] algorithms.

We continue this section by pointing out some important special cases of our result. We first show in the next corollary how our algorithm is universally optimal and capable of achieving the lower bounds (5) and (6) simultaneously. The proof follows from (19) and $n - n_1 \geq \frac{n}{2} \geq \sqrt{\kappa}$.

Corollary 3.8. *Under the premise of Theorem 3.6, consider a computational budget of $n \geq 2\sqrt{\kappa}$ iterations. By setting $n_1 = \frac{n}{C}$ for some positive constant $C \geq 2$, we obtain a bound matching the lower bounds in (5) and (6), i.e.,*

$$\mathbb{E}[f(x_n)] - f^* \leq \mathcal{O}(1) \left(\exp(-\frac{n}{C\sqrt{\kappa}})(f(x_0^0) - f^*) + \frac{\sigma^2}{n\mu} \right).$$

We note that achieving the lower bound through the M-ASG algorithm requires the knowledge or estimation of the strong convexity constant μ . In some applications, μ may not be known a priori. However, for regularized risk minimization problems, the regularization parameter is known and it determines the strong convexity constant. It is also worth noting that, even for the deterministic case, [1] has shown that for a wide class of algorithms including ASG, it is *not possible* to obtain the lower bound (5) without knowing the strong convexity parameter. In addition, in Appendix L, we show how our framework can be extended to obtain nearly optimal results in the merely convex setting; i.e. when $\mu = 0$. Finally, note that the Lipschitz constant L can be estimated from data using standard line search techniques in practice, see [5] and [32, Alg. 2].

The lower bound can also be stated as the minimum number of iterations needed to find an ϵ -solution, i.e. to find x_ϵ such that $\mathbb{E}[f(x_\epsilon)] - f^* \leq \epsilon$, for any given $\epsilon > 0$. In the following corollary, and with the additional assumption of knowing the bound Δ on the initial optimality gap $f(x_0^0) - f^*$, we state this version of lower bound. The proof is provided in Appendix J.

Corollary 3.9. *Let $f \in S_{\mu, L}(\mathbb{R}^d)$. Given $\Delta \geq f(x_0^0) - f^*$, for any $\epsilon \in (0, \Delta)$, running M-ASG, Algorithm 1, with parameters given in Theorem 3.4, $p = 1$, and $n_1 = \lceil \sqrt{\kappa} \log(\frac{4\Delta}{\epsilon}) \rceil$, one can compute ϵ -solution within n_ϵ iterations in total, where*

$$n_\epsilon = \lceil \sqrt{\kappa} \log \left(\frac{4\Delta}{\epsilon} \right) \rceil + \lceil 16(1 + \log(8)) \frac{\sigma^2}{\mu\epsilon} \rceil. \quad (20)$$

Recall that we presented a comparison with other state-of-the-art algorithms in Table 1. In particular, this table shows that Multistage AC-SA [17] and Multistage ORDA [7] also achieve the lower bounds provided that noise parameters are known – *note we do not make this extra assumption for M-ASG*. It is also worth noting that the idea of restart, which plays a key role in achieving the lower bounds, has been studied before in the context of deterministic accelerated methods [29, 39]. However, a naive extension of these restart methods to the stochastic setting leads to a two-stage algorithm which switches from constant step-size to diminishing step-size when the variance term dominates the bias term. Nevertheless, implementing this technique requires the knowledge of σ^2 and optimality gap to tune algorithms for achieving optimal rates in both bias and variance terms. M-ASG, on the other hand, achieves the optimal rates using a specific multistage scheme that does not require the knowledge of the parameter σ^2 . In the supplementary material, we also discuss how M-ASG is related to AC-SA and Multistage AC-SA algorithms proposed in [16, 17].

4 M-ASG*: An improved bias-variance trade-off

In section 3, we described a universal algorithm that do not require the knowledge of neither initial suboptimality gap Δ nor the noise magnitude σ^2 to operate. However, as we will argue in this section, our framework is flexible in the sense that additional information about the magnitude of Δ or σ^2 can be leveraged to improve practical performance. We first note that several algorithms in the literature assume that an upper bound on Δ is known or can be estimated, as summarized in Table 1. This assumption is reasonable in a variety of applications when there is a natural lower bound on f . For example, in supervised learning scenarios such as support vector machines, regression or logistic regression problems, the loss function f has non-negative values [35]. Similarly, the noise level σ^2 may be known or estimated, e.g., in private risk minimization [4], the noise is added by the user to ensure privacy; therefore, it is a known quantity.

There is a natural well-known trade-off between constant and decaying stepsizes (decaying with the number of iterations n) in stochastic gradient algorithms. Since the noise is multiplied with the stepsize, a stepsize that is decaying with the number of iterations n leads to a decay in the variance term; however, this will slow down the decay of the bias term, which is controlled essentially by the behavior of the underlying *deterministic accelerated gradient* algorithm (AG) that will give the best performance with the constant stepsize (note that when $\sigma = 0$, the bias term gives the known performance bounds for the AG algorithm). The main idea behind the M-ASG algorithm (which allows it to achieve the lower bounds) is to exploit this trade-off to decide on the *right time*, n_1 , to *switch* to decaying stepsizes, i.e., when the bias term is sufficiently small so that the variance term dominates and should be handled with the decaying stepsize. This insight is visible from the results of Theorem 3.4 which gives further insights on the choice of the stepsize at every stage to achieve the lower bounds. Theorem 3.4 shows that if M-ASG is run with a constant stepsize $\alpha_1 = \frac{1}{L}$ in the first stage, then the variance term admits the bound $\frac{\sigma^2 \sqrt{\kappa}}{L}$ which does not decay with the number of iterations n_1 in the first stage. However, in later stages, when $n > n_1$, the stepsize α_k is decreased as the number of iterations grows and this results in a decay of the variance term. Overall, the choice of the length of the first stage n_1 , has a major impact in practice which we will highlight in our numerical experiments.

If an estimate of Δ or σ^2 is known, it is desirable to choose n_1 as small as possible such that it ensures the bias term becomes smaller than the variance term at the end of the first stage. More specifically, applying Theorem 3.1 for $c = 1$, one can choose n_1 to balance the variance $\frac{\sigma^2 \sqrt{\kappa}}{L}$ and the bias $\exp(-n_1 \frac{1}{\sqrt{\kappa}}) \mathbb{E} [V_{P_{\alpha_1}}(\xi_1^1)]$ terms. The term $\mathbb{E} [V_{P_{\alpha_1}}(\xi_1^1)]$, as shown in the proof of Lemma 3.3, can be bounded by $\mathbb{E} [V_{P_{\alpha_1}}(\xi_1^1)] = \mu \|x_0^0 - x^*\|_2^2 + f(x_0^0) - f^* \leq 2(f(x_0^0) - f^*)$. Therefore, by having an estimate of an upper bound for Δ , n_1 can be set to be the smallest number such that $2\Delta \exp(-n_1 \frac{1}{\sqrt{\kappa}}) \leq \frac{\sigma^2 \sqrt{\kappa}}{L}$, i.e.,

$$n_1 = \lceil \sqrt{\kappa} \log \left(\frac{2L\Delta}{\sigma^2 \sqrt{\kappa}} \right) \rceil. \quad (21)$$

This result allows one to fine-tune the switching point to start using the decaying stepsizes within our framework as a function of σ^2 and Δ . In scenarios, when the noise level σ is small or the initial gap Δ is large, n_1 is chosen large enough to guarantee a fast decay in the bias term. We would like to emphasize that this modified M-ASG algorithm only requires the knowledge of σ and Δ for selecting




Figure 1: Comparison on a quadratic function for $n = 1000$ iterations with different level of noise.




Figure 2: Comparison on a quadratic function for $n = 10000$ iterations with different level of noise.

n_1 and the rest of the parameters can be chosen as in Theorem 3.4 which are independent of both σ and Δ . Finally, the following theorem provides theoretical guarantees of our framework for this choice of n_1 . The proof is omitted as it is similar to the proofs of Theorems 3.4 and 3.6.

Theorem 4.1. *Let $f \in S_{\mu,L}(\mathbb{R}^d)$. Consider running Algorithm 1 for n iterations and with parameters given in Theorem 3.4, $p = 1$, and n_1 set as (21). Then, the expected suboptimality in function values admits the bound $\mathbb{E}[f(x_n)] - f^* \leq 36(1 + \log(8)) \frac{\sigma^2}{(n-n_1)\mu}$ for all $n \geq n_1$.*

5 Numerical experiments

In this section, we demonstrate the numerical performance of Algorithm 1 with parameters specified by Corollary 3.7 (M-ASG) and Theorem 4.1 (M-ASG*) and compare with other methods from the literature. In our first experiment, we consider the strongly convex quadratic objective $f(x) = \frac{1}{2}x^\top Qx - bx + \lambda\|x\|^2$ where Q is the Laplacian of a cycle graph⁶, b is a random vector and $\lambda = 0.01$ is a regularization parameter. We assume the gradients $\nabla f(x)$ are corrupted by additive noise with a Gaussian distribution $\mathcal{N}(0, \sigma_n^2)$ where $\sigma_n^2 \in \{10^{-6}, 10^{-4}, 10^{-2}\}$. We note that this example has been previously considered in the literature as a problem instance where *Standard ASG* (ASG iterations with standard choice of parameters $\alpha = \frac{1}{L}$ and $\beta = \frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}$) perform badly compared to *Standard GD* (Gradient Descent with standard choice of the stepsize $\alpha = 1/L$) [18]. In Figures 1 and 2, we compare M-ASG and M-ASG* with Standard GD, Standard AG, μ AGD+ [8], and Multistage AC-SA [17]. We consider dimension $d = 100$ and initialize all the methods from $x_0^0 = 0$. We run the algorithms Multistage AC-SA, and M-ASG*, having access to the same estimate of Δ . Figures 1-2 show the average performance of all the algorithms along with the 95% confidence interval over 50 sample runs while the total number of iterations $n = 1000$ and $n = 10000$ respectively as the noise level σ^2 is varied. The simulation results reveal that both M-ASG and M-ASG* have typically a faster decay of the error in the beginning and outperforms the other algorithms in general when

⁶All diagonal entries of Q are 2, $Q_{i,j} = -1$ if $|i - j| \equiv 1 \pmod{d}$, and the remaining entries are zero.




Figure 3: Comparison on logistic regression with $n = 10000$ iterations and with different batch sizes.

the number of iterations is small to moderate. In this case, the speed-up obtained by M-ASG and M-ASG* is more prominent if the noise level σ^2 is smaller. However, as the number of iterations grows, the performance of the algorithms become similar as the variance term dominates. In addition, we would like to highlight that when the noise is small, using n_1 as suggested in (21), M-ASG* runs stage one longer than M-ASG; hence, enjoys the linear rate of decay for more iterations before the variance term becomes the dominant term.

For the second set of experiments, we consider a regularized logistic regression problem for binary classification. In particular, we read 10000 images from the M-NIST [23] data-set, and our goal is to distinguish the image of digit zero from that of digit eight.⁷ The number of samples is $N = 1945$, and the size of each image is 20 by 20 after removing the margins (hence $d = 400$ after vectorizing the images). At each iteration, we randomly choose a batch size b of images to compute an estimate of the gradient.⁸ We choose the regularization parameter equal to $\frac{1}{\sqrt{N}}$ following the standard practice (see e.g. [34]). In Figure 3, we compare M-ASG with Standard GD, Standard AG, μ AGD+ [8], and AC-SA [17] for $b \in \{50, 100, 500\}$. The batch size controls the noise level, with larger batches leading to smaller σ . We run each of these algorithms for 50 times, and plot their average performance and 95% confidence intervals. It can be seen that M-ASG usually start faster, and achieves the asymptotic rate of other algorithms for all different batch sizes.

6 Conclusion

In this work, we consider strongly convex smooth optimization problems where we have access to noisy estimates of the gradients. We proposed a multistage method that adapts the choice of the parameters of the Nesterov’s accelerated gradient at each stage to achieve the optimal rate. Our method is universal in the sense that it does not require the knowledge of the noise characteristics to operate and can achieve the optimal rate both in the deterministic and stochastic settings. We provided numerical experiments that compare our method with existing approaches in the literature, illustrating that our method performs well in practice.

Acknowledgements

The work of Necdet Serhat Aybat is partially supported by NSF Grant CMMI-1635106. Alireza Fallah is partially supported by Siebel Scholarship. Mert Gürbüzbalaban acknowledges support from the grants NSF DMS-1723085 and NSF CCF-1814888.

⁷We provide an experiment with synthetic data for logistic loss in Appendix N.

⁸This is an unbiased estimate of the gradient with finite but unknown variance, and therefore we do not use M-ASG* or other algorithms that need the knowledge of variance.

References

- [1] Yossi Arjevani and Ohad Shamir. On the iteration complexity of oblivious first-order optimization algorithms. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 908–916, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [2] Necdet Serhat Aybat, Alireza Fallah, Mert Gurbuzbalaban, and Asuman Ozdaglar. Robust accelerated gradient methods for smooth strongly convex functions. *arXiv preprint arXiv:1805.10579*, 2018.
- [3] Francis Bach and Eric Moulines. Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning. In *Neural Information Processing Systems (NIPS)*, Spain, 2011.
- [4] R. Bassily, A. Smith, and A. Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 464–473. IEEE, 2014.
- [5] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- [6] Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- [7] Xi Chen, Qihang Lin, and Javier Pena. Optimal regularized dual averaging methods for stochastic optimization. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 395–403. Curran Associates, Inc., 2012.
- [8] Michael Cohen, Jelena Diakonikolas, and Lorenzo Orecchia. On acceleration with noise-corrupted gradients. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1019–1028, Stockholmsmässan, Stockholm Sweden, 2018. PMLR.
- [9] A. d’Aspremont. Smooth optimization with approximate gradient. *SIAM Journal on Optimization*, 19(3):1171–1183, 2008.
- [10] E de Klerk. *Aspects of Semidefinite Programming: Interior Point Algorithms and Selected Applications*, volume 65. Springer Science & Business Media, 2002.
- [11] O. Devolder, F. Glineur, and Y. Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146(1-2):37–75, 2014.
- [12] Aymeric Dieuleveut, Nicolas Flammarion, and Francis Bach. Harder, better, faster, stronger convergence rates for least-squares regression. *The Journal of Machine Learning Research*, 18(1):3520–3570, 2017.
- [13] N. Flammarion and F. Bach. From averaging to acceleration, there is only a step-size. In *Conference on Learning Theory*, pages 658–695, 2015.
- [14] X. Gao, M. Gürbüzbalaban, and L. Zhu. Global Convergence of Stochastic Gradient Hamiltonian Monte Carlo for Non-Convex Stochastic Optimization: Non-Asymptotic Performance Bounds and Momentum-Based Acceleration. *ArXiv e-prints*, September 2018.
- [15] Xuefeng Gao, Mert Gurbuzbalaban, and Lingjiong Zhu. Breaking Reversibility Accelerates Langevin Dynamics for Global Non-Convex Optimization. *arXiv e-prints*, page arXiv:1812.07725, December 2018.
- [16] Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization i: A generic algorithmic framework. *SIAM Journal on Optimization*, 22(4):1469–1492, 2012.

- [17] Saeed Ghadimi and Guanghai Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization, ii: shrinking procedures and optimal algorithms. *SIAM Journal on Optimization*, 23(4):2061–2089, 2013.
- [18] M. Hardt. Robustness versus acceleration. August 18th, 2014. <http://blog.mrtz.org/2014/08/18/robustness-versus-acceleration.html>, August 2014.
- [19] Bin Hu and Laurent Lessard. Dissipativity theory for Nesterov’s accelerated method. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1549–1557, International Convention Centre, Sydney, Australia, 2017. PMLR.
- [20] Chonghai Hu, Weike Pan, and James T. Kwok. Accelerated gradient methods for stochastic optimization and online learning. In *Advances in Neural Information Processing Systems 22*, pages 781–789. Curran Associates, Inc., 2009.
- [21] Prateek Jain, Sham M. Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford. Accelerating stochastic gradient descent for least squares regression. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 545–604. PMLR, 2018.
- [22] Guanghai Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1):365–397, Jun 2012.
- [23] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [24] Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.
- [25] Arvind Neelakantan, Luke Vilnis, Quoc V Le, Ilya Sutskever, Lukasz Kaiser, Karol Kurach, and James Martens. Adding gradient noise improves learning for very deep networks. *arXiv preprint arXiv:1511.06807*, 2015.
- [26] Arkadii Semenovich Nemirovsky and David Borisovich Yudin. *Problem complexity and method efficiency in optimization*. Wiley, 1983.
- [27] Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer, 2004.
- [28] Atsushi Nitanda. Stochastic proximal gradient descent with acceleration techniques. In *Advances in Neural Information Processing Systems*, pages 1574–1582, 2014.
- [29] B. O’Donoghue and E. Candès. Adaptive restart for accelerated gradient schemes. *Foundations of Computational Mathematics*, 15(3):715–732, Jun 2015.
- [30] M. Raginsky, A. Rakhlin, and M. Telgarsky. Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. *arXiv preprint arXiv:1702.03849*, 2017.
- [31] Maxim Raginsky and Alexander Rakhlin. Information-based complexity, feedback and dynamics in convex programming. *IEEE Transactions on Information Theory*, 57(10):7036–7056, 2011.
- [32] Mark Schmidt, Reza Babanezhad, Mohamed Ahmed, Aaron Defazio, Ann Clifton, and Anoop Sarkar. Non-uniform stochastic average gradient method for training conditional random fields. In *artificial intelligence and statistics*, pages 819–828, 2015.
- [33] Bin Shi, Simon S Du, Michael I Jordan, and Weijie J Su. Understanding the acceleration phenomenon via high-resolution differential equations. *arXiv preprint arXiv:1810.08907*, 2018.
- [34] Karthik Sridharan, Shai Shalev-Shwartz, and Nathan Srebro. Fast rates for regularized objectives. In *Advances in Neural Information Processing Systems*, pages 1545–1552, 2009.

- [35] Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- [36] Sharan Vaswani, Francis Bach, and Mark Schmidt. Fast and faster convergence of sgd for over-parameterized models and an accelerated perceptron. *arXiv preprint arXiv:1810.07288*, 2018.
- [37] Hoi-To Wai, Wei Shi, Cesar A Uribe, Angelia Nedich, and Anna Scaglione. On curvature-aided incremental aggregated gradient methods. *arXiv preprint arXiv:1806.00125*, 2018.
- [38] Lin Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11(Oct):2543–2596, 2010.
- [39] Zeyuan Allen Zhu and Lorenzo Orecchia. Linear coupling: An ultimate unification of gradient and mirror descent. In *ITCS*, 2017.

A Proof of Lemma 2.1

Let us denote the asymptotic convergence rate of the ASG method as a function of α and β by $\rho(\alpha, \beta)$. It is well-known that $\rho(\alpha, \beta)$ has the following characterization (see e.g. [24], [29]):

$$\rho(\alpha, \beta) = \max\{\rho_\mu(\alpha, \beta), \rho_L(\alpha, \beta)\}, \quad (22)$$

where $\lambda \in \{\mu, L\}$ and ρ_λ is defined as:

$$\rho_\lambda(\alpha, \beta) = \begin{cases} \frac{1}{2}|(1+\beta)(1-\alpha\lambda)| + \frac{1}{2}\sqrt{\Delta_\lambda} & \text{if } \Delta_\lambda \geq 0, \\ \sqrt{\beta(1-\alpha\lambda)} & \text{otherwise,} \end{cases} \quad (23)$$

with $\Delta_\lambda = (1+\beta)^2(1-\alpha\lambda)^2 - 4\beta(1-\alpha\lambda)$. Note that, since $\alpha \leq \frac{1}{L}$, we have $1-\alpha\lambda \geq 0$ for $\lambda \in \{\mu, L\}$; therefore, $\Delta_\lambda \geq 0$ if and only if $(1+\beta)^2(1-\alpha\lambda) \geq 4\beta$, which is equivalent to $\frac{1-\sqrt{\alpha\lambda}}{1+\sqrt{\alpha\lambda}} \geq \beta$.

Using the fact that $\mu \leq L$ and $\frac{1-\sqrt{\alpha\lambda}}{1+\sqrt{\alpha\lambda}}$ is decreasing in $\lambda > 0$, we obtain $\frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}} \geq \frac{1-\sqrt{\alpha L}}{1+\sqrt{\alpha L}}$; hence, for $\beta > \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$, we have both $\Delta_\mu < 0$ and $\Delta_L < 0$. As a consequence, (22) implies that for $\beta > \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$, we have

$$\rho(\alpha, \beta) = \max\{\sqrt{\beta(1-\alpha\mu)}, \sqrt{\beta(1-\alpha L)}\} = \sqrt{\beta(1-\alpha\mu)}. \quad (24)$$

Moreover, for $\beta = \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$, the two branches in (23) take the same value for $\lambda = \mu$ and $\alpha \in (0, 1/L]$; therefore, when β is set to this critical value, we also get $\rho(\alpha, \beta) = \sqrt{\beta(1-\alpha\mu)}$ for $\alpha \in (0, 1/L]$. Note (24) is an increasing function of β for any $\alpha \in (0, 1/L]$; thus, given $\alpha \in (0, 1/L]$, the smallest rate possible is equal to $\inf\{\rho(\alpha, \beta) : \beta \geq \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}\} = 1 - \sqrt{\alpha\mu}$, which is the rate given in the statement of the lemma and it is achieved by $\beta = \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$.

Now, we consider the case $\beta \leq \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$. From (22), if $\rho_\mu(\alpha, \beta) \geq 1 - \sqrt{\alpha\mu}$, then we also have $\rho(\alpha, \beta) \geq 1 - \sqrt{\alpha\mu}$. Thus, showing $\rho_\mu(\alpha, \beta) \geq 1 - \sqrt{\alpha\mu}$ suffices us to claim that for any $\alpha \in (0, 1/L]$, the best possible rate is $1 - \sqrt{\alpha\mu}$ and this can be achieved by setting $\beta = \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$. Indeed, as we discussed above, for the case $\beta \leq \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$, we have $\Delta_\mu \geq 0$; thus,

$$\begin{aligned} \rho_\mu(\alpha, \beta) &= \frac{1}{2}(1+\beta)(1-\alpha\mu) + \frac{1}{2}\sqrt{\Delta_\mu} \\ &= \frac{1-\sqrt{\alpha\mu}}{2} \left((1+\beta)(1+\sqrt{\alpha\mu}) + \sqrt{(1+\beta)^2(1+\sqrt{\alpha\mu})^2 - \frac{4\beta(1+\sqrt{\alpha\mu})}{1-\sqrt{\alpha\mu}}} \right). \end{aligned}$$

Therefore, to show $\rho_\mu(\alpha, \beta) \geq 1 - \sqrt{\alpha\mu}$, we just need to prove

$$\sqrt{(1+\beta)^2(1+\sqrt{\alpha\mu})^2 - \frac{4\beta(1+\sqrt{\alpha\mu})}{1-\sqrt{\alpha\mu}}} \geq 2 - (1+\beta)(1+\sqrt{\alpha\mu}). \quad (25)$$

Taking the square of both sides of (25), it follows that (25) is equivalent to

$$(1+\beta)(1+\sqrt{\alpha\mu}) \geq 1 + \frac{\beta(1+\sqrt{\alpha\mu})}{1-\sqrt{\alpha\mu}}$$

and this holds when $\beta \leq \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$. Therefore, for any $\alpha \in (0, 1/L]$, we have $\rho_\mu(\alpha, \beta) \geq 1 - \sqrt{\alpha\mu}$ for $\beta \leq \frac{1-\sqrt{\alpha\mu}}{1+\sqrt{\alpha\mu}}$, which completes the proof.

B Proof of Lemma 2.2

We first state the following lemma which is an extension of Lemma 4.1 in [2] for ASG.

Lemma B.1. Let $P = \tilde{P} \otimes I_d$ where $\tilde{P} \in \mathbb{S}_+^2$ and consider the function $\mathcal{W}_P(\xi) = (\xi - \xi^*)^\top P(\xi - \xi^*)$. Then we have

$$\mathbb{E}[\mathcal{W}_P(\xi_{k+1})] \leq \mathbb{E} \left[\begin{bmatrix} \xi_k - \xi^* \\ \nabla f(y_k) \end{bmatrix}^\top \begin{bmatrix} A^\top P A & A^\top P B \\ B^\top P A & B^\top P B \end{bmatrix} \begin{bmatrix} \xi_k - \xi^* \\ \nabla f(y_k) \end{bmatrix} \right] + \sigma^2 \alpha^2 \tilde{P}_{11}. \quad (26)$$

Proof. Let $\tilde{\xi}_k = \xi_k - \xi^*$ for any $k \geq 0$. Since $\xi^* = A\xi^*$, (8) implies $\tilde{\xi}_{k+1} = A\tilde{\xi}_k + B(\tilde{\nabla} f(y_k, w_k))$ for $k \geq 0$. Note that, for any $k \geq 1$, $\tilde{\xi}_k$ and y_k are deterministic functions of $\xi_0, \{w_i\}_{i=0}^{k-1}$. Using this fact, along with knowing that w_k is independent of $\xi_0, \{w_i\}_{i=0}^{k-1}$, implies that

$$\begin{aligned} \mathbb{E}[\mathcal{W}_P(\xi_{k+1})] &= \mathbb{E}[\tilde{\xi}_{k+1}^\top P \tilde{\xi}_{k+1}] \\ &= \mathbb{E} \left[(A\tilde{\xi}_k + B\tilde{\nabla} f(y_k, w_k))^\top P (A\tilde{\xi}_k + B\tilde{\nabla} f(y_k, w_k)) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[(A\tilde{\xi}_k + B\tilde{\nabla} f(y_k, w_k))^\top P (A\tilde{\xi}_k + B\tilde{\nabla} f(y_k, w_k)) \middle| \xi_0, \{w_i\}_{i=0}^{k-1} \right] \right] \\ &= \mathbb{E} \left[(A\tilde{\xi}_k + B\nabla f(y_k))^\top P (A\tilde{\xi}_k + B\nabla f(y_k)) \right] \\ &\quad + \mathbb{E} \left[\mathbb{E} \left[\tilde{\nabla} f(y_k, w_k)^\top B^\top P B \tilde{\nabla} f(y_k, w_k) \middle| y_k \right] - \nabla f(y_k)^\top B^\top P B \nabla f(y_k) \right] \end{aligned} \quad (27)$$

$$\begin{aligned} &= \mathbb{E} \left[(A\tilde{\xi}_k + B\nabla f(y_k))^\top P (A\tilde{\xi}_k + B\nabla f(y_k)) \right] \\ &\quad + \alpha^2 \tilde{P}_{11} \mathbb{E} \left[\mathbb{E} \left[\tilde{\nabla} f(y_k, w_k)^\top \tilde{\nabla} f(y_k, w_k) \middle| y_k \right] - \nabla f(y_k)^\top \nabla f(y_k) \right] \end{aligned} \quad (28)$$

$$\leq \mathbb{E} \left[(A\tilde{\xi}_k + B\nabla f(y_k))^\top P (A\tilde{\xi}_k + B\nabla f(y_k)) \right] + \sigma^2 \alpha^2 \tilde{P}_{11} \quad (29)$$

$$= \mathbb{E} \left[\begin{bmatrix} \tilde{\xi}_k \\ \nabla f(y_k) \end{bmatrix}^\top \begin{bmatrix} A^\top P A & A^\top P B \\ B^\top P A & B^\top P B \end{bmatrix} \begin{bmatrix} \tilde{\xi}_k \\ \nabla f(y_k) \end{bmatrix} \right] + \sigma^2 \alpha^2 \tilde{P}_{11}. \quad (30)$$

where in (27) we used the equality in (3) and the facts that we mentioned above. Also, (28) comes from the fact that $B^\top P B = \alpha^2 \tilde{P}_{11} I_d$ which can be shown by substituting B from (9) and using the assumption $P = \tilde{P} \otimes I_d$. Finally (29) follows from the inequality in (3), and (30) is obtained by writing the first term of (29) in matrix format. \square

Similarly, by extending Lemma 4.5 in [2] to the noise setting (3), for every $k \geq 0$ we obtain

$$\begin{aligned} \mathbb{E} \left[\begin{bmatrix} \xi_k - \xi^* \\ \nabla f(y_k) \end{bmatrix}^\top (\rho^2 X_1 + (1 - \rho^2) X_2) \begin{bmatrix} \xi_k - \xi^* \\ \nabla f(y_k) \end{bmatrix} \right] &\leq \rho^2 \mathbb{E}[f(x_k) - f^*] - \mathbb{E}[f(x_{k+1}) - f^*] \\ &\quad + \frac{L\alpha^2}{2} \sigma^2 \end{aligned}$$

where $X_1 = \tilde{X}_1 \otimes I_d$ and $X_2 = \tilde{X}_2 \otimes I_d$. The rest of the proof of Lemma 2.2 is very similar to the proof of Theorem 4.6 in [2], and we just need to use the fact that the Kronecker product of two positive semidefinite matrices is positive semidefinite [10].

C Proof of Theorem 2.3

Let

$$\Gamma \triangleq \rho^2 \tilde{X}_1 + (1 - \rho^2) \tilde{X}_2 - \begin{bmatrix} \tilde{A}^\top \tilde{P} \tilde{A} - \rho^2 \tilde{P} & \tilde{A}^\top \tilde{P} \tilde{B} \\ \tilde{B}^\top \tilde{P} \tilde{A} & \tilde{B}^\top \tilde{P} \tilde{B} \end{bmatrix}$$

with $\rho^2 = 1 - \sqrt{\alpha\mu}$ and $P = \tilde{P}_\alpha \otimes I_d$. According to Lemma 2.2, it suffices to show that $\Gamma \succeq 0_3$. Using the Symbolic toolbox in MATLAB, we see that Γ has the following properties

- (i) $\Gamma_{3,3} = \frac{\alpha(1 - L\alpha)}{2}$,
- (ii) $\Gamma_{2,3} = \Gamma_{3,2} = 0$,

$$\begin{aligned}
\text{(iii)} \quad \Gamma_{2,2} &= \frac{\mu\sqrt{\mu\alpha}(1-\sqrt{\mu\alpha})^2}{2(1+\sqrt{\mu\alpha})^2} + \frac{\mu(1-\sqrt{\mu\alpha})^3}{2(1+\sqrt{\mu\alpha})^2} - \frac{(1-\sqrt{\mu\alpha})^2}{2\alpha(1+\sqrt{\mu\alpha})^2} + \frac{(1-\sqrt{\mu\alpha})^3}{2\alpha} \\
&= \frac{(1-\sqrt{\mu\alpha})^2}{2\alpha(1+\sqrt{\mu\alpha})^2} \left(\alpha(\mu\sqrt{\mu\alpha} + \mu(1-\sqrt{\mu\alpha})) - 1 + (1-\sqrt{\mu\alpha})(1+\sqrt{\mu\alpha})^2 \right) \\
&= \frac{(1-\sqrt{\mu\alpha})^3\sqrt{\mu}}{2\sqrt{\alpha}(1+\sqrt{\mu\alpha})} \geq 0 \\
\text{(iv)} \quad \det(\Gamma) &= 0.
\end{aligned}$$

In fact, if $\alpha = \frac{1}{L}$, then

$$\Gamma = \frac{(1-\sqrt{\alpha\mu})^3\sqrt{\mu}}{2\sqrt{\alpha}(1+\sqrt{\alpha\mu})} \begin{bmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

which is positive semidefinite. Now, consider the case that $\alpha < \frac{1}{L}$. For any $\epsilon > 0$, let

$$\Gamma^\epsilon \triangleq \Gamma + \epsilon \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Note that, for any $\epsilon > 0$, (i) implies $\Gamma_{3,3}^\epsilon > 0$. This fact, along with (ii) and (iii), indicates that $\det(\Gamma_{[2:3],[2:3]}^\epsilon) > 0$. Hence,

$$\det(\Gamma^\epsilon) = \det(\Gamma) + \epsilon \det(\Gamma_{[2:3],[2:3]}^\epsilon) = \epsilon \det(\Gamma_{[2:3],[2:3]}^\epsilon) > 0$$

where the second equality comes from (iv). Therefore, the determinant of Γ^ϵ , itself, and two submatrices $\Gamma_{3,3}^\epsilon$ and $\Gamma_{[2:3],[2:3]}^\epsilon$ are all positive. Thus, by Sylvester's criterion, Γ^ϵ is positive definite for any $\epsilon > 0$. As a consequence, since $\Gamma = \lim_{\epsilon \rightarrow 0} \Gamma^\epsilon$, Γ is positive semidefinite.

D Proof of Theorem 3.1

Using Theorem 2.3, for every $k \geq 1$, we have

$$\begin{aligned}
\mathbb{E}[V_{P_\alpha}(\xi_{k+1})] &\leq (1 - \frac{c}{\sqrt{\kappa}})\mathbb{E}[V_{P_\alpha}(\xi_k)] + \frac{\sigma^2}{2L}c^2(1 + c^2) \\
&\leq (1 - \frac{c}{\sqrt{\kappa}})\mathbb{E}[V_{P_\alpha}(\xi_k)] + \frac{\sigma^2}{L}c^2
\end{aligned}$$

where in the last inequality we used the fact that $c^2 \leq 1$. Using this bound recursively for n times, we obtain

$$\begin{aligned}
\mathbb{E}[V_{P_\alpha}(\xi_{n+1})] &\leq (1 - \frac{c}{\sqrt{\kappa}})^n \mathbb{E}[V_{P_\alpha}(\xi_1)] + \frac{\sigma^2}{L}c^2 \sum_{i=1}^n (1 - \frac{c}{\sqrt{\kappa}})^{n-i} \\
&\leq \exp(-n \frac{c}{\sqrt{\kappa}}) \mathbb{E}[V_{P_\alpha}(\xi_1)] + \frac{\sigma^2}{L}c^2 \frac{1 - (1 - \frac{c}{\sqrt{\kappa}})^n}{1 - (1 - \frac{c}{\sqrt{\kappa}})} \\
&\leq \exp(-n \frac{c}{\sqrt{\kappa}}) \mathbb{E}[V_{P_\alpha}(\xi_1)] + \frac{\sigma^2}{L}c\sqrt{\kappa}
\end{aligned}$$

where the second inequality follows from the inequality that $1 - t \leq \exp(-t)$ for every $t \geq 0$, and the third inequality is obtained by replacing $1 - (1 - \frac{c}{\sqrt{\kappa}})^n$ by 1.

E Proof of Corollary 3.2

We first show $\alpha_1 \leq \frac{1}{L}$. Note that, by assumption, n can be written as $p\sqrt{\kappa}n_0$ where $n_0 \geq 2\log(p\sqrt{\kappa})$ and $n_0 \geq e$. This assumption, along with the fact that $\frac{\log n}{n}$ is a decreasing function of n as $n \geq e$,

implies

$$\begin{aligned} \frac{p\sqrt{\kappa} \log n}{n} &= \frac{p\sqrt{\kappa} \log(p\sqrt{\kappa}n_0)}{p\sqrt{\kappa}n_0} \\ &= \frac{\log(p\sqrt{\kappa}) + \log n_0}{n_0} \\ &\leq \frac{1}{2} + \frac{\log n_0}{n_0} \end{aligned} \quad (31)$$

$$\leq 1 \quad (32)$$

where in (31) we used the assumption $n_0 \geq 2 \log(p\sqrt{\kappa})$, and (32) follows from the fact that $n_0 \geq e$, and therefore, $\frac{\log n_0}{n_0} \leq 1/e \leq 1/2$.

Next, using Theorem 3.1 with $c = \frac{p\sqrt{\kappa} \log n}{n}$ immediately gives the desired bound.

F Proof of Lemma 3.3

First, note that $\xi_1^{k+1} = [x_{n_k+1}^k]^\top, x_{n_k+1}^k]^\top$, and therefore,

$$\begin{aligned} &(\xi_1^{k+1} - \xi^*)^\top P_{\alpha_{k+1}} (\xi_1^{k+1} - \xi^*) \\ &= (\xi_1^{k+1} - \xi^*)^\top \begin{bmatrix} \sqrt{\frac{1}{2\alpha_{k+1}}} \\ \sqrt{\frac{\mu}{2}} - \sqrt{\frac{1}{2\alpha_{k+1}}} \end{bmatrix} \begin{bmatrix} \sqrt{\frac{1}{2\alpha_{k+1}}} & \sqrt{\frac{\mu}{2}} - \sqrt{\frac{1}{2\alpha_{k+1}}} \end{bmatrix} (\xi_1^{k+1} - \xi^*) \\ &= \|x_{n_k+1}^k - x^*\|^2 \left(\frac{1}{2\alpha_{k+1}} + \left(\sqrt{\frac{\mu}{2}} - \sqrt{\frac{1}{2\alpha_{k+1}}} \right)^2 + \frac{\sqrt{2}}{\sqrt{\alpha_{k+1}}} \left(\sqrt{\frac{\mu}{2}} - \sqrt{\frac{1}{2\alpha_{k+1}}} \right) \right) \\ &= \frac{\mu}{2} \|x_{n_k+1}^k - x^*\|^2. \end{aligned} \quad (33)$$

Plugging (33) into (10) for $V_{P_{\alpha_{k+1}}}(\xi_1^{k+1})$ yields

$$\begin{aligned} V_{P_{\alpha_{k+1}}}(\xi_1^{k+1}) &= (\xi_1^{k+1} - \xi^*)^\top P_{\alpha_{k+1}} (\xi_1^{k+1} - \xi^*) + f(x_{n_k+1}^k) - f^* \\ &= \frac{\mu}{2} \|x_{n_k+1}^k - x^*\|^2 + f(x_{n_k+1}^k) - f^* \\ &\leq 2(f(x_{n_k+1}^k) - f^*) \end{aligned} \quad (34)$$

$$\leq 2V_{P_{\alpha_k}}(\xi_{n_k+1}^k) \quad (35)$$

where (34) follows from (2) with $x = x_{n_k+1}^k$ and $y = x^*$. Finally, taking expectation from (35) completes the proof.

G Proof of Theorem 3.4

We claim that for every $k \geq 1$

$$\mathbb{E} [V_{P_{\alpha_k}}(\xi_{n_k+1}^k)] \leq \frac{2}{2^{(p+1)(k-1)}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^0) - f^*)) + \frac{\sigma^2 \sqrt{\kappa}}{L 2^{k-1}}. \quad (36)$$

which implies (17) as $V_{P_{\alpha_k}}(\xi_{n_k+1}^k) \geq f(x_{n_k+1}^k) - f^*$. We show (36) by induction on k . For $k = 1$, using Theorem 3.1, we obtain

$$\begin{aligned} \mathbb{E} [V_{P_{\alpha_1}}(\xi_{n_1+1}^1)] &\leq \exp\left(-\frac{n_1}{\sqrt{\kappa}}\right) \mathbb{E} [V_{P_{\alpha_1}}(\xi_1)] + \frac{\sigma^2 \sqrt{\kappa}}{L} \\ &\leq 2 \exp\left(-\frac{n_1}{\sqrt{\kappa}}\right) (f(x_0^0) - f^*) + \frac{\sigma^2 \sqrt{\kappa}}{L} \end{aligned} \quad (37)$$

where the second inequality comes from the inequality $V_{P_{\alpha_1}}(\xi_1) \leq 2(f(x_0^0) - f^*)$ which can be derived similar to (34).

Next, we assume (36) holds for k and show it also holds for $k + 1$. Note that

$$\mathbb{E} [V_{P_{\alpha_{k+1}}}(\xi_{n_{k+1}+1}^{k+1})] \leq \exp \left(-n_{k+1} \sqrt{\frac{\alpha_{k+1}L}{\kappa}} \right) \mathbb{E} [V_{P_{\alpha_{k+1}}}(\xi_1^{k+1})] + \frac{\sigma^2 \sqrt{\kappa} \sqrt{\alpha_{k+1}L}}{L} \quad (38)$$

$$\begin{aligned} &= \exp(-\log(2^{p+2})) \mathbb{E} [V_{P_{\alpha_{k+1}}}(\xi_1^{k+1})] + \frac{\sigma^2 \sqrt{\kappa}}{2^{k+1}L} \\ &\leq \frac{2}{2^{p+2}} \mathbb{E} [V_{P_{\alpha_k}}(\xi_{n_k+1}^k)] + \frac{\sigma^2 \sqrt{\kappa}}{2^{k+1}L} \end{aligned} \quad (39)$$

$$\leq \frac{1}{2^{p+1}} \cdot \frac{2}{2^{(p+1)(k-1)}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^0) - f^*)) + \frac{\sigma^2 \sqrt{\kappa}}{L} \left(\frac{1}{2^{k+1}} + \frac{1}{2^{p+k}} \right) \quad (40)$$

$$\leq \frac{2}{2^{(p+1)k}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^0) - f^*)) + \frac{\sigma^2 \sqrt{\kappa}}{L2^k} \quad (41)$$

where, (38) and (39) are obtained by using Theorem 3.1 and Lemma 3.3, respectively, and in (40) we used the assumption that (36) holds for k .

H Proof of Theorem 3.6

First, we will show that for every $k \geq 1$ and $0 \leq m \leq n_k + 1$, we have

$$\mathbb{E} [f(x_m^{k+1})] - f^* \leq \frac{4}{2^{(p+1)(k-1)}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^0) - f^*)) + \frac{9\sigma^2 \sqrt{\kappa}}{4L2^{k-1}}. \quad (42)$$

Indeed,

$$\begin{aligned} \mathbb{E} [f(x_m^{k+1})] - f^* &\leq \mathbb{E} [V_{P_{\alpha_{k+1}}}(\xi_m^{k+1})] \\ &\leq \exp \left(-m \sqrt{\frac{\alpha_{k+1}L}{\kappa}} \right) \mathbb{E} [V_{P_{\alpha_{k+1}}}(\xi_1^{k+1})] + \frac{\sigma^2 \sqrt{\kappa} \sqrt{\alpha_{k+1}L}}{L} \end{aligned} \quad (43)$$

$$\leq 2\mathbb{E} [V_{P_{\alpha_k}}(\xi_{n_k+1}^k)] + \frac{\sigma^2 \sqrt{\kappa}}{2^{k+1}L} \quad (44)$$

$$\leq \frac{4}{2^{(p+1)(k-1)}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^0) - f^*)) + \frac{9\sigma^2 \sqrt{\kappa}}{4L2^{k-1}} \quad (45)$$

where, (43) and (44) follows again from Theorem 3.1 and Lemma 3.3, and we obtain (45) using (36).

Recall the definition $N_K(p, n_1) \triangleq \sum_{k=1}^K n_k$ which denotes the total number of stochastic gradient iterations required to complete K stages of M-ASG for parameter p and first-stage iteration number n_1 fixed. Given the computational budget of n iterations such that $n \geq n_1$, let K be the largest number such that $n \geq N_K \triangleq N_K(p, n_1)$. As a result, at iteration n we are in stage $K + 1$. Note that (18) implies $2^{K-1} \geq \frac{N_K - n_1 + \Psi}{\Psi}$ with $\Psi = 4[(p+2)\sqrt{\kappa} \log(2)]$; therefore

$$\frac{1}{2^{K-1}} \leq \frac{\Psi}{N_K - n_1 + \Psi}. \quad (46)$$

Thus, we get the following upper bound on the suboptimality:

$$\begin{aligned} \mathbb{E} [f(x_n)] - f^* &= \mathbb{E} [f(x_{n-N_K}^{K+1})] - f^* \\ &\leq \frac{4}{2^{(p+1)(K-1)}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^0) - f^*)) + \frac{9\sigma^2 \sqrt{\kappa}}{4L2^{K-1}}, \end{aligned} \quad (47)$$

and by substituting (46) in (47) we obtain the bound

$$\begin{aligned} &\mathbb{E} [f(x_n)] - f^* \\ &\leq \mathcal{O}(1) \left(\frac{(4(p+1)\sqrt{\kappa} \log(2))^{p+1}}{(N_K - n_1 + \Psi)^{p+1}} (\exp(-n_1/\sqrt{\kappa})(f(x_0^0) - f^*)) + \frac{(p+1)\sigma^2}{(N_K - n_1 + \Psi)\mu} \right). \end{aligned} \quad (48)$$

Next, by (18), $N_{K+1} - n_1 \leq 2(N_K - n_1 + \Psi)$, and thus, $n - n_1 \leq 2(N_K - n_1 + \Psi)$. Replacing $\frac{1}{N_K - n_1 + \Psi}$ by $\frac{2}{n - n_1}$ in (48) completes the proof of (19).

I Proof of Corollary 3.7

Note that by setting $n_1 = \lceil (p+1)\sqrt{\kappa} \log(12(p+1)\kappa) \rceil$, we have $\exp(-n_1/\sqrt{\kappa}) \leq \frac{1}{(16 \log(2)(p+1)\sqrt{\kappa})^{p+1}}$; hence, plugging $n_1 = \lceil (p+1)\sqrt{\kappa} \log(12(p+1)\kappa) \rceil$ in (19) implies the following bound with an $\mathcal{O}(1)$ constant that does not depend on μ , L and x_0^0 :

$$\mathcal{O}(1) \left(\frac{2^{-(p+1)}}{(n-n_1)^{p+1}} (f(x_0) - f^*) + \frac{(p+1)\sigma^2}{(n-n_1)\mu} \right). \quad (49)$$

Finally, note that $n \geq 2n_1$; therefore, $n - n_1 \geq n/2$ and using it in (49) completes the proof.

J Proof of Corollary 3.9

By plugging $p = 1$ and $n_1 = \lceil \sqrt{\kappa} \log(\frac{4\Delta}{\epsilon}) \rceil$ in (17), it is straightforward to check the bias term is bounded by $\frac{\epsilon}{2}$. Next, consider running M-ASG with given parameters, possibly without knowing and/or specifying the exact number of stages. Consider the end of the K -th stage, where $K \triangleq \lceil \log_2(\frac{\sigma^2 \sqrt{\kappa}}{L\epsilon}) \rceil + 2$. Since $\frac{1}{2^{K-1}} \leq \frac{L\epsilon}{2\sigma^2 \sqrt{\kappa}}$, the variance term in (17) is also bounded by $\frac{\epsilon}{2}$, and as a result $x_{n_K+1}^K$ is an ϵ -solution.

Now, by using (18), we can bound the number of iterations for completing K stages:

$$N_K = n_1 + (2^{K+1} - 4) \lceil \sqrt{\kappa} \log(8) \rceil \leq n_1 + 2(1 + \log(8)) (2^K \sqrt{\kappa}) \quad (50)$$

$$\leq \lceil \sqrt{\kappa} \log\left(\frac{4\Delta}{\epsilon}\right) \rceil + 16(1 + \log(8)) \frac{\sigma^2}{\mu\epsilon} \quad (51)$$

where in (50) we used the fact that $\lceil \sqrt{\kappa} \log(8) \rceil \leq (1 + \frac{1}{\log(8)}) \sqrt{\kappa} \log(8)$ since $\kappa \geq 1$, and (51) follows from the bound $K \leq 3 + \log_2(\frac{\sigma^2 \sqrt{\kappa}}{L\epsilon})$.

K Results for More General Noise Setting

In this section, we show how our analysis can be extended to a more general noise setting where the bound on the variance can depend on the distance to the optimal solution. More formally, we assume that at $x \in \mathbb{R}^d$, we have access to the noisy gradient $\tilde{\nabla} f(x, w)$ such that for some $\sigma, \eta \geq 0$,

$$\begin{aligned} \mathbb{E}[\tilde{\nabla} f(x, w)|x] &= \nabla f(x) \\ \mathbb{E}[\|\tilde{\nabla} f(x, w) - \nabla f(x)\|^2|x] &\leq \sigma^2 + \eta^2 \|x - x^*\|^2, \end{aligned} \quad (52)$$

where w is a random variable independent of previous iterates.

In what follows, we first show how the results of Theorem 2.3 and Lemma 3.3 extends to this setting, and then briefly discuss the results of our multistage scheme for this noise setting.

Theorem K.1. *Let $f \in S_{\mu, L}(\mathbb{R}^d)$ with $\kappa \geq 4$. Consider the ASG iterations given in (7) under noise model in (52). For $\alpha \in (0, \bar{\alpha}]$ and $\beta = \frac{1 - \sqrt{\alpha\mu}}{1 + \sqrt{\alpha\mu}}$ with*

$$\bar{\alpha} := \begin{cases} \min\{\frac{1}{L}, \frac{\mu^3}{(60\eta^2)^2}\} & \text{if } \eta > 0, \\ \frac{1}{L} & \text{if } \eta = 0, \end{cases} \quad (53)$$

it follows that

$$\mathbb{E}[V_{Q_\alpha}(\xi_{k+1})] \leq (1 - \sqrt{\alpha\mu}/3) \mathbb{E}[V_{Q_\alpha}(\xi_k)] + 2\sigma^2\alpha, \quad (54)$$

for every $k \geq 0$, where $Q_\alpha = \tilde{Q}_\alpha \otimes I_d$ with

$$\tilde{Q}_\alpha = \tilde{P}_\alpha + 2\alpha\eta^2 \tilde{C}^\top \tilde{C} = \begin{bmatrix} \sqrt{\frac{1}{2\alpha}} \\ \sqrt{\frac{\mu}{2}} - \sqrt{\frac{1}{2\alpha}} \end{bmatrix} \begin{bmatrix} \sqrt{\frac{1}{2\alpha}} & \sqrt{\frac{\mu}{2}} - \sqrt{\frac{1}{2\alpha}} \end{bmatrix} + 2\alpha\eta^2 \begin{bmatrix} 1 + \beta \\ -\beta \end{bmatrix} \begin{bmatrix} 1 + \beta & -\beta \end{bmatrix}.$$

Proof. First, note that similar to the proof of Lemma 2.2 and by using $\alpha L \leq 1$, we can show

$$\mathbb{E}[V_{P_\alpha}(\xi_{k+1})] \leq (1 - \sqrt{\alpha\mu})\mathbb{E}[V_{P_\alpha}(\xi_k)] + \alpha\sigma^2 + \alpha\eta^2\mathbb{E}[\|y_k - x^*\|^2]. \quad (55)$$

Using $y_k = C\xi_k$, we can substitute $\|y_k - x^*\|^2$ by $(\xi_k - \xi^*)^\top C^\top C(\xi_k - \xi^*)$ in (55); hence,

$$\begin{aligned} \mathbb{E}[V_{P_\alpha}(\xi_{k+1})] &\leq (1 - \sqrt{\alpha\mu})\mathbb{E}[V_{P_\alpha}(\xi_k)] + \frac{1}{2}\mathbb{E}[(\xi_k - \xi^*)^\top 2\alpha\eta^2 C^\top C(\xi_k - \xi^*)] + \alpha\sigma^2 \\ &\leq (1 - \sqrt{\alpha\mu})\mathbb{E}[V_{Q_\alpha}(\xi_k)] + \alpha\sigma^2, \end{aligned} \quad (56)$$

where the last inequality follows from $1 - \sqrt{\alpha\mu} \geq 1/2$ which is true since $\alpha \leq 1/L$ and $\kappa \geq 4$. Also note that

$$\begin{aligned} (\xi_k - \xi^*)^\top C^\top C(\xi_k - \xi^*) &= \|(1 + \beta)(x_{k+1} - x^*) - \beta(x_k - x^*)\|^2 \\ &\leq 2(1 + \beta)^2\|x_{k+1} - x^*\|^2 + 2\beta^2\|x_k - x^*\|^2 \end{aligned} \quad (57)$$

$$\leq \frac{16}{\mu}(f(x_{k+1}) - f^*) + \frac{4}{\mu}(f(x_k) - f^*) \quad (58)$$

$$\leq \frac{16}{\mu}V_{P_\alpha}(\xi_{k+1}) + \frac{4}{\mu}V_{Q_\alpha}(\xi_k). \quad (59)$$

where (57) follows from $(a + b)^2 \leq 2a^2 + 2b^2$ and (58) follows from $\beta \leq 1$ and the strong convexity assumption, i.e., $f(x) - f^* \geq \frac{\mu}{2}\|x - x^*\|^2$. Finally, (59) is obtained using $\min\{V_{P_\alpha}(\xi_k), V_{Q_\alpha}(\xi_k)\} \geq f(x_k) - f^*$. Plugging (59) into the definition of $V_{Q_\alpha}(\xi_{k+1})$ implies

$$\begin{aligned} \mathbb{E}[V_{Q_\alpha}(\xi_{k+1})] &= \mathbb{E}[V_{P_\alpha}(\xi_{k+1})] + 2\alpha\eta^2\mathbb{E}[(\xi_k - \xi^*)^\top C^\top C(\xi_k - \xi^*)] \\ &\leq (1 + \frac{32\alpha\eta^2}{\mu})\mathbb{E}[V_{P_\alpha}(\xi_{k+1})] + \frac{8\alpha\eta^2}{\mu}V_{Q_\alpha}(\xi_k). \end{aligned} \quad (60)$$

Using this result along with (56) yields

$$\begin{aligned} \mathbb{E}[V_{Q_\alpha}(\xi_{k+1})] &\leq \left((1 - \sqrt{\alpha\mu})(1 + \frac{32\alpha\eta^2}{\mu}) + \frac{8\alpha\eta^2}{\mu} \right) \mathbb{E}[V_{Q_\alpha}(\xi_k)] + \alpha(1 + \frac{32\alpha\eta^2}{\mu})\sigma^2 \\ &\leq \left(1 - \sqrt{\alpha\mu} + \frac{40\alpha\eta^2}{\mu} \right) \mathbb{E}[V_{Q_\alpha}(\xi_k)] + \alpha(1 + \frac{32\alpha\eta^2}{\mu})\sigma^2 \\ &\leq (1 - \sqrt{\alpha\mu}/3)\mathbb{E}[V_{Q_\alpha}(\xi_k)] + (1 + \frac{32\alpha\eta^2}{\mu})\alpha\sigma^2 \end{aligned} \quad (61)$$

where the last inequality follows from the assumption $\alpha \leq \frac{\mu^3}{(60\eta^2)^2}$ which implies

$$\frac{2}{3}\sqrt{\alpha\mu} \geq \frac{40\alpha\eta^2}{\mu}. \quad (62)$$

Finally, note that, (62) along with $\alpha \leq 1/L$ also implies $1 \geq 60\alpha\eta^2/\mu$; thus, we can bound $1 + 32\alpha\eta^2/\mu$ in (61) by 2 which gives us the desired result. \square

Next, note that we can also extend Lemma 3.3 to the new Lyapunov function $\mathbb{E}[V_{Q_\alpha}(\xi_k)]$ as well:

Lemma K.2. *Let $f \in S_{\mu,L}(\mathbb{R}^d)$. Consider M-ASG, i.e., Algorithm 1, assuming the noise model in (52), with $\alpha \in (0, \bar{\alpha}]$, where $\bar{\alpha}$ is defined by (53). Then, for every $1 \leq k \leq K - 1$,*

$$\mathbb{E}[V_{Q_{\alpha_{k+1}}}(\xi_1^{k+1})] \leq 3\mathbb{E}[V_{Q_{\alpha_k}}(\xi_{n_k+1}^k)]. \quad (63)$$

Proof. The proof is very similar to the arguments in Appendix F. In particular,

$$\begin{aligned} V_{Q_{\alpha_{k+1}}}(\xi_1^{k+1}) &= (\xi_1^{k+1} - \xi^*)^\top Q_{\alpha_{k+1}}(\xi_1^{k+1} - \xi^*) + f(x_{n_k+1}^k) - f^* \\ &= \frac{\mu}{2}\|x_{n_k+1}^k - x^*\|^2 + 2\alpha\eta^2(\xi_1^{k+1} - \xi^*)^\top C^\top C(\xi_1^{k+1} - \xi^*) + f(x_{n_k+1}^k) - f^* \\ &= (\frac{\mu}{2} + 2\alpha\eta^2)\|x_{n_k+1}^k - x^*\|^2 + f(x_{n_k+1}^k) - f^* \\ &\leq (2 + \frac{4\alpha\eta^2}{\mu})(f(x_{n_k+1}^k) - f^*) \end{aligned} \quad (64)$$

$$\leq 3V_{Q_{\alpha_k}}(\xi_{n_k+1}^k) \quad (65)$$

where (64) follows from (2) with $x = x_{n_k+1}^k$ and $y = x^*$. Finally, (65) follows from $1/60 \geq \alpha\eta^2/\mu$, which holds due to (62) and $\alpha \leq \frac{1}{L}$, along with $V_{Q_{\alpha_k}}(\xi_{n_k+1}^k) \geq f(x_{n_k+1}^k) - f^*$. Taking expectations of both sides of (65) completes the proof. \square

Using the results in Theorem K.1 and Lemma K.2, we can analyze M-ASG for this more general noise setting in (52) as well and extend our complexity result in Corollary 3.8 as follows:

$$\mathbb{E}[f(x_n)] - f^* \leq \mathcal{O}(1) \left(\exp(-n/(\Theta(1)(\sqrt{\kappa} + \eta^2/\mu^2))) (f(x_0^0) - f^*) + \frac{\sigma^2}{n\mu} \right)$$

for n sufficiently large and known in advance. It is worth noting that we can also derive similar results to Theorems 3.4 and 3.6 when n is not known. We skip the details as all the arguments follow very similar to our analysis in Section 3.

L M-ASG for Convex Objective Functions

For merely convex objective functions, as discussed in [22], the suboptimality $\mathbb{E}[f(x_n)] - f^*$ admits the lower bound given below:

$$\Theta(1) \left(\frac{L}{n^2} \|x_0 - x^*\|_2^2 + \frac{\sigma^2}{\sqrt{n}} \right). \quad (66)$$

The author of [22] obtains this lower bound for the case of compact domain with the additional knowledge of noise parameter σ . For unconstrained optimization, and without using the information on the noise parameter, σ^2 , it is shown in [8] that one can achieve the rate $\mathcal{O}(\frac{1}{\sqrt{n}})$ in both bias and variance terms (see last part of Corollary 3.9 and also Corollary 4.1 in [8]). As we state below, a direct application of our current results recovers a similar result up to a log factor.

Theorem L.1. *Let f be a merely convex function, i.e., $f \in S_{0,L}(\mathbb{R}^d)$ with $\mu = 0$, and let $n \geq 2$ be the given iteration budget. Define $f_\lambda(x) \triangleq f(x) + \frac{\lambda}{2} \|x - x_0\|^2$ with $\lambda \triangleq L/(\sqrt{n} - 1)$. Consider running ASG, given in (7), with stepsize $\alpha = \frac{(\log n)^2}{n^{3/2}L}$ for solving $\min_x f_\lambda(x)$. Then,*

$$\mathbb{E}[f(x_{n+1})] - f^* \leq \frac{2}{n} (f(x_0) - f^*) + \frac{L}{\sqrt{n}} \|x_0 - x^*\|^2 + \frac{\sigma^2 \log n}{\sqrt{n}L}. \quad (67)$$

Proof. Define $f_\lambda^* \triangleq \min_x f_\lambda(x)$. Note $f_\lambda \in S_{\lambda,L+\lambda}(\mathbb{R}^d)$; thus, using Theorem 3.1 with $c = \log n/n^{3/4}$ and $\kappa = (L + \lambda)/\lambda = \sqrt{n}$ implies

$$\begin{aligned} \mathbb{E}[f_\lambda(x_{n+1}^1)] - f_\lambda^* &\leq \mathbb{E}[V_{P_\alpha}(\xi_{n+1})] \\ &\leq \exp(-n \frac{c}{\sqrt{\kappa}}) \mathbb{E}[V_{P_\alpha}(\xi_1)] + \frac{\sigma^2 \sqrt{\kappa} c}{L + \lambda} \\ &\leq \frac{1}{n} \mathbb{E}[V_{P_\alpha}(\xi_1)] + \frac{\sigma^2 \log n}{\sqrt{n}L}. \end{aligned} \quad (68)$$

Now, using the fact that $x_0 = x_{-1}$, and similar to the proof of Lemma 3.3, we can show

$$\mathbb{E}[V_{P_\alpha}(\xi_1)] \leq 2(f_\lambda(x_0) - f_\lambda^*) = 2(f(x_0) - f_\lambda^*).$$

Therefore, plugging this into (68), we obtain

$$\mathbb{E}[f_\lambda(x_{n+1})] - f_\lambda^* \leq \frac{2}{n} (f(x_0) - f_\lambda^*) + \frac{\sigma^2 \log n}{\sqrt{n}L}, \quad (69)$$

which is equivalent to

$$\mathbb{E}[f_\lambda(x_{n+1})] - (1 - \frac{2}{n})f_\lambda^* \leq \frac{2}{n} f(x_0) + \frac{\sigma^2 \log n}{\sqrt{n}L}. \quad (70)$$

This result along with $f(x_{n+1}) \leq f_\lambda(x_{n+1})$ implies

$$\mathbb{E}[f(x_{n+1})] - (1 - \frac{2}{n})f_\lambda^* \leq \frac{2}{n} f(x_0) + \frac{\sigma^2 \log n}{\sqrt{n}L}. \quad (71)$$

Finally, using the bound

$$f_\lambda^* \leq f_\lambda(x^*) = f^* + \frac{\lambda}{2} \|x_0 - x^*\|^2$$

completes the proof. \square

M AC-SA from the perspective of Nesterov's Accelerated Method

Recall that AC-SA [16] with initial point x_0 and sequence of stepsize parameters $\{\eta_t\}_{t \geq 1}$ and $\{\gamma_t\}_{t \geq 1}$ has the following update rule:

- (i) Set $x_0^{ag} = x_0$ and $t = 1$;
- (ii) Set $x_t^{md} = \frac{(1-\eta_t)(\mu+\gamma_t)x_{t-1}^{ag} + \eta_t[(1-\eta_t)\mu+\gamma_t]x_{t-1}}{\gamma_t + (1-\eta_t^2)\mu}$;
- (iii) Set $x_t = \frac{\eta_t \mu x_t^{md} + [(1-\eta_t)\mu+\gamma_t]x_{t-1} - \eta_t G_t}{\mu+\gamma_t}$ where $G_t = \tilde{\nabla} f(x_t^{md}, w_t)$;
- (iv) Set $x_t^{ag} = \eta_t x_t + (1-\eta_t)x_{t-1}^{ag}$;
- (v) Set $t \leftarrow t + 1$ and go to step (ii).

We claim that this algorithm can be cast as an ASG method in (7) with a specific varying stepsize rule. In fact, we show it can be represented as

$$x_t^{md} = (1 + \tilde{\beta}_t)x_{t-1}^{ag} - \tilde{\beta}_t x_{t-2}^{ag} \quad (72a)$$

$$x_t^{ag} = x_t^{md} - \tilde{\alpha}_t \tilde{\nabla} f(x_t^{md}, w_t). \quad (72b)$$

with

$$\tilde{\alpha}_t = \frac{\eta_t^2}{\mu + \gamma_t}, \quad \tilde{\beta}_t = \frac{\eta_t(1 - \eta_{t-1})[(1 - \eta_t)\mu + \gamma_t]}{\eta_{t-1}[\gamma_t + (1 - \eta_t^2)\mu]}.$$

To show this, first, multiplying both sides of ((ii)) by $\frac{\gamma_t + (1 - \eta_t^2)\mu}{\mu + \gamma_t}$ implies

$$\frac{\gamma_t + (1 - \eta_t^2)\mu}{\mu + \gamma_t} x_t^{md} = (1 - \eta_t)x_{t-1}^{ag} + \frac{\eta_t[(1 - \eta_t)\mu + \gamma_t]}{\mu + \gamma_t} x_{t-1}, \quad (73)$$

and by substituting $(1 - \eta_t)x_{t-1}^{ag}$ by $x_{t-1}^{ag} - \eta_t x_t$ from ((iv)) we obtain

$$x_t^{ag} = \eta_t \left(x_t - \frac{(1 - \eta_t)\mu + \gamma_t}{\mu + \gamma_t} x_{t-1} \right) + \frac{\gamma_t + (1 - \eta_t^2)\mu}{\mu + \gamma_t} x_t^{md}. \quad (74)$$

Note that, by ((iii)), we have

$$x_t - \frac{(1 - \eta_t)\mu + \gamma_t}{\mu + \gamma_t} x_{t-1} = \frac{\eta_t \mu}{\mu + \gamma_t} x_t^{md} - \frac{\eta_t}{\mu + \gamma_t} G_t, \quad (75)$$

and therefore, (74) and (75) yield

$$\begin{aligned} x_t^{ag} &= \eta_t \left(\frac{\eta_t \mu}{\mu + \gamma_t} x_t^{md} - \frac{\eta_t}{\mu + \gamma_t} G_t \right) + \frac{\gamma_t + (1 - \eta_t^2)\mu}{\mu + \gamma_t} x_t^{md} \\ &= \left(\frac{\eta_t^2 \mu}{\mu + \gamma_t} + \frac{\gamma_t + (1 - \eta_t^2)\mu}{\mu + \gamma_t} \right) x_t^{md} - \frac{\eta_t^2}{\mu + \gamma_t} G_t \\ &= x_t^{md} - \frac{\eta_t^2}{\mu + \gamma_t} G_t \end{aligned}$$

which implies (72b).




Figure 4: Comparison of GD, AGD, μ AGD+, and MASG for logistic regression with $n = 1000$ iterations with different level of noise.




Figure 5: Comparison of GD, AGD, μ AGD+, and MASG for logistic regression with $n = 10000$ iterations with different level of noise.

To show (72a), first note that by (iv) for $t - 1$, we obtain $x_{t-1} = \frac{1}{\eta_{t-1}}(x_{t-1}^{ag} - (1 - \eta_{t-1})x_{t-2}^{ag})$. Plugging in this in (ii), leads to

$$\begin{aligned}
 x_t^{md} &= \frac{(1 - \eta_t)(\mu + \gamma_t)}{\gamma_t + (1 - \eta_t^2)\mu} x_{t-1}^{ag} + \frac{\eta_t[(1 - \eta_t)\mu + \gamma_t]}{\eta_{t-1}[\gamma_t + (1 - \eta_t^2)\mu]} (x_{t-1}^{ag} - (1 - \eta_{t-1})x_{t-2}^{ag}) \\
 &= \left(\frac{(1 - \eta_t)(\mu + \gamma_t)}{\gamma_t + (1 - \eta_t^2)\mu} + \frac{\eta_t[(1 - \eta_t)\mu + \gamma_t]}{\eta_{t-1}[\gamma_t + (1 - \eta_t^2)\mu]} \right) x_{t-1}^{ag} - \frac{\eta_t(1 - \eta_{t-1})[(1 - \eta_t)\mu + \gamma_t]}{\eta_{t-1}[\gamma_t + (1 - \eta_t^2)\mu]} x_{t-2}^{ag} \\
 &= \frac{\eta_{t-1}(1 - \eta_t)(\mu + \gamma_t) + \eta_t[(1 - \eta_t)\mu + \gamma_t]}{\eta_{t-1}[\gamma_t + (1 - \eta_t^2)\mu]} x_{t-1}^{ag} - \tilde{\beta}_t x_{t-2}^{ag} \\
 &= \frac{\eta_{t-1}(1 - \eta_t)(\mu + \gamma_t) + \eta_t \eta_{t-1}[(1 - \eta_t)\mu + \gamma_t]}{\eta_{t-1}[\gamma_t + (1 - \eta_t^2)\mu]} x_{t-1}^{ag} + \tilde{\beta}_t (x_{t-1}^{ag} - x_{t-2}^{ag}) \\
 &= x_{t-1}^{ag} + \tilde{\beta}_t (x_{t-1}^{ag} - x_{t-2}^{ag})
 \end{aligned}$$

which is (72a) and the proof is complete.

As a consequence, Multistage AC-SA is a variant of M-ASG Algorithm that has a different length n_k for each stage $k \geq 1$ and employs a specific varying stepsize rule together with a different selection for the momentum parameter at each stage.

N Additional Numerical Experiments

In this section, we study another classification problem using logistic regression, but with a synthesized data. In particular, we generate a random matrix $M \in \mathbb{R}^{2000 \times 100}$ and a random vector $w \in \mathbb{R}^{100}$ and compute $y = \text{sign}(Mw)$ which is the vector that contains the sign of the inner product with the rows of M and the vector w . Our goal is to recover w by optimizing a regularized logistic

objective when the gradient of the loss function is corrupted with additive Gaussian noise. We compare M-ASG and M-ASG* with Standard GD, Standard AG, μ AGD+ [8], and Multistage AC-SA [17]. We note that the condition number of the problem $\kappa \sim 1000$ for this problem. Figures 4–5 illustrate the behavior of the algorithms for $n = 1000$ and $n = 10000$ iterations for the noise level $\sigma_n^2 \in \{10^{-6}, 10^{-4}, 10^{-2}\}$ as before. It can be seen that both M-ASG and M-ASG* usually start faster, and do not perform worse than other algorithms in different scenarios; moreover, they outperform other algorithms when the iteration budget is limited or the noise level is small. Furthermore, note that in the setting where the noise is large, M-ASG* behaves better than M-ASG, as it terminates the first stage earlier, which is helpful as the noise is large; hence, the variance becomes term dominant in the first stage just after a few iterations.