

# A UNIFIED THEORY OF STOCHASTIC PROXIMAL POINT METHODS WITHOUT SMOOTHNESS

Anonymous authors

Paper under double-blind review

## ABSTRACT

This paper presents a comprehensive analysis of a broad range of variations of the stochastic proximal point method (SPPM). Proximal point methods have attracted considerable interest owing to their numerical stability and robustness against imperfect tuning, a trait not shared by the dominant stochastic gradient descent (SGD) algorithm. A framework of assumptions that we introduce encompasses methods employing techniques such as variance reduction and arbitrary sampling. A cornerstone of our general theoretical approach is a parametric assumption on the iterates, correction and control vectors. We establish a single theorem that ensures linear convergence under this assumption and the  $\mu$ -strong convexity of the loss function, and without the need to invoke smoothness. This integral theorem reinstates best known complexity and convergence guarantees for several existing methods which demonstrate the robustness of our approach. We expand our study by developing three new variants of SPPM, and through numerical experiments we elucidate various properties inherent to them.

## 1 INTRODUCTION

In this paper we consider the stochastic optimization problem

$$\min_{x \in \mathbb{R}^d} \{f(x) := \mathbb{E}_{\xi \sim \mathcal{D}} [f_{\xi}(x)]\}, \quad (1)$$

where  $\xi \sim \mathcal{D}$  is a random variable following distribution  $\mathcal{D}$ , and  $\mathbb{E}[\cdot]$  denotes mathematical expectation. Problems of this type are fundamental for statistical supervised learning theory. Here,  $x$  is a machine learning model of  $d \in \mathbb{N}$  parameters,  $\mathcal{D}$  is an unknown distribution of labeled examples, samples  $\xi \sim \mathcal{D}$  are available,  $f_{\xi}$  is the loss on datapoint  $\xi$ ,  $f$  is the generalization error. In such a setup, an unbiased estimator of the gradient  $\nabla f_{\xi}(x)$  is computed instead of the gradient  $\nabla f(x)$ . We rely on two assumptions, presented next.

**Assumption 1.** *Function  $f_{\xi} : \mathbb{R}^d \rightarrow \mathbb{R}$  is differentiable for all samples  $\xi \sim \mathcal{D}$ .*

We implicitly assume that the order of differentiation and expectation can be swapped, which means that  $\nabla f(x) \stackrel{(1)}{=} \nabla \mathbb{E}_{\xi \sim \mathcal{D}} [f_{\xi}(x)] = \mathbb{E}_{\xi \sim \mathcal{D}} [\nabla f_{\xi}(x)]$ . This implies that  $f$  is differentiable.

**Assumption 2.** *Function  $f_{\xi} : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $\mu$ -strongly convex for all samples  $\xi \sim \mathcal{D}$ , where  $\mu > 0$ :*

$$f_{\xi}(y) + \langle \nabla f_{\xi}(y), x - y \rangle + \frac{\mu}{2} \|x - y\|^2 \leq f_{\xi}(x), \quad \forall x, y \in \mathbb{R}^d. \quad (2)$$

This implies that  $f$  is  $\mu$ -strongly convex, and hence  $f$  has a unique minimizer, which we denote by  $x_{\star}$ . We know that  $\nabla f(x_{\star}) = 0$ . Notably, we do *not* assume  $f$  to be  $L$ -smooth.

Another type of problem considered in the paper is a minimization of functions  $f$  that are averages of a large number of differentiable functions:

$$\min_{x \in \mathbb{R}^d} \left\{ f(x) := \frac{1}{n} \sum_{i=1}^n f_i(x) \right\}. \quad (3)$$

This is a special case of (1), with  $\mathcal{D}$  being the uniform distribution over the finite set  $[n]$ :  $\xi = i$  with probability  $\frac{1}{n}$ . Problems with this structure commonly emerge in practice during the training of supervised machine learning models via empirical risk minimization. They are known as finite-sum optimization problems. We rely on two assumptions, presented next.

**Assumption 3.** Function  $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$  is differentiable for all  $i \in [n]$ .

This implies that  $f$  is differentiable.

**Assumption 4.** Function  $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $\mu_i$ -strongly convex for all  $i \in [n]$ , where  $\mu_i > 0$ . That is,

$$f_i(y) + \langle \nabla f_i(y), x - y \rangle + \frac{\mu_i}{2} \|x - y\|^2 \leq f_i(x), \quad x, y \in \mathbb{R}^d. \quad (4)$$

This refines Assumption 2, with the the same strong convexity parameter for all functions. This implies that  $f$  is  $\mu$ -strongly convex with  $\mu = \min_i \mu_i$ . Hence,  $f$  has a unique minimizer  $x_*$ .

## 2 VARIATIONS OF SPPM

The most common algorithm for finding an  $\varepsilon$ -accurate solution of problems (1) and (3) is stochastic gradient descent (SGD) Robbins and Monro (1951); Nemirovski et al. (2009); Gower et al. (2019). In general, SGD updates have the form of  $x^{k+1} = x^k - \gamma g^k$ , where  $g^k$  is an unbiased stochastic gradient estimator:  $\mathbb{E}[g^k | x^k] = \nabla f(x^k)$ . There are numerous approaches to constructing an estimator with the aim of attaining preferable algorithmic features such as faster convergence rate, lower iteration cost, parallelization, and generalization. One of the most significant issues in SGD variations is the difficulties with a suitable selection of the stepsize. Theoretical results highlight that the stepsize is restricted to small values (Bach and Moulines, 2011).

Whenever we are able to evaluate the stochastic proximity operator, another option to consider in place of SGD is the stochastic proximal point method (SPPM), whose iterations have the form

$$\text{prox}_{\gamma f_{\xi_k}}(x_k) := \arg \min_{x \in \mathbb{R}^d} \left\{ f_{\xi_k}(x) + \frac{1}{2\gamma} \|x - x_k\|^2 \right\}, \quad \xi_k \sim \mathcal{D},$$

where  $\gamma > 0$  is a stepsize. Clearly, the proximity operator is well-defined due to the strong convexity Assumption 2 on  $f_{\xi}(x)$ . Given the error tolerance  $\varepsilon > 0$ , we derive in Appendix B.2 (see (23); also, see (Asi and Duchi, 2019, Proposition 5.3), (Khaled and Jin, 2023, Theorem 1)) that with  $\sigma_*^2 := \mathbb{E}_{\xi \sim \mathcal{D}} [\|\nabla f_{\xi}(x_*)\|^2]$ , the stepsize  $\gamma = \frac{\mu \varepsilon}{\sigma_*^2}$ , we get  $\mathbb{E}[\|x_k - x_*\|^2] \leq \varepsilon$  provided that

$$k \geq \left( \frac{1}{2} + \frac{\sigma_*^2}{2\mu^2\varepsilon} \right) \log \left( \frac{2\|x_0 - x_*\|^2}{\varepsilon} \right).$$

As shown by Gower et al. (2019), in the same setting, SGD with fixed stepsize reaches an  $\varepsilon$ -accurate solution after

$$k \geq \left( \frac{2L}{\mu} + \frac{2\sigma_*^2}{\mu^2\varepsilon} \right) \log \left( \frac{4\|x_0 - x_*\|^2}{\varepsilon} \right)$$

iterations ( $L$  is a bound on the smoothness constant of stochastic functions). Note that although both iteration complexities are dependent on the stochastic noise term, the iteration complexity of SGD additionally hinges on the condition number.  $\kappa := \frac{L}{\mu} \geq 1$  ( $\mu$  is a strong convexity parameter of  $f$  here). In contrast, the iteration complexity of SPPM remains unaffected by the smoothness constant  $L$ . Consequently, if we have access to stochastic proximal operator evaluations, we can achieve a faster convergence rate than SGD. Another important aspect is that SPPM still works for large stepsizes. *The primary distinction* with the result of Khaled and Jin (2023) is that the neighborhood guaranteed in our analysis for SPPM does not blow up to infinity as the stepsize  $\gamma$  grows to infinity (see (21) in Theorem 2 and Commentary 2 after it).

Ryu and Boyd (2016); Asi and Duchi (2019) demonstrated that SPPM exhibits greater resilience in terms of the choice of stepsize compared to SGD. Ryu and Boyd (2016) furnish convergence rates for SPPM and note its stability against learning rate misspecification, a trait not shared by SGD. Asi and Duchi (2019) examine a broader method (AProx) which encompasses SPPM as a particular instance, providing both stability and convergence rates under convexity. Additionally, the convergence rates of SPPM remain consistent with those of SGD across different versions of the algorithms.

We have discussed the versatility in designing the unbiased stochastic estimator  $g^k$  in SGD, which can be accomplished in various manners. Among these, popular sampling strategies include importance sampling and mini-batching. A comprehensive analysis of these methods within the framework of arbitrary sampling was presented by Gower et al. (2019). Similarly, an analogous effort is

made for SPPM in our paper, where sampling schemes and methods such as SPPM-NS and SPPM-AS are proposed (see Appendix B.3 and B.4). While sampling strategies offer significant utility, both for SGD and SPPM variations, they tend to converge to the vicinity of the solution for fixed stepsizes, reaching the exact solution only in the overparameterized regime.

That being said, the issue of SGD iterates not converging to the optimum prompted the development of variance-reduced methods. These methods enhance the convergence rate of SGD for finite-sum problems by constructing gradient estimators with diminishing variance over time (e.g., SAGA (Defazio et al., 2014) and SVRG (Johnson and Zhang, 2013)). In other words, variance-reduced methods progressively acquire knowledge of the stochastic gradients at the optimum and mitigate the influence of gradient noise. Consequently, for strongly convex  $f$ , these variance-reduced methods exhibit linear convergence towards  $x_*$  with a fixed step size (see a survey by Gower et al. (2020)). In the discussion above we concluded that SPPM has advantages over SGD. This motivated Khaled and Jin (2023) to explore variance-reduced variants of SPPM. They proposed the SVRP algorithm with an approximate calculation of the proximal operator and demonstrated that variance reduced variants of SPPM have better convergence guarantees under second-order similarity assumption for the finite-sum setting (3). We present and employ in this paper a more general form of this assumption for the stochastic setting (1) in order to analyze L-SVRP with the proximal operator in the updates. We stipulate the similarity assumption to be valid exclusively at  $x_*$ , whereas Khaled and Jin (2023) require it to hold for any  $y \in \mathbb{R}^d$  instead.

**Assumption 5** (Similarity). *There exists  $\delta \geq 0$  such that*

$$\mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|\nabla f_{\xi}(x) - \nabla f(x) - \nabla f_{\xi}(x_*)\|^2 \right] \leq \delta^2 \|x - x_*\|^2, \quad \forall x \in \mathbb{R}^d. \quad (5)$$

Let us also note here that the standard  $\delta$ -smoothness assumption implies the second-order similarity assumption by Khaled and Jin (2023), which in turn implies Assumption 5. In Appendix B.6 we discuss the generality of Assumption 5 in detail. In the studies of Szlendak et al. (2022); Panferov et al. (2024), federated optimization with compression under the similarity assumption is explored, leading to improved convergence rates achieved through specially designed compression operators.

Point SAGA, also a variant of SPPM, was proposed by Defazio (2016), where its convergence was analyzed under the individual smoothness assumption. The algorithm requires a large amount of memory for its execution. It inherits this shortcoming from its SGD variance-reduced archetype SAGA by Defazio et al. (2014). The structure of Point SAGA will not allow us to perform the analysis under the similarity assumption (Assumption 5). Instead, we will rely on this stronger assumption.

**Assumption 6.** *We assume that there exists  $\nu > 0$  such that the inequality*

$$\frac{1}{n} \sum_{j=1}^n \left\| \nabla f_j(x^j) - \frac{1}{n} \sum_{i=1}^n \nabla f_i(x^i) - \nabla f_j(x_*) \right\|^2 \leq \nu^2 \frac{1}{n} \sum_{j=1}^n \|x^j - x_*\|^2 \quad (6)$$

*holds for all  $x^1, \dots, x^n \in \mathbb{R}^d$ .*

Let us note that this condition is weaker than the individual  $\nu$ -smoothness assumption. In Appendix B.8 we discuss the generality of Assumption 6 in more detail.

Traoré et al. (2023) analyze L-SVRP and Point SAGA algorithms in the setting (3) when each  $f_i$  is smooth, convex, and either  $f$  is convex or satisfies the PL-condition. In contrast, in our work we consider a different setting with individually strongly convex functions under Assumption 5 or Assumption 6 (both are weaker than the individual smoothness assumption). The convergence theory of variance-reduced SPPM methods significantly differs from that of standard SPPM. We suggest the possibility of a unified theory that encompasses both SPPM and its variance-reduced counterparts.

### 3 CONTRIBUTIONS

Numerous efficient adaptations of the SPPM algorithm have emerged, each with its specific applications. Our research stems from the absence of a comprehensive and universally applicable theory. While some connections among existing methods have been established (as demonstrated

by Traoré et al. (2023), who link the classical SPPM method with its variance-reduced counterparts), a cohesive theoretical framework without the smoothness assumption is missing. Understanding the connections among various algorithms rooted in SPPM is becoming increasingly challenging for the community, both in theory and practical applications. While new variants are yet to be discovered, determining concrete principles beyond intuition to guide their discovery remains challenging. Complicating matters further is the use of various assumptions regarding the correction vectors across different literature, each with differing levels of rigor. Key contributions of this study comprise:

- **A universal algorithm.** We design a universal SPPM-LC algorithm (stochastic proximal point method with learned correction, Algorithm 1) that encompasses 7 variants of stochastic proximal point method (SPPM, SPPM-NS, SPPM-AS, SPPM\*, SPPM-GC, L-SVRP and Point-SAGA; see Table 1) through (*random*) *correction vectors*  $h_k$ . A specific choice of correction vectors allows to retrieve any particular of the mentioned algorithms. Motivated by the work of Gorbunov et al. (2020), a similar universal algorithm was proposed by Traoré et al. (2023), but under different assumptions.

- **Comprehensive analysis.** We introduce a cohesive theoretical framework of assumptions that includes three restrictions imposed on correction vectors  $h_k$  of the SPPM-LC algorithm (Assumption 7). These restrictions connect the correction vectors  $h_k$  to the iterates of the algorithm and have the forms of parametric recursions. The development of this framework constitutes one of our major contributions. Under Assumption 1, Assumption 2 on the functions and Assumption 7 on the correction vectors, we analyze the convergence of SPPM-LC in the case when it is applied to find an  $\varepsilon$ -accurate solution of the general stochastic optimization problem (1). This is the first comprehensive analysis of the original, sampling-based, and variance-reduced versions of SPPM. It implies the convergence results for all 7 variations of stochastic proximal point method (SPPM, SPPM-NS, SPPM-AS, SPPM\*, SPPM-GC, L-SVRP and Point-SAGA). The sampling-based methods SPPM-NS and SPPM-AS are employed for a less general setting (3), as well as the variance-reduced method Point SAGA, they are analyzed under Assumption 3 and Assumption 4 on the functions. In addition, the analysis of L-SVRP is based on the similarity Assumption 5, and the analysis of Point-SAGA is based on Assumption 6. In order to obtain the convergence result for any particular of the mentioned algorithms, one needs to check that (the recursive parametric) Assumption 7 holds for it. Traoré et al. (2023) introduce a slightly different theoretical framework of assumptions and analyze SPPM and its variance-reduced versions in a finite-sum setting (3) in another setup when functions are individually convex and  $L$ -smooth.

- **Best known rates retrieved.** The rates derived from our comprehensive Theorem 1, under Assumption 1, Assumption 2 on the functions for SPPM, SPPM\*, SPPM-GC, L-SVRP (under additional Assumption 5) and Point-SAGA (under additional Assumption 6); under Assumption 3 and Assumption 4 on the functions for SPPM-NS, SPPM-AS, represent the sharpest rates for these methods. Notably, the neighborhood that we guarantee for SPPM in our analysis does not blow up for the stepsize  $\gamma \rightarrow \infty$  in comparison to the result of Khaled and Jin (2023) (see (21) in Theorem 2 and Commentary 2 after it). Also, we present the analysis for L-SVRP in a simplified way and obtain slightly better bounds on the iteration complexity up to a constant factor than Khaled and Jin (2023).

- **Analysis under general similarity assumptions.** The analysis of L-SVRP is based on the similarity Assumption 5, and the analysis of Point-SAGA is based on Assumption 6. Both of these assumptions are very general and distinguish our approach from the previous ones. Defazio (2016) analyzed Point SAGA under the individual smoothness assumption, which implies Assumption 6. Khaled and Jin (2023) analyzed L-SVRP under a more restrictive similarity assumption than Assumption 5. Traoré et al. (2023) analyzed L-SVRP and Point SAGA under the individual smoothness assumption, which is much more restrictive than our Assumption 5 and Assumption 6. Their results were obtained prior to our work, but we conducted our analysis independently: we already had our result when we discovered their paper.

- **Analysis in a general setting.** We analyze L-SVRP in the stochastic optimization setting (1) while previous works of Khaled and Jin (2023); Traoré et al. (2023) do this in the less general setting (3).

- **New methods.** Our comprehensive theory offers complexity bounds for a range of novel (SPPM-LC, SPPM-NS, SPPM-AS, SPPM-GC, SPPM\*) and upcoming variations of SPPM. It suffices to confirm that Assumption 7 holds, and a complexity estimate is readily provided by Theorem 1. Selected existing and new methods that align with our framework are outlined in Table 1.

216 • **Experiments.** Through extensive experimentation, we demonstrate that several of the newly introduced methods, analyzed within our framework, exhibit compelling empirical properties when  
 217 compared to natural baselines.  
 218  
 219

## 220 4 MAIN RESULT

221  
 222 We are now ready to present our general Algorithm 1, which we call Stochastic Proximal Point  
 223 Method with Learned Correction (SPPM-LC). Subsequently, we introduce the core assumption on  
 224 the correction vectors, iterates and control vectors of Algorithm 1 enabling our general analysis, and  
 225 further state and comment on our unified convergence result (Theorem 1).  
 226

---

### 227 **Algorithm 1** Stochastic Proximal Point Method with Learned Correction (SPPM-LC)

---

- 228 1: **Parameters:** learning rate  $\gamma > 0$ , starting point  $x_0 \in \mathbb{R}^d$ , **control vector**  $\phi_0 \in \mathbb{R}^m$
  - 229 2: **for**  $k = 0, 1, 2, \dots$  **do**
  - 230 3:   Sample  $\xi_k \sim \mathcal{D}$
  - 231 4:   Form correction vector  $h_k$  as a function of the iterate  $x_k$ , **control vector**  $\phi_k$ , and sample  $\xi_k$
  - 232 5:    $x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k + \gamma h_k)$
  - 233 6:   **Construct a new control vector**  $\phi_{k+1}$
  - 234 7: **end for**
- 

#### 236 4.1 KEY ASSUMPTION

237  
 238 We assume that the (random) correction vector  $h_k$  has zero mean (conditioned on  $x_k$  and  $\phi_k$ , the  
 239  $k$ -th iterate and control vector, respectively), and that it is connected with the iterates of SPPM-LC  
 240 via two parametric recursions/inequalities, described next. We introduce versatility by expressing  
 241 these inequalities parametrically.  
 242

243 **Assumption 7** (Parametric recursions). *Let  $\{x_k, \phi_k\}_{k \geq 0}$  be the random iterates produced by SPPM-*  
 244 *LC. Assume that*

$$245 \mathbb{E}[h_k | x_k, \phi_k] = 0. \quad (7)$$

246 *Further, assume that there exists a function  $\sigma^2 : \mathbb{R}^m \rightarrow \mathbb{R}_+$  and nonnegative constants*  
 247  *$A_1, B_1, C_1, A_2, B_2, C_2$ , with  $B_2 < 1$ , such that*

$$248 \mathbb{E} \left[ \|h_k - \nabla f_{\xi_k}(x_*)\|^2 \mid x_k, \phi_k \right] \leq A_1 \|x_k - x_*\|^2 + B_1 \sigma_k^2 + C_1, \quad (8)$$

$$249 \mathbb{E} [\sigma_{k+1}^2 \mid x_{k+1}, \phi_k] \leq A_2 \|x_{k+1} - x_*\|^2 + B_2 \sigma_k^2 + C_2, \quad (9)$$

250 where  $\sigma_k^2 := \sigma^2(\phi_k)$ .  
 251

252 For brevity, we refer to this assumption as the “ $\sigma_k^2$  assumption”. If Assumption 7 holds, then by  
 253 taking expectation on both sides of (8) and (9) and applying the tower property in each case, we get

$$254 \mathbb{E} \left[ \|h_k - \nabla f_{\xi_k}(x_*)\|^2 \right] \leq A_1 \mathbb{E} \left[ \|x_k - x_*\|^2 \right] + B_1 \mathbb{E} [\sigma_k^2] + C_1, \quad (10)$$

$$255 \mathbb{E} [\sigma_{k+1}^2] \leq A_2 \mathbb{E} \left[ \|x_{k+1} - x_*\|^2 \right] + B_2 \mathbb{E} [\sigma_k^2] + C_2, \quad (11)$$

256 The novelty of our approach lies in the introduction of inequalities (8) and (9). We support and  
 257 validate this assertion by providing numerous examples (in Section 5), demonstrating that these  
 258 inequalities encapsulate the nature of a broad range of existing SPPM methods as well as some new  
 259 ones, including standard SPPM alongside its arbitrary sampling and variance-reduced variants. In  
 260 its essence, we generalize, parameterize and establish as an independent assumption the conditions  
 261 on correction vectors for SPPM-type methods present in the literature, regardless of the specifics  
 262 defining the base method from which they stem. Traoré et al. (2023) propose different parameterized  
 263 recursive inequalities and analyze SPPM and its variance-reduced versions in a finite-sum setting (3)  
 264 under the condition where functions are individually convex and  $L$ -smooth. Similar inequalities  
 265 can be found in the analysis of SGD-type methods (a unified theory developed by Gorbunov et al.  
 266 (2020)).  
 267  
 268  
 269

Table 1: Compilation of both existing and novel methods that align with our comprehensive analytical framework. Problem (1) encompasses a broader scope compared to problem (3). VR = variance reduced method, AS = arbitrary sampling. Thm  $x$  = Thm 1 + Lemma  $x$ , Lemma  $x$  is in Section A. $x$ ,  $x \in [8]$ . The last column indicates whether the analysis is new or not and whether we recover the previously established rate.

Problem	Method	Alg #	VR?	AS?	Section	Result / Rate
(1)	SPPM-LC [NEW]	Alg 1	✓✗	✗	B.1	Thm 1 [NEW]
(1)	SPPM <sup>(1)</sup> (Bertsekas, 2011)	Alg 2	✗	✗	B.2	Thm 2 [NEW]
(3)	SPPM-NS [NEW]	Alg 3	✗	✗	B.3	Thm 3 [NEW]
(3)	SPPM-AS [NEW]	Alg 4	✗	✓	B.4	Thm 4 [NEW]
(1)	SPPM* [NEW]	Alg 5	✓	✓	B.5	Thm 5 [NEW]
(1)	SPPM-GC [NEW]	Alg 6	✓	✗	B.6	Thm 6 [NEW]
(1)	L-SVRP <sup>(2)</sup> (Khaled and Jin, 2023)	Alg 7	✓	✗	B.7	Thm 7 [NEW]
(3)	Point SAGA <sup>(3)</sup> (Defazio, 2016)	Alg 8	✓	✗	B.8	Thm 9 [NEW]

<sup>(1)</sup> SPPM was studied by Khaled and Jin (2023) under the less general similarity assumption than ours. Bertsekas (2011) proposed incremental proximal point method (SPPM for problem (3)) and analyzed it under assumptions that each  $f_i$  is Lipschitz. We guarantee in our analysis that the neighborhood does not blow up when the stepsize is large and consider the general setting (1).

<sup>(2)</sup> The L-SVRP method was proposed by Khaled and Jin (2023) and called SVRP therein. It was inspired by the L-SVRG method of Hofmann et al. (2015); Kovalev et al. (2020), who were in turn inspired by the SVRG method of Johnson and Zhang (2013). Following Khaled and Jin (2023), we use the name L-SVRP to highlight the loopless nature of the update of control vector. The theoretical results presented here are a minor adaptation of the results of Khaled and Jin (2023). We present the analysis in a simplified way, and hence obtain slightly better bounds up to a constant factor. Khaled and Jin (2023) employ an approximation of the proximal operator in the updates of L-SVRP while we use the operator itself.

<sup>(3)</sup> Point SAGA was proposed by Defazio (2016). The main difference between our form and the original one is in the control vectors. Defazio (2016) updates a table with gradients, while we update a table with points at which we compute the gradients.

## 4.2 MAIN THEOREM

We are now prepared to introduce our main convergence result.

**Theorem 1.** *Let Assumption 1 (differentiability) and Assumption 2 ( $\mu$ -strong convexity) hold. Let  $\{x_k, h_k\}$  be the iterates produced by SPPM-LC (Algorithm 1), and assume that they satisfy Assumption 7 ( $\sigma_k^2$ -assumption). Choose any  $\gamma > 0$  and  $\alpha > 0$  satisfying the inequalities*

$$\frac{(1 + \gamma^2 A_1)(1 + \alpha A_2)}{(1 + \gamma\mu)^2} < 1, \quad \frac{\gamma^2 B_1(1 + \alpha A_2)}{\alpha(1 + \gamma\mu)^2} + B_2 < 1, \quad (12)$$

and define the Lyapunov function

$$\Psi_k := \|x_k - x_*\|^2 + \alpha\sigma_k^2. \quad (13)$$

Then for all iterates  $k \geq 0$  of SPPM-LC we have

$$\mathbb{E}[\Psi_k] \leq \theta^k \Psi_0 + \frac{\zeta}{1 - \theta}, \quad (14)$$

where the parameters  $0 \leq \theta < 1$  and  $\zeta \geq 0$  are given by

$$\theta = \max \left\{ \frac{(1 + \gamma^2 A_1)(1 + \alpha A_2)}{(1 + \gamma\mu)^2}, \frac{\gamma^2 B_1(1 + \alpha A_2)}{\alpha(1 + \gamma\mu)^2} + B_2 \right\}, \quad (15)$$

$$\zeta = \frac{\gamma^2 C_1(1 + \alpha A_2)}{(1 + \gamma\mu)^2} + \alpha C_2. \quad (16)$$

Theorem 1 proves a linear rate of convergence for a number of stochastic proximal point methods towards a fluctuation neighborhood around the solution, regulated by the additive term in (14). It depends on parameters  $C_1$  and  $C_2$ . The neighborhood vanishes (i.e.,  $\zeta = 0$ ) iff  $C_1 = C_2 = 0$ . If this happens, then SPPM-LC converges linearly to the solution as  $k \rightarrow \infty$  for any fixed  $\gamma > 0$ , satisfying the conditions of Theorem 1. Notice that  $\theta \geq B_2$ , and hence the linear rate can not be faster than  $(B_2)^k$ . That is, as shown in Appendix B (also, see Table 2), the main difference between variance-reduced versions of SPPM and its other variants is that the former methods satisfy  $\sigma_k^2$ -assumption with  $C_1 = C_2 = 0$  (and reach the optimum  $x_*$ ), whilst the latter have either  $C_1 > 0$  or  $C_2 > 0$ .

## 5 OVERVIEW OF SPECIFIC METHODS AND OF THE FRAMEWORK

In this section we demonstrate the descriptive power of our new framework. We show that most popular existing methods can be expressed in terms of our framework and, in addition to that, we describe several new methods in more detail (see Table 1).

### 5.1 A BRIEF OVERVIEW

As asserted, the proposed framework is powerful enough to include methods without variance reduction (✗ in the “VR” column) alongside variance-reduced methods (✓ in the “VR” column), methods that fall under the arbitrary sampling paradigm (✓ in the “AS” column). All novel methods introduced in this paper are clearly designated with the label [NEW]. Additionally, to facilitate a thorough understanding of all algorithms discussed, detailed explanations are included in the Appendix. A link is provided for convenient navigation to the supplementary section. The generality of our framework is reflected in Table 1. The “Result / Rate” column of Table 1 refers to a Theorem  $x$  which follows from Theorem 1 and Lemma  $x$ ,  $x \in [8]$ . The convergence results of the algorithms considered in the paper are outlined in these theorems, offering insights into their performance characteristics. Importantly, in instances where established methods are recovered, the best-known convergence rates are reaffirmed or better results are obtained. This underscores the robustness and reliability of our analytical framework in accurately capturing the behavior of established algorithms.

### 5.2 PARAMETERS OF THE FRAMEWORK

The algorithms in Table 1 demonstrate specific patterns in relation to the parameters in Assumption 7. To elucidate this observation, we provide a summary of these parameters in Table 2. As predicted by Theorem 1, when  $C_1 = C_2 = 0$ , the corresponding method does not oscillate and converges to the optimum  $x_*$ , indicating the variance-reduced nature of the algorithm. All parameters referenced in Table 2 are defined in the Appendix, alongside descriptions and analyses of the respective methods.

### 5.3 FIVE NOVEL ALGORITHMS

To showcase the efficiency of our comprehensive framework, we develop three new variants of SPPM which have not previously been addressed in the literature (see Table 1). In this section, we briefly outline the reasoning behind their implementation. Further specifics are available in the Appendix.

SPPM-NS (Algorithm 3). The method is designed for solving the problem (3). Let positive numbers  $p_1, \dots, p_n$  sum up to 1, set  $i_k = i$  with probability  $p_i$ ,  $i \in [n]$ . The step of the method has the form  $x_{k+1} = \text{prox}_{\frac{\gamma}{np_i} f_{i_k}}(x_k)$ . It unifies several powerful sampling strategies (as, e.g., importance sampling). Sampling allows to improve the convergence rate and modify the neighborhood.

SPPM-AS (Algorithm 4). The method is also designed for the problem (3). The arbitrary sampling framework was developed for SGD by Gower et al. (2019). It allows to conduct a sharp unified convergence analysis for various effective sampling and mini-batch strategies. For strongly convex functions, the method with constant stepsize converges linearly to the neighborhood of the solution.

SPPM\* (Algorithm 5). This novel algorithm links conventional and variance-reduced SPPM methods. Although not immediately practical, it offers valuable insights into the inner workings of variance reduction. This method addresses the fundamental question: assuming that the gradients  $\nabla f_\xi(x_*)$  are available, can they be leveraged to devise a more potent variant of SPPM? The affirmative answer culminates in the development of SPPM\*. The construction of updates in SPPM\* involves correction vectors of the form  $h_k = \nabla f_{\xi_k}(x_*)$ . In essence, this implies augmenting  $x_k$  with gradients of the same functions at the optimal point  $x_*$ , with respect to which the proximal operator is calculated. As evidenced in Table 2, where  $C_1 = C_2 = 0$ , this method converges directly to  $x_*$  without oscillation, rather than converging to a neighborhood of the solution. Practical variance-reduced methods operate by iteratively refining estimates of  $\nabla f_{\xi_k}(x_*)$ . Notably, the term  $\sigma_k^2$  in the Lyapunov function of variance-reduced methods converges to zero.

SPPM-GC (Algorithm 6). This method can be viewed as a practical variant of SPPM\* featuring a computable version of the correction vector,  $h_k = \nabla f_{\xi_k}(x_k) - \nabla f(x_k)$ , instead of the incomputable correction  $\nabla f_{\xi_k}(x_*)$ . This method follows the same paradigm of iteratively constructing increasingly more refined estimates of  $\nabla f_{\xi_k}(x_*)$  and is also variance reduced. As indicated in Table 2, where  $C_1 = C_2 = 0$ , this method also converges to  $x_*$ , and not to some neighborhood of the solution only.

SPPM-LC (Algorithm 1). This new generic algorithm constructs the updates with correction vectors  $h_k$  of a universal form. We analyze its behavior under a new parametric Assumption 7. The convergence result then follows. This is a unified analysis of all stochastic proximal point methods we have encountered so far: SPPM, SPPM-NS, SPPM-AS, SPPM\*, SPPM-GC, L-SVRP and Point-SAGA. One only needs to check that this parametric assumption holds in each particular case (see Table 2).

Method	$A_1$	$B_1$	$C_1$	$A_2$	$B_2$	$C_2$	Lemma
SPPM-LC	$A_1$	$B_1$	$C_1$	$A_2$	$B_2$	$C_2$	Lemma 1
SPPM	0	0	$\sigma_*^2$	0	0	0	Lemma 2
SPPM-NS	0	0	$\sigma_{*,NS}^2$	0	0	0	Lemma 3
SPPM-AS	0	0	$\sigma_{*,AS}^2$	0	0	0	Lemma 4
SPPM*	0	0	0	0	0	0	Lemma 5
SPPM-GC	$\delta^2$	0	0	0	0	0	Lemma 6
L-SVRP	0	$\delta^2$	0	$p$	$1-p$	0	Lemma 7
Point SAGA	0	$\nu^2$	0	$\frac{1}{n}$	$\frac{n-1}{n}$	0	Lemma 8

Table 2: The parameters determining the compliance of the methods listed in Table 1 to Assumption 7. Detailed explanations of the expressions featured in the table are provided in the Appendix.

## 6 EXPERIMENTS

In this section we describe numerical experiments conducted for the linear regression problem

$$\min_{x \in \mathbb{R}^d} \left\{ \frac{1}{2n} \sum_{i=1}^n (a_i^\top x - b_i)^2 + \lambda_i \|x\|^2 \right\}, \quad (17)$$

where  $a_i \in \mathbb{R}^d$ ,  $b_i \in \mathbb{R}$  is the  $i$ -th data pair, each  $\lambda_i$  is a  $\ell_2$ -regularization parameter. We provide 4 sets of experiments. The first one is devoted to the comparison of different sampling strategies for SPPM-NS : uniform sampling US, importance sampling IS, variance sampling VS (see Section B.3), with three selected stepsizes. The second set of experiments demonstrates the change of the radius of neighborhood for SPPM-AS with  $\tau$ -nice sampling. In the next two sets of experiments we illustrate the main difference between SPPM and SPPM with variance reduction. Also in practice we show the relationship between SPPM-GC, L-SVRP, Poin-SAGA as we did in theory. For the first bunch of experiments we set  $n = 10$ ,  $d = 3$  and the regularization parameters  $\lambda_i = 1/2^i$ , where  $i \in [d]$ . Looking at Figure 1, we can see the different behavior of the methods. In the first two plots with the smallest stepsizes, SPPM-IS has a faster start, but a larger neighborhood than the other considered methods. Unfortunately, we cannot say that SPPM-VS has a much smaller neighborhood radius than SPPM-US or SPPM-IS, but theoretically, the variance is smaller. In the second set of numerical experiments (see Figure 2), we observe a clear correlation between the neighborhood radius and the cardinality of the sampled subset  $\tau$ . More precisely, the larger  $\tau$  is, the smaller the neighborhood radius is. In the third set of numerical experiments (see Figure 3), we set  $n = 1000$ ,  $d = 10$  and each  $\lambda_i = 1$  for the problem (17). On all three plots we observe the superiority of SPPM-star in terms of the stepsize choice. For example, in the third plot with  $\gamma = 10^2$ , SPPM-GC diverges. In the fourth set of numerical experiments (see Figure 4) with parameters  $n = 1000$ ,  $d = 10$  and each  $\lambda_i = 1$  for the problem (17), we observe that the performances of SPPM-GC and L-SVRP with  $p = 1$  are identical, which is supported by our theoretical findings. With decreasing  $p$  from 1 to  $1/n$  we see how the behavior of L-SVRP worsens and matches with the performance of Point-SAGA when  $p = 1/n$ .






Figure 1: Comparison of the performance of SPPM-US, SPPM-IS, SPPM-VS and SPPM-AS with  $\tau = 9$ -nice sampling for different selections of stepsize  $\gamma \in \{10^{-4}, 10^{-2}, 1, 10^2\}$ .




Figure 2: Comparison of the performance of SPPM-AS with  $\tau$ -nice sampling with different selections of cardinality  $\tau \in \{1, 2, 5, 9, n = 10\}$  and stepsize  $\gamma \in \{10^{-2}, 10^{-1}, 1\}$ .




Figure 3: Comparison of the performance of SPPM-US, SPPM-GC and SPPM-star with different selections of stepsize  $\gamma \in \{10^{-2}, 1, 10^2\}$ .




Figure 4: Comparison of the performance of SPPM-GC, Point-SAGA, L-SVRP with different selection of probabilities  $\gamma \in \{1/n = 10^{-3}, 5 \cdot 10^{-3}, 10^2, 5 \cdot 10^2, 10^{-1}, 1\}$ . The stepsizes are taken according to the theory.

## 7 FURTHER DISCUSSION

Although our approach is general, we still see some limitations, open problems and several possible directions for future extensions. Generating a similar result in the nonconvex case continues to be an unsolved challenge. Expanding Assumption 7 to incorporate iteration-dependent parameters  $A_1, B_1, C_1, A_2, B_2, C_2$  could facilitate the development of various novel methods as SPPM with

486 decreasing stepsizes. Khaled and Jin (2023) explore federated learning, investigating client sam-  
 487 pling to enhance communication efficiency. Another way to do that is to incorporate the utilization  
 488 of compressed vectors exchanged between a server and the clients. This motivates the problem of  
 489 providing the analysis for such SPPM-type methods and incorporating them into our framework.  
 490 It would be interesting to build theory for algorithms with correction vectors  $h_k$  with a non-zero  
 491 expected value and unify with our theory. Another potential avenue for future research involves of-  
 492 fering a comprehensive analysis of SPPM-type methods incorporating acceleration and momentum.

493  
 494

## 495 REFERENCES

- 496 Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. QSGD:  
 497 Communication-efficient SGD via gradient quantization and encoding. In *Advances in Neural*  
 498 *Information Processing Systems*, pages 1709–1720, 2017.
- 499  
 500 Dan Alistarh, Torsten Hoefer, Mikael Johansson, Sarit Khirirat, Nikola Konstantinov, and Cédric  
 501 Renggli. The convergence of sparsified gradient methods. In *Advances in Neural Information*  
 502 *Processing Systems*, 2018.
- 503  
 504 Hilal Asi and John C. Duchi. Stochastic (approximate) proximal point methods: Convergence,  
 505 optimality, and adaptivity. *SIAM Journal on Optimization*, 29(3):2257–2290, 2019. doi: 10.1137/  
 506 18M1230323.
- 507  
 508 Francis R. Bach and Éric Moulines. Non-asymptotic analysis of stochastic approximation algorithms  
 509 for machine learning. In *Neural Information Processing Systems*, 2011. URL [https://api.  
 510 semanticscholar.org/CorpusID:3806935](https://api.semanticscholar.org/CorpusID:3806935).
- 511  
 512 Dimitri P Bertsekas. Incremental proximal methods for large scale convex optimization. *Mathemat-*  
*ical programming*, 129(2):163–195, 2011.
- 513  
 514 Aleksandr Beznosikov and Alexander V. Gasnikov. Compression and data similarity: Combina-  
 515 tion of two techniques for communication-efficient solving of distributed variational inequal-  
 516 ities. In *OPTIMA*, 2022. URL [https://api.semanticscholar.org/CorpusID:  
 517 249889654](https://api.semanticscholar.org/CorpusID:249889654).
- 518  
 519 Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *International Con-*  
*ference on Computational Statistics*, 2010. URL [https://api.semanticscholar.org/  
 520 CorpusID:115963355](https://api.semanticscholar.org/CorpusID:115963355).
- 521  
 522 Gong Chen and Marc Teboulle. Convergence analysis of a proximal-like minimization algorithm  
 523 using bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, 1993. doi: 10.1137/  
 524 0803026.
- 525  
 526 Aaron Defazio. A simple practical accelerated method for finite sums. In *Advances in Neural*  
*Information Processing Systems 29*, 2016.
- 527  
 528 Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. SAGA: A fast incremental gradient  
 529 method with support for non-strongly convex composite objectives. In *Advances in Neural Infor-*  
 530 *mation Processing Systems 27*, 2014.
- 531  
 532 Yury Demidovich, Grigory Malinovsky, Igor Sokolov, and Peter Richtárik. A guide through the zoo  
 533 of biased SGD. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.  
 534 URL <https://openreview.net/forum?id=OCtv4NyahI>.
- 535  
 536 Eduard Gorbunov, Filip Hanzely, and Peter Richtárik. A unified theory of sgd: Variance reduction,  
 537 sampling, quantization and coordinate descent. In *The 23rd International Conference on Artificial*  
*Intelligence and Statistics*, 2020.
- 538  
 539 Robert M. Gower, Mark Schmidt, Francis Bach, and Peter Richtárik. Variance-reduced methods for  
 machine learning. *Proceedings of the IEEE*, 108(11):1968–1983, 2020.

- 540 Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter  
541 Richtárik. SGD: General analysis and improved rates. In Kamalika Chaudhuri and Ruslan  
542 Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*,  
543 volume 97 of *Proceedings of Machine Learning Research*, pages 5200–5209, Long Beach, Cal-  
544 ifornia, USA, 09–15 Jun 2019. PMLR. URL [http://proceedings.mlr.press/v97/  
545 qian19b.html](http://proceedings.mlr.press/v97/qian19b.html).
- 546 Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with  
547 limited numerical precision. In Francis Bach and David Blei, editors, *Proceedings of the 32nd  
548 International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning  
549 Research*, pages 1737–1746, Lille, France, 07–09 Jul 2015. PMLR.
- 550 Thomas Hofmann, Aurelien Lucchi, Simon Lacoste-Julien, and Brian McWilliams.  
551 Variance reduced stochastic gradient descent with neighbors. In C. Cortes,  
552 N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural  
553 Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL  
554 [https://proceedings.neurips.cc/paper\\_files/paper/2015/file/  
555 effc299aladb07e7089f9b269c31f2f-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2015/file/effc299aladb07e7089f9b269c31f2f-Paper.pdf).
- 556 Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance  
557 reduction. In *Advances in Neural Information Processing Systems 26*, pages 315–323, 2013.
- 558 Sai Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Suresh.  
559 SCAFFOLD: Stochastic controlled averaging for on-device federated learning. In *ICML, 2020*.
- 560 Ahmed Khaled and Chi Jin. Faster federated optimization under second-order similarity. In *ICLR,*  
561 *2023*.
- 562 Ahmed Khaled and Peter Richtárik. Better theory for SGD in the nonconvex world. *Transactions*  
563 *on Machine Learning Research*, 2023. ISSN 2835-8856. URL [https://openreview.net/  
564 forum?id=AU4qHN2Vks](https://openreview.net/forum?id=AU4qHN2Vks). Survey Certification.
- 565 Junhyung Lyle Kim, Panos Toulis, and Anastasios Kyrillidis. Convergence and stability of the  
566 stochastic proximal point algorithm with momentum. In Roya Firoozi, Negar Mehr, Esen Yel,  
567 Rika Antonova, Jeannette Bohg, Mac Schwager, and Mykel Kochenderfer, editors, *Proceedings  
568 of The 4th Annual Learning for Dynamics and Control Conference*, volume 168 of *Proceedings  
569 of Machine Learning Research*, pages 1034–1047. PMLR, 23–24 Jun 2022.
- 570 Jakub Konečný and Peter Richtárik. Randomized distributed mean estimation: accuracy vs commu-  
571 nication. *Frontiers in Applied Mathematics and Statistics*, 4(62):1–11, 2018.
- 572 Jakub Konečný, Jie Lu, Peter Richtárik, and Martin Takáč. Mini-batch semi-stochastic gradient  
573 descent in the proximal setting. *IEEE Journal of Selected Topics in Signal Processing*, 10(2):  
574 242–255, 2016.
- 575 Dmitry Kovalev, Samuel Horváth, and Peter Richtárik. Don’t jump through hoops and remove  
576 those loops: SVRG and Katyusha are better without the outer loop. In *Proceedings of the 31st  
577 International Conference on Algorithmic Learning Theory*, 2020.
- 578 Konstantin Mishchenko, Ahmed Khaled, and Peter Richtárik. Proximal and federated random  
579 reshuffling. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu,  
580 and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learn-  
581 ing*, volume 162 of *Proceedings of Machine Learning Research*, pages 15718–15749. PMLR,  
582 17–23 Jul 2022.
- 583 Deanna Needell, Nathan Srebro, and Rachel Ward. Stochastic gradient descent, weighted sampling,  
584 and the randomized Kaczmarz algorithm. *Mathematical Programming*, 155(1–2):549–573, 2015.
- 585 Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic  
586 approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–  
587 1609, 2009.
- 588 Yurii Nesterov. *Lectures on Convex Optimization*. Springer Publishing Company, Incorporated, 2nd  
589 edition, 2018. ISBN 3319915770.

- 594 Lam Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. SARAH: A novel method for machine  
595 learning problems using stochastic recursive gradient. In *The 34th International Conference on*  
596 *Machine Learning*, 2017.
- 597 Andrei Panferov, Yury Demidovich, Ahmad Rammal, and Peter Richtárik. Correlated quantization  
598 for faster nonconvex distributed optimization. *ArXiv*, abs/2401.05518, 2024.
- 600 Andrei Pătrașcu and Ion Necoara. Nonasymptotic convergence of stochastic proximal point meth-  
601 ods for constrained convex optimization. *J. Mach. Learn. Res.*, 18:198:1–198:42, 2017. URL  
602 <https://api.semanticscholar.org/CorpusID:49679465>.
- 603 Herbert Robbins and Sutton Monro. A stochastic approximation method. *Annals of Mathematical*  
604 *Statistics*, 22:400–407, 1951.
- 606 Nicolas Le Roux, Mark Schmidt, and Francis Bach. A stochastic gradient method with an expo-  
607 nential convergence rate for finite training sets. In *Advances in Neural Information Processing*  
608 *Systems*, pages 2663–2671, 2012.
- 609 Ernest Ryu and Stephen Boyd. Stochastic proximal iteration: A non-asymptotic improvement upon  
610 stochastic gradient descent. Technical report, Stanford University, 2016.
- 612 Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 1-bit stochastic gradient descent and its  
613 application to data-parallel distributed training of speech dnns. In *Fifteenth Annual Conference of*  
614 *the International Speech Communication Association*, 2014.
- 615 Ohad Shamir, Nati Srebro, and Tong Zhang. Communication-efficient distributed optimization using  
616 an approximate newton-type method. In Eric P. Xing and Tony Jebara, editors, *Proceedings of*  
617 *the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine*  
618 *Learning Research*, pages 1000–1008, Beijing, China, 22–24 Jun 2014. PMLR.
- 620 Rafał Szlendak, Alexander Tyurin, and Peter Richtárik. Permutation compressors for provably faster  
621 distributed nonconvex optimization. In *10th International Conference on Learning Representa-*  
622 *tions*, 2022.
- 623 Cheik Traoré, Vassilis Apidopoulos, Saverio Salzo, and Silvia Villa. Variance reduction techniques  
624 for stochastic proximal point algorithms. *ArXiv*, abs/2308.09310, 2023. URL <https://api.semanticscholar.org/CorpusID:261031238>.
- 626 Sharan Vaswani, Francis Bach, and Mark Schmidt. Fast and faster convergence of SGD for over-  
627 parameterized models and an accelerated perceptron. In *22nd International Conference on Artifi-*  
628 *cial Intelligence and Statistics*, volume 89 of *PMLR*, pages 1195–1204, 2019.
- 630 Xiao-Tong Yuan and Ping Li. On convergence of distributed approximate newton methods: Global-  
631 ization, sharper bounds and beyond. *Journal of Machine Learning Research*, 21(206):1–51, 2020.  
632 URL <http://jmlr.org/papers/v21/19-764.html>.
- 633 Yuchen Zhang and Lin Xiao. DiSCO: Distributed optimization for self-concordant empirical loss.  
634 In *Proceedings of the 32nd International Conference on Machine Learning*, *PMLR*, volume 37,  
635 pages 362–370, 2015.
- 636 Peilin Zhao and Tong Zhang. Stochastic optimization with importance sampling for regularized  
637 loss minimization. In *Proceedings of the 32nd International Conference on Machine Learning*,  
638 *PMLR*, volume 37, pages 1–9, 2015.
- 640  
641  
642  
643  
644  
645  
646  
647

648	CONTENTS	
649		
650		
651	<b>1 Introduction</b>	<b>1</b>
652		
653	<b>2 Variations of SPPM</b>	<b>2</b>
654		
655		
656	<b>3 Contributions</b>	<b>3</b>
657		
658		
659	<b>4 Main result</b>	<b>5</b>
660	4.1 Key assumption . . . . .	5
661	4.2 Main theorem . . . . .	6
662		
663		
664	<b>5 Overview of specific methods and of the framework</b>	<b>7</b>
665		
666	5.1 A brief overview . . . . .	7
667	5.2 Parameters of the framework . . . . .	7
668	5.3 Five novel algorithms . . . . .	7
669		
670		
671	<b>6 Experiments</b>	<b>8</b>
672		
673		
674	<b>7 Further discussion</b>	<b>9</b>
675		
676		
677	<b>A Extended literature overview</b>	<b>14</b>
678		
679	<b>B Special cases</b>	<b>14</b>
680	B.1 Stochastic proximal point method with learned correction (SPPM-LC) . . . . .	14
681	B.2 Stochastic proximal point method (SPPM) . . . . .	16
682	B.3 Stochastic proximal point method with nonuniform sampling (SPPM-NS) . . . . .	18
683	B.4 Stochastic proximal point method with arbitrary sampling (SPPM-AS) . . . . .	20
684	B.5 Stochastic proximal point method with optimal gradient correction (SPPM*) . . . . .	23
685	B.6 Stochastic proximal point method with gradient correction (SPPM-GC) . . . . .	25
686	B.7 Loopless stochastic variance reduced proximal point method (L-SVRP/ SPPM-LGC) . . . . .	28
687	B.8 Point SAGA (Point SAGA) . . . . .	33
688		
689		
690		
691		
692		
693	<b>C Auxiliary lemma for SPPM-LC</b>	<b>36</b>
694		
695		
696	<b>D Auxiliary lemmas</b>	<b>37</b>
697	D.1 SPPM-AS . . . . .	37
698	D.2 L-SVRP . . . . .	37
699		
700		
701	<b>E Auxiliary facts</b>	<b>38</b>

## A EXTENDED LITERATURE OVERVIEW

Stochastic gradient descent (SGD) Robbins and Monro (1951); Nemirovski et al. (2009); Bottou (2010) is a contemporary common algorithm for solving optimization problems (1) and (3) with updates of the form  $x^{k+1} = x^k - \gamma g^k$ , where  $g^k$  is an unbiased stochastic gradient estimator:  $\mathbb{E}[g^k | x^k] = \nabla f(x^k)$ . Theoretical properties of SGD are nowadays studied well in works by Vaswani et al. (2019); Gower et al. (2019); Gorbunov et al. (2020); Khaled and Richtárik (2023); Demidovich et al. (2023). The versatility in the design of the estimator  $g^k$  has led to a development of highly effective variants of SGD based on importance sampling (see works by Needell et al. (2015); Zhao and Zhang (2015)), mini-batching (see a work of Konečný et al. (2016)). These strategies are unified in the arbitrary sampling paradigm proposed by Gower et al. (2019). The iterates of standard SGD as well as of its sampling variants do not converge to the optimum due to the presence of the stochastic gradient noise. The existence of this issue led to the construction of variance reduction methods in works by Roux et al. (2012); Defazio et al. (2014); Johnson and Zhang (2013); Kovalev et al. (2020); Nguyen et al. (2017). Such methods are able to sequentially learn the stochastic gradients at the optimum which allows them to reach a linear convergence to  $x_*$  when equipped with constant stepsize in case of strongly convex  $f$ . The work by Gorbunov et al. (2020) provided a unified theory for SGD and its variance-reduced, sampling, quantization and coordinate sub-sampling modifications. Problem setting (3) can also be used to describe a Federated Learning setup where  $n$  is the number of workers, and  $f_i$  stands for the loss dependent on the data of worker  $i \in [n]$ . Workers compute individual stochastic gradients and aggregate them on a master node (see a paper by Konečný and Richtárik (2018)). A primary bottleneck in such settings is communication. To address it, many techniques are used as quantization (see works by Gupta et al. (2015); Seide et al. (2014)), sparsification (see a paper by Alistarh et al. (2018)), dithering (see a work by Alistarh et al. (2017)). Many distributed optimization works employ variants of Assumption 5 to analyze methods that solve the problem (3) in the strongly convex case (see works by Shamir et al. (2014); Zhang and Xiao (2015); Yuan and Li (2020); Karimireddy et al. (2020)). Szlendak et al. (2022); Beznosikov and Gasnikov (2022); Panferov et al. (2024) achieve better communication complexity guarantees in a Federated Learning setting under the similarity assumption employing specific quantizers and sparsifiers. Selecting the right stepsize is a critical aspect in the implementation of SGD, see the work of Bach and Moulines (2011). In Section 2 we show that the iteration complexity of SGD depends on condition number  $\kappa = \frac{L}{\mu}$ . If we are able to evaluate stochastic proximal operator, another algorithm to use instead of SGD is stochastic proximal point method (SPPM). Its iteration complexity does not depend on the smoothness constant  $L$ . Bertsekas (2011) calls it the incremental proximal point method and considers it for solving the problem (3), proving a convergence to the neighborhood when each  $f_i$  is Lipschitz. In the study by Ryu and Boyd (2016), convergence rates are presented for SPPM, emphasizing its robustness against learning rate misspecification, a characteristic not exhibited by SGD. Pătrașcu and Necoara (2017) consider SPPM to solve a general stochastic optimization problem (1) in the convex case and provide nonasymptotic convergence guarantees. Asi and Duchi (2019) delve into the exploration of a broader method termed AProx, which includes SPPM as a particular case. Their research entails analysis of convergence rates for convex functions. Variance-reduced versions of SPPM should also have advantages over their SGD counterparts. It motivated Khaled and Jin (2023) to analyze L-SVRP in the finite-sum and federated settings under a stronger version of Assumption 5 and strong convexity. Defazio (2016) analyzes Point SAGA under the individual smoothness and strong convexity assumptions. Traoré et al. (2023) consider variance-reduced methods L-SVRP and Point SAGA when each  $f_i(x)$  is  $L$ -smooth, convex, and either  $f(x)$  is convex or satisfies PL-condition. Kim et al. (2022) consider SPPM with momentum.

## B SPECIAL CASES

### B.1 STOCHASTIC PROXIMAL POINT METHOD WITH LEARNED CORRECTION (SPPM-LC)

This section is devoted to a unified analysis of all stochastic proximal point methods we have encountered so far: SPPM, SPPM-NS, SPPM-AS, SPPM\*, SPPM-GC, L-SVRP and Point-SAGA. The analysis is based on formulating a new parametric assumption (Assumption 7) on the behavior of the method, and the convergence result then follows. All that has to be done is to check that this parametric assumption holds in each particular case of interest.

**Lemma 1.** Assumption 7 holds with constants  $A_1, B_1, C_1$  and  $A_2, B_2, C_2$ .

*Proof of Lemma 1.* Since we assume the iterates produced by SPPM-LC (Algorithm 1) satisfy Assumption 7, the statement of Lemma 1 holds automatically.  $\square$

We can now present the main result of this section.

**Theorem 1.** Let Assumption 1 (differentiability) and Assumption 2 ( $\mu$ -strong convexity) hold. Let  $\{x_k, h_k\}$  be the iterates produced by SPPM-LC (Algorithm 1), and assume that they satisfy Assumption 7 ( $\sigma_k^2$ ). Choose any  $\gamma > 0$  and  $\alpha > 0$  satisfying the inequalities

$$\frac{(1 + \gamma^2 A_1)(1 + \alpha A_2)}{(1 + \gamma\mu)^2} < 1, \quad \frac{\gamma^2 B_1(1 + \alpha A_2)}{\alpha(1 + \gamma\mu)^2} + B_2 < 1, \quad (18)$$

and define the Lyapunov function

$$\Psi_k := \|x_k - x_*\|^2 + \alpha\sigma_k^2. \quad (19)$$

Then for all iterates  $k \geq 0$  of SPPM-LC we have

$$\mathbb{E}[\Psi_k] \leq \theta^k \Psi_0 + \frac{\zeta}{1 - \theta},$$

where the parameters  $0 \leq \theta < 1$  and  $\zeta \geq 0$  are given by

$$\theta = \max \left\{ \frac{(1 + \gamma^2 A_1)(1 + \alpha A_2)}{(1 + \gamma\mu)^2}, \frac{\gamma^2 B_1(1 + \alpha A_2)}{\alpha(1 + \gamma\mu)^2} + B_2 \right\}, \quad \zeta = \frac{\gamma^2 C_1(1 + \alpha A_2)}{(1 + \gamma\mu)^2} + \alpha C_2.$$

*Proof of Theorem 1.* We use Lemma 9 (see Appendix C) and  $\sigma_k^2$ -assumption. In particular, by combining inequality Assumption 7 (see (9)) and Lemma 9 (see (70)), for all  $\gamma > 0$  we get

$$\begin{aligned} \mathbb{E}[\Psi_{k+1}] &\stackrel{(19)}{=} \mathbb{E} \left[ \|x_{k+1} - x_*\|^2 + \alpha\sigma_{k+1}^2 \right] \\ &= \mathbb{E} \left[ \|x_{k+1} - x_*\|^2 \right] + \alpha \mathbb{E} \left[ \sigma_{k+1}^2 \right] \\ &\stackrel{(70)}{\leq} \frac{(1 + \gamma^2 A_1)}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \|x_k - x_*\|^2 \right] + \frac{\gamma^2 B_1}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \sigma_k^2 \right] + \frac{\gamma^2 C_1}{(1 + \gamma\mu)^2} + \alpha \mathbb{E} \left[ \sigma_{k+1}^2 \right] \\ &\stackrel{(8)}{\leq} \frac{(1 + \gamma^2 A_1)}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \|x_k - x_*\|^2 \right] + \frac{\gamma^2 B_1}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \sigma_k^2 \right] + \frac{\gamma^2 C_1}{(1 + \gamma\mu)^2} \\ &\quad + \alpha A_2 \mathbb{E} \left[ \|x_{k+1} - x_*\|^2 \right] + \alpha B_2 \mathbb{E} \left[ \sigma_k^2 \right] + \alpha C_2 \\ &\stackrel{(70)}{\leq} \frac{(1 + \gamma^2 A_1)}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \|x_k - x_*\|^2 \right] + \frac{\gamma^2 B_1}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \sigma_k^2 \right] + \frac{\gamma^2 C_1}{(1 + \gamma\mu)^2} \\ &\quad + \frac{(1 + \gamma^2 A_1)\alpha A_2}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \|x_k - x_*\|^2 \right] + \frac{\gamma^2 B_1 \alpha A_2}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \sigma_k^2 \right] + \frac{\gamma^2 C_1 \alpha A_2}{(1 + \gamma\mu)^2} \\ &\quad + \alpha B_2 \mathbb{E} \left[ \sigma_k^2 \right] + \alpha C_2 \\ &= \frac{(1 + \gamma^2 A_1)(1 + \alpha A_2)}{(1 + \gamma\mu)^2} \mathbb{E} \left[ \|x_k - x_*\|^2 \right] + \left( \frac{\gamma^2 B_1(1 + \alpha A_2)}{\alpha(1 + \gamma\mu)^2} + B_2 \right) \alpha \mathbb{E} \left[ \sigma_k^2 \right] \\ &\quad + \frac{\gamma^2 C_1(1 + \alpha A_2)}{(1 + \gamma\mu)^2} + \alpha C_2 \\ &\leq \underbrace{\max \left\{ \frac{(1 + \gamma^2 A_1)(1 + \alpha A_2)}{(1 + \gamma\mu)^2}, \frac{\gamma^2 B_1(1 + \alpha A_2)}{\alpha(1 + \gamma\mu)^2} + B_2 \right\}}_{:=\theta} \mathbb{E}[\Psi_k] \\ &\quad + \underbrace{\frac{\gamma^2 C_1(1 + \alpha A_2)}{(1 + \gamma\mu)^2} + \alpha C_2}_{:=\zeta} \\ &= \theta \mathbb{E}[\Psi_k] + \zeta. \end{aligned}$$

By unrolling the recurrence (we can also just apply Fact 3), we get

$$\mathbb{E}[\Psi_k] \leq \theta^k \Psi_0 + \sum_{t=0}^{k-1} \theta^{k-1-t} \zeta = \theta^k \Psi_0 + \sum_{t=0}^{k-1} \theta^t \zeta \leq \theta^k \Psi_0 + \sum_{t=0}^{\infty} \theta^t \zeta = \theta^k \Psi_0 + \frac{1}{1-\theta} \zeta.$$

□

## B.2 STOCHASTIC PROXIMAL POINT METHOD (SPPM)

We aim to solve problem (1) via the stochastic proximal point method (SPPM), formalized as Algorithm 2. The main step has the form

$$\text{prox}_{\gamma f_{\xi_k}}(x_k) := \arg \min_{x \in \mathbb{R}^d} \left\{ f_{\xi_k}(x) + \frac{1}{2\gamma} \|x - x_k\|^2 \right\},$$

where  $\gamma > 0$  is a stepsize.

---

### Algorithm 2 Stochastic Proximal Point Method (SPPM)

---

- 1: **Parameters:** learning rate  $\gamma > 0$ , starting point  $x_0 \in \mathbb{R}^d$
  - 2: **for**  $k = 0, 1, 2, \dots$  **do**
  - 3:   Sample  $\xi_k \sim \mathcal{D}$
  - 4:    $x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k)$
  - 5: **end for**
- 

Notice that if  $\gamma$  is kept “too large”, then  $x_{k+1} \approx \arg \min_x f_{\xi_k}(x)$ , and hence the method just “samples” the minimizers of the stochastic functions, and does not progress towards finding  $x_*$  (unless, of course,  $x_*$  happens to be shared by all functions  $f_{\xi_k}$ ). As we shall see, the situation is different if  $\gamma$  is kept sufficiently small.

**Lemma 2 (SPPM).** *Assumption 7 holds for the iterates of SPPM (Algorithm 2) with*

$$A_1 = 0, B_1 = 0, C_1 = \sigma_*^2, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

*Proof of Lemma 2.* Recall that the iterates of SPPM have the form

$$x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k).$$

Thus,  $h_k = 0$ . Let  $\sigma^2 \equiv 0$ . Clearly, (7) holds. Furthermore,

$$\mathbb{E} \left[ \|h_k - \nabla f_{\xi_k}(x_*)\|^2 \middle| x_k, \sigma_k^2 \right] = \mathbb{E} \left[ \|\nabla f_{\xi_k}(x_*)\|^2 \right] := \sigma_*^2.$$

Hence, Assumption 7 holds with

$$A_1 = 0, B_1 = 0, C_1 = \sigma_*^2, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

□

The convergence of SPPM is captured by the following theorem.

**Theorem 2.** *Let Assumption 1 (differentiability) and Assumption 2 ( $\mu$ -strong convexity) hold and define*

$$\sigma_*^2 := \mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|\nabla f_{\xi}(x_*)\|^2 \right]. \quad (20)$$

*Let  $x_0 \in \mathbb{R}^d$  be an arbitrary starting point. Then for any  $k \geq 0$  and any  $\gamma > 0$ , the iterates of SPPM (Algorithm 2) satisfy*

$$\mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1 + \gamma\mu} \right)^{2k} \|x_0 - x_*\|^2 + \frac{\gamma\sigma_*^2}{\gamma\mu^2 + 2\mu}. \quad (21)$$

Commentary:



- 864  
865  
866  
867  
868  
869
1. **Interpolation regime.** Consider the interpolation regime, characterized by  $\sigma_*^2 = 0$ . Since we can use arbitrarily large  $\gamma > 0$ , we obtain an arbitrarily fast convergence rate:

$$\mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1 + \gamma\mu} \right)^{2k} \|x_0 - x_*\|^2. \quad (22)$$

870  
871  
872  
873  
874

Indeed,  $\left( \frac{1}{1 + \gamma\mu} \right)^{2k}$  can be made arbitrarily small for any fixed  $k \geq 1$ , even  $k = 1$ , by choosing  $\gamma$  large enough. However, this is not surprising, since now  $f$  and all functions  $f_\xi$  share a single minimizer,  $x_*$ , and hence it is possible to find it by sampling a single function  $f_\xi$  and minimizing it, which is what the prox does, as long as  $\gamma$  is large enough.

- 875  
876  
877  
878
2. **A single step travels far.** Observe that for  $\gamma = \frac{1}{\mu}$ , we have  $\frac{\gamma\sigma_*^2}{\gamma\mu^2 + 2\mu} = \frac{\sigma_*^2}{3\mu^2}$ . In fact, the convergence neighborhood  $\frac{\gamma\sigma_*^2}{\gamma\mu^2 + 2\mu}$  is bounded above by three times this quantity irrespective of the choice of the stepsize. Indeed,

$$\frac{\gamma\sigma_*^2}{\gamma\mu^2 + 2\mu} \leq \min \left\{ \frac{\sigma_*^2}{\mu^2}, \frac{\gamma\sigma_*^2}{\mu} \right\} \leq \frac{\sigma_*^2}{\mu^2}.$$

879  
880  
881  
882  
883  
884

This means that no matter how far the starting point  $x_0$  is from the optimal solution  $x_*$ , if we choose the stepsize  $\gamma$  to be large enough, then we can get a decent-quality solution after a single iteration of SPPM already! Indeed, if we choose  $\gamma$  large enough so that

$$\left( \frac{1}{1 + \gamma\mu} \right)^2 \|x_0 - x_*\|^2 \leq \delta,$$

885  
886  
887  
888

where  $\delta > 0$  is chosen arbitrarily, then for  $k = 1$  we get

$$\mathbb{E} \left[ \|x_1 - x_*\|^2 \right] \leq \delta + \frac{\sigma_*^2}{\mu^2}.$$

- 889  
890  
891  
892  
893  
894  
895  
896
3. **Iteration complexity.** We have seen above that accuracy arbitrarily close to (but not reaching)  $\frac{\sigma_*^2}{\mu^2}$  can be achieved via a single step of the method, provided the stepsize  $\gamma$  is large enough. Assume now that we aim for  $\varepsilon$  accuracy where  $\varepsilon \leq \frac{\sigma_*^2}{\mu^2}$ . Using the inequality  $1 - t \leq \exp(-t)$  which holds for all  $t > 0$ , we get

$$\left( \frac{1}{1 + \gamma\mu} \right)^{2k} = \left( 1 - \frac{\gamma\mu}{1 + \gamma\mu} \right)^{2k} \leq \exp \left( -\frac{2\gamma\mu k}{1 + \gamma\mu} \right).$$

897  
898  
899  
900

Therefore, provided that

$$k \geq \frac{1 + \gamma\mu}{2\gamma\mu} \log \left( \frac{2\|x_0 - x_*\|^2}{\varepsilon} \right),$$

901  
902  
903  
904  
905  
906  
907  
908

we get  $\left( \frac{1}{1 + \gamma\mu} \right)^{2k} \|x_0 - x_*\|^2 \leq \frac{\varepsilon}{2}$ . Furthermore, as long as  $\gamma \leq \frac{2\varepsilon\mu}{2\sigma_*^2 - \varepsilon\mu^2}$  (this is true provided that the more restrictive but also more elegant-looking condition  $\gamma \leq \varepsilon \frac{\mu}{\sigma_*^2}$  holds), we get

$$\frac{\gamma\sigma_*^2}{\gamma\mu^2 + 2\mu} \leq \frac{\varepsilon}{2}.$$

909  
910  
911  
912

Putting these observations together, we conclude that with the stepsize  $\gamma = \varepsilon \frac{\mu}{\sigma_*^2}$ , we get

$$\mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \varepsilon$$

913  
914  
915

provided that

$$k \geq \frac{1 + \gamma\mu}{2\gamma\mu} \log \frac{2\|x_0 - x_*\|^2}{\varepsilon} = \left( \frac{\sigma_*^2}{2\varepsilon\mu^2} + \frac{1}{2} \right) \log \left( \frac{2\|x_0 - x_*\|^2}{\varepsilon} \right). \quad (23)$$

*Proof of Theorem 2.* Recall that in the case of SPPM, we have  $A_1 = 0, B_1 = 0, C_1 = \sigma_*^2, A_2 = 0, B_2 = 0, C_2 = 0$  (it is the result of Lemma 2). From Theorem 1, choosing any  $\alpha > 0, \theta = \frac{1}{(1+\gamma\mu)^2}, \zeta = \frac{\gamma^2\sigma_*^2}{(1+\gamma\mu)^2}$  (see (15) and (16)), we obtain

$$\frac{\zeta}{1-\theta} = \frac{\gamma^2\sigma_*^2}{(1+\gamma\mu)^2(1-\theta)} = \frac{\gamma^2\sigma_*^2}{(1+\gamma\mu)^2-1} = \frac{\gamma\sigma_*^2}{\gamma\mu^2+2\mu}.$$

For any  $k \geq 0$  and any  $\gamma > 0$ , the iterates of SPPM (Algorithm 2) satisfy

$$\mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1+\gamma\mu} \right)^{2k} \|x_0 - x_*\|^2 + \frac{\gamma\sigma_*^2}{\gamma\mu^2+2\mu}.$$

□

### B.3 STOCHASTIC PROXIMAL POINT METHOD WITH NONUNIFORM SAMPLING (SPPM-NS)

Applying SPPM directly to the optimization formulation (3), Theorem 2 applies with  $\mu = \min_i \mu_i$  since  $f$  is  $\mu$ -strongly convex for  $\mu = \min_i \mu_i$ . If  $\min_i \mu_i$  is small, the convergence rate becomes weak. In this section we shall describe a trick which enables us to obtain a dependence on the average of the strong convexity constants instead of their minimum.

Choose positive numbers  $p_1, \dots, p_n$  summing up to 1, and let

$$\tilde{f}_i(x) := \frac{1}{np_i} f_i(x), \quad i \in [n].$$

Note that (3) can be reformulated in the form

$$\min_{x \in \mathbb{R}^d} \left\{ f(x) := \sum_{i=1}^n p_i \tilde{f}_i(x) \right\}. \quad (24)$$

We rely on Assumptions 3 and 4.

This is a more refined version of Assumption 2, where we assumed the strong convexity parameter was the same for all functions. This implies that  $f$  is  $\mu$ -strongly convex with  $\mu = \min_i \mu_i$ . Hence,  $f$  has a unique minimizer, which we shall denote  $x_*$ .

We can now apply SPPM to the reformulated problem (24) instead, with  $\mathcal{D}$  being the nonuniform distribution over the finite set  $[n]$  given by the parameters  $p_1, \dots, p_n$  as follows:  $\xi = i$  with probability  $p_i > 0$ . This method is called *stochastic proximal point method with nonuniform sampling* (SPPM-NS):

$$x_{k+1} = \text{prox}_{\gamma \tilde{f}_{i_k}}(x_k) = \text{prox}_{\frac{\gamma}{np_i} f_{i_k}}(x_k),$$

where  $i_k = i$  with probability  $p_i > 0$ .

---

#### Algorithm 3 Stochastic Proximal Point Method with Nonuniform Sampling (SPPM-NS)

---

- 1: **Parameters:** learning rate  $\gamma > 0$ , starting point  $x_0 \in \mathbb{R}^d$ , positive probabilities  $p_1, \dots, p_n$  summing up to 1
  - 2: **for**  $k = 0, 1, 2, \dots$  **do**
  - 3:   Choose  $i_k = i$  with probability  $p_i > 0$
  - 4:    $x_{k+1} = \text{prox}_{\frac{\gamma}{np_i} f_{i_k}}(x_k)$
  - 5: **end for**
- 

Define

$$\mu_{\text{NS}} := \min_i \frac{\mu_i}{np_i}, \quad \sigma_{*,\text{NS}}^2 := \frac{1}{n} \sum_{i=1}^n \frac{1}{np_i} \|\nabla f_i(x_*)\|^2. \quad (25)$$

**Lemma 3** (SPPM-NS). *Assumption 7 holds for the iterates of SPPM-NS (Algorithm 3) with*

$$A_1 = 0, B_1 = 0, C_1 = \sigma_{*,\text{NS}}^2, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

972 *Proof.* Recall that the iterates of (SPPM-NS) have the form

$$973 \quad x_{k+1} = \text{prox}_{\gamma \tilde{f}_{i_k}}(x_k) = \text{prox}_{\frac{\gamma}{np_i} f_{i_k}}(x_k),$$

974 where  $i_k = i$  with probability  $p_i > 0$ . Therefore, that  $h_k = 0$ . Let  $\sigma^2 \equiv 0$ . Clearly, (7) holds.  
975 Furthermore,

$$976 \quad \mathbb{E} \left[ \left\| h_k - \nabla \tilde{f}_{i_k}(x_*) \right\|^2 \middle| x_k, \sigma_k^2 \right] = \mathbb{E} \left[ \left\| \nabla \tilde{f}_{i_k}(x_*) \right\|^2 \right] := \sigma_{*,\text{NS}}^2.$$

977 Hence, Assumption 7 holds with

$$978 \quad A_1 = 0, B_1 = 0, C_1 = \sigma_{*,\text{NS}}^2, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

981  $\square$

982 **Theorem 3.** *Let Assumption 3 and Assumption 4 hold. Let  $x_0 \in \mathbb{R}^d$  be an arbitrary starting point. Let  $\mu_{\text{NS}}$  and  $\sigma_{*,\text{NS}}^2$  be as in (25). Then for any  $k \geq 0$  and any  $\gamma > 0$ , the iterates of SPPM-NS (Algorithm 3) satisfy*

$$983 \quad \mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1 + \gamma \mu_{\text{NS}}} \right)^{2k} \|x_0 - x_*\|^2 + \frac{\gamma \sigma_{*,\text{NS}}^2}{\gamma \mu_{\text{NS}}^2 + 2\mu_{\text{NS}}}. \quad (26)$$

984 Commentary:

985 (a) **Uniform sampling.** If we choose  $p_i = \frac{1}{n}$  for all  $i \in [n]$ , we shall refer to Algorithm 3 as *stochastic proximal point method with uniform sampling* (SPPM-US). In this case,

$$986 \quad \mu_{\text{NS}} = \mu_{\text{US}} := \min_i \mu_i, \quad \sigma_{*,\text{NS}}^2 = \sigma_{*,\text{US}}^2 := \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(x_*)\|^2.$$

987 (b) **Importance sampling: optimizing the linear rate.** If we choose  $p_i = \frac{\mu_i}{\sum_{j=1}^n \mu_j}$  for all  $i \in [n]$ , we shall refer to Algorithm 3 as *stochastic proximal point method with importance sampling* (SPPM-IS). In this case,

$$988 \quad \mu_{\text{NS}} = \mu_{\text{IS}} := \frac{1}{n} \sum_{i=1}^n \mu_i, \quad \sigma_{*,\text{NS}}^2 = \sigma_{*,\text{IS}}^2 := \frac{\sum_{i=1}^n \mu_i \sum_{i=1}^n \|\nabla f_i(x_*)\|^2}{n \mu_i}.$$

989 This choice maximizes the value of  $\mu_{\text{NS}}$  (and hence the first part of the convergence rate) over the choice of the probabilities.

990 (c) **Variance sampling: optimizing the variance.** If we choose  $p_i = \frac{\|\nabla f_i(x_*)\|}{\sum_{j=1}^n \|\nabla f_j(x_*)\|}$  for all  $i \in [n]$ , we shall refer to Algorithm 3 as *stochastic proximal point method with variance sampling* (SPPM-VS). In this case,

$$991 \quad \mu_{\text{NS}} = \mu_{\text{VS}} := \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(x_*)\| \left( \min_i \frac{\mu_i}{\|\nabla f_i(x_*)\|} \right),$$

$$992 \quad \sigma_{*,\text{NS}}^2 = \sigma_{*,\text{VS}}^2 := \left( \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(x_*)\| \right)^2.$$

993 This choice minimizes the value of  $\sigma_{*,\text{NS}}$  (and hence the second part of the convergence rate) over the choice of the probabilities.

994 *Proof of Theorem 3.* From Lemma 3, we have that Assumption 7 holds for the iterates of SPPM-NS (Algorithm 3) with

$$995 \quad A_1 = 0, B_1 = 0, C_1 = \sigma_{*,\text{NS}}^2, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

996 From Theorem 1, choosing any  $\alpha > 0$ ,  $\theta = \frac{1}{(1 + \gamma \mu_{\text{NS}})^2}$ ,  $\zeta = \frac{\gamma^2 \sigma_{*,\text{NS}}^2}{(1 + \gamma \mu_{\text{NS}})^2}$  (see (15) and (16)), we obtain that, for any  $k \geq 0$  and any  $\gamma > 0$ , the iterates of SPPM-NS (Algorithm 3) satisfy

$$997 \quad \mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1 + \gamma \mu_{\text{NS}}} \right)^{2k} \|x_0 - x_*\|^2 + \frac{\gamma \sigma_{*,\text{NS}}^2}{\gamma \mu_{\text{NS}}^2 + 2\mu_{\text{NS}}}.$$

998 Alternatively, we can derive Theorem 3 from Theorem 2. It is easy to show that each function  $\tilde{f}_i$  is  $\tilde{\mu}_i$ -strongly convex with  $\tilde{\mu}_i := \frac{\mu_i}{np_i}$ . Hence,  $\tilde{f}_i$  is also  $\mu_{\text{NS}}$ -strongly convex for all  $i$ . It now only remains to apply Theorem 2.  $\square$

#### B.4 STOCHASTIC PROXIMAL POINT METHOD WITH ARBITRARY SAMPLING (SPPM-AS)

We consider the optimization problem (3) and rely on Assumption 3 (differentiability) and Assumption 4 (strong convexity).

Let  $\mathcal{S}$  be a probability distribution over the  $2^n$  subsets of  $[n]$ . Given a random set  $S \sim \mathcal{S}$ , we define

$$p_i := \mathbf{Prob}(i \in S), \quad i \in [n]. \quad (27)$$

We will restrict our attention to proper and nonvacuous random sets.

**Assumption 8.**  $S$  is proper (i.e.,  $p_i > 0$  for all  $i \in [n]$ ) and nonvacuous (i.e.,  $\mathbf{Prob}(S = \emptyset) = 0$ ).

Given  $\emptyset \neq C \subseteq [n]$  and  $i \in [n]$ , we define

$$v_i(C) := \begin{cases} \frac{1}{p_i} & i \in C \\ 0 & i \notin C \end{cases}, \quad (28)$$

and

$$f_C(x) := \frac{1}{n} \sum_{i=1}^n v_i(C) f_i(x) \stackrel{(28)}{=} \sum_{i \in C} \frac{1}{np_i} f_i(x). \quad (29)$$

Note that  $v_i(S)$  is a random variable and  $f_S$  is a random function. By construction,  $\mathbb{E}_{S \sim \mathcal{S}} [v_i(S)] = 1$  for all  $i \in [n]$ , and hence

$$\mathbb{E}_{S \sim \mathcal{S}} [f_S(x)] = \mathbb{E}_{S \sim \mathcal{S}} \left[ \frac{1}{n} \sum_{i=1}^n v_i(S) f_i(x) \right] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{S \sim \mathcal{S}} [v_i(S)] f_i(x) = \frac{1}{n} \sum_{i=1}^n f_i(x) = f(x).$$

Therefore, the optimization problem (3) is equivalent to the stochastic optimization problem

$$\min_{x \in \mathbb{R}^d} \{f(x) := \mathbb{E}_{S \sim \mathcal{S}} [f_S(x)]\}. \quad (30)$$

Further, if for each  $C \subset [n]$  we let  $p_C := \mathbf{Prob}(S = C)$ ,  $f$  can be written in the equivalent form

$$f(x) = \mathbb{E}_{S \sim \mathcal{S}} [f_S(x)] = \sum_{C \subseteq [n]} p_C f_C(x) = \sum_{C \subseteq [n], p_C > 0} p_C f_C(x). \quad (31)$$

Applying SPPM to (30), we arrive at *stochastic proximal point method with arbitrary sampling* (SPPM-AS) (Algorithm 4):

$$x_{k+1} = \text{prox}_{\gamma f_{S_k}}(x_k),$$

where  $S_k \sim \mathcal{S}$ .

---

#### Algorithm 4 Stochastic Proximal Point Method with Arbitrary Sampling (SPPM-AS)

---

- 1: **Parameters:** learning rate  $\gamma > 0$ , starting point  $x_0 \in \mathbb{R}^d$ , distribution  $\mathcal{S}$  over the subsets of  $[n]$
  - 2: **for**  $k = 0, 1, 2, \dots$  **do**
  - 3:   Sample  $S_k \sim \mathcal{S}$
  - 4:    $x_{k+1} = \text{prox}_{\gamma f_{S_k}}(x_k)$
  - 5: **end for**
- 

Define

$$\mu_{\text{AS}} := \min_{C \subseteq [n], p_C > 0} \sum_{i \in C} \frac{\mu_i}{np_i}, \quad \sigma_{\star, \text{AS}}^2 := \sum_{C \subseteq [n], p_C > 0} p_C \left\| \sum_{i \in C} \frac{1}{np_i} \nabla f_i(x_\star) \right\|^2. \quad (32)$$

**Lemma 4.** *Assumption 7 holds for the iterates of SPPM-AS (Algorithm 4) with*

$$A_1 = 0, B_1 = 0, C_1 = \sigma_{\star, \text{AS}}^2, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

1080 *Proof.* Recall that the iterates of SPPM-AS (Algorithm 4) have the form

$$1081 \quad x_{k+1} = \text{prox}_{\gamma f_{S_k}}(x_k).$$

1082 Thus,  $h_k = 0$ . Let  $\sigma^2 \equiv 0$ . Clearly, (7) holds. Furthermore,

$$1083 \quad \mathbb{E} \left[ \|h_k - \nabla f_C(x_*)\|^2 \mid x_k, \sigma_k^2 \right] = \mathbb{E} \left[ \|\nabla f_C(x_*)\|^2 \right] := \sigma_{*,AS}^2,$$

1084 since

$$\begin{aligned} 1085 \quad \sigma_*^2 &:= \mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|\nabla f_\xi(x_*)\|^2 \right] \\ 1086 &\stackrel{(31)}{=} \sum_{C \subseteq [n], p_C > 0} p_C \|\nabla f_C(x_*)\|^2 \\ 1087 &\stackrel{(29)}{=} \sum_{C \subseteq [n], p_C > 0} p_C \left\| \sum_{i \in C} \frac{1}{np_i} \nabla f_i(x_*) \right\|^2 \\ 1088 &:= \sigma_{*,AS}^2. \end{aligned}$$

1089 Hence, Assumption 7 holds with

$$1090 \quad A_1 = 0, B_1 = 0, C_1 = \sigma_{*,AS}^2, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

1091  $\square$

1092 The convergence of SPPM-AS is captured by the following theorem.

1093 **Theorem 4.** *Let Assumption 3 (differentiability) and Assumption 4 (strong convexity) hold. Let  $S$*   
1094 *be a random set satisfying Assumption 8. Let  $x_0 \in \mathbb{R}^d$  be an arbitrary starting point,  $\mu_{AS}$  and  $\sigma_{*,AS}^2$*   
1095 *be as in (32). Then for any  $k \geq 0$  and any  $\gamma > 0$ , the iterates of SPPM-AS (Algorithm 4) satisfy*

$$1096 \quad \mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1 + \gamma \mu_{AS}} \right)^{2k} \|x_0 - x_*\|^2 + \frac{\gamma \sigma_{*,AS}^2}{\gamma \mu_{AS}^2 + 2\mu_{AS}}. \quad (33)$$

1097 **Commentary:**

- 1098 (a) **Full sampling.** Let  $S = [n]$  with probability 1 (“full sampling”; abbreviated as “FS”).  
1099 Then SPPM-AS applied to (30) becomes PPM for minimizing  $f$ . So, besides Theorem 2,  
1100 Theorem 4 is the third theorem capturing the convergence of PPM in the differentiable and  
1101 strongly convex regime.

1102 Moreover, we have  $p_i = 1$  for all  $i \in [n]$  and the expressions (32) take on the form

$$1103 \quad \mu_{AS} = \mu_{FS} := \frac{1}{n} \sum_{i=1}^n \mu_i, \quad \sigma_{*,AS}^2 = \sigma_{*,FS}^2 := 0.$$

1104 Note that  $\mu_{FS}$  a lower bound on the the strong convexity constant of  $f$  – one that can be  
1105 computed from the strong convexity constants  $\mu_i$  of the constituent functions  $f_i$ .

- 1106 (b) **Nonuniform sampling.** Let  $S = \{i\}$  with probability  $q_i > 0$ , where  $\sum_i q_i = 1$ . This leads  
1107 to SPPM-NS. Then  $p_i := \mathbf{Prob}(i \in S) = q_i$  for all  $i \in [n]$ , and the expressions (32) take  
1108 on the form

$$1109 \quad \mu_{AS} = \mu_{NS} := \min_i \frac{\mu_i}{np_i}, \quad \sigma_{*,AS}^2 = \sigma_{*,NS}^2 := \frac{1}{n} \sum_{i=1}^n \frac{1}{np_i} \|\nabla f_i(x_*)\|^2,$$

1110 which recovers the quantities from Section B.3; see (25).

- 1111 (c) **Nice sampling.** Choose  $\tau \in [n]$  and let  $S$  be a random subset of  $[n]$  of size  $\tau$  chosen uni-  
1112 formly at random. We call the resulting method SPPM-NICE. Then  $p_i := \mathbf{Prob}(i \in S) =$   
1113  $\frac{\tau}{n}$  for all  $i \in [n]$ . Moreover,  $p_C = \frac{1}{\binom{n}{\tau}}$  whenever  $|C| = \tau$  and  $p_C = 0$  otherwise. So,

$$1114 \quad \mu_{AS} = \mu_{NICE}(\tau) := \min_{C \subseteq [n], |C|=\tau} \frac{1}{\tau} \sum_{i \in C} \mu_i \quad (34)$$

and

$$\sigma_{\star,AS}^2 = \sigma_{\star,NICE}^2(\tau) := \sum_{C \subseteq [n], |C|=\tau} \frac{1}{\binom{n}{\tau}} \left\| \frac{1}{\tau} \sum_{i \in C} \nabla f_i(x_\star) \right\|^2. \quad (35)$$

It can be shown that  $\mu_{NICE}(\tau)$  is a nondecreasing function of  $\tau$ . So, as the minibatch size  $\tau$  increases, the strong convexity constant  $\mu_{NICE}(\tau)$  can only improve. Since  $\mu_{NICE}(1) = \min_i \mu_i$  and  $\mu_{NICE}(n) = \frac{1}{n} \sum_{i=1}^n \mu_i$ , the value of  $\mu_{NICE}(\tau)$  interpolates these two extreme cases as  $\tau$  varies between 1 and  $n$ .

- (d) **Block sampling.** Let  $C_1, \dots, C_b$  be a partition of  $[n]$  into  $b$  nonempty blocks. For each  $i \in [n]$ , let  $B(i)$  indicate which block does  $i$  belong to. In other words,  $i \in C_j$  iff  $B(i) = j$ . Let  $S = C_j$  with probability  $q_j > 0$ , where  $\sum_j q_j = 1$ . We call the resulting method SPPM-BS. Then  $p_i := \mathbf{Prob}(i \in S) = q_{B(i)}$ , and hence the expressions (32) take on the form

$$\mu_{AS} = \mu_{BS} := \min_{j \in [b]} \frac{1}{nq_j} \sum_{i \in C_j} \mu_i, \quad \sigma_{\star,AS}^2 = \sigma_{\star,BS}^2 := \sum_{j \in [b]} q_j \left\| \sum_{i \in C_j} \frac{1}{np_j} \nabla f_i(x_\star) \right\|^2.$$

We now consider two extreme cases:

- If  $b = 1$ , then SPPM-BS = SPPM-FS = PPM. Let's see, as a sanity check, whether we recover the right rate as well. We have  $q_1 = 1$ ,  $C_1 = [n]$ ,  $p_i := \mathbf{Prob}(i \in S) = 1$  for all  $i \in [n]$ , and the expressions for  $\mu_{AS}$  and  $\sigma_{\star,BS}^2$  simplify to

$$\mu_{BS} = \mu_{FS} := \frac{1}{n} \sum_{i=1}^n \mu_i, \quad \sigma_{\star,BS}^2 = \sigma_{\star,FS}^2 := 0.$$

So, indeed, we recover the same rate as SPPM-FS.

- If  $b = n$ , then SPPM-BS = SPPM-NS. Let's see, as a sanity check, whether we recover the right rate as well. We have  $C_i = \{i\}$  and  $p_i := \mathbf{Prob}(i \in S) = q_i$  for all  $i \in [n]$ , and the expressions for  $\mu_{AS}$  and  $\sigma_{\star,BS}^2$  simplify to

$$\mu_{BS} = \mu_{NS} := \min_{i \in [n]} \frac{\mu_i}{np_i}, \quad \sigma_{\star,BS}^2 = \sigma_{\star,NS}^2 := \frac{1}{n} \sum_{i=1}^n \frac{1}{np_i} \|\nabla f_i(x_\star)\|^2.$$

So, indeed, we recover the same rate as SPPM-NS.

- (e) **Stratified sampling.** Let  $C_1, \dots, C_b$  be a partition of  $[n]$  into  $b$  nonempty blocks, as before. For each  $i \in [n]$ , let  $B(i)$  indicate which block  $i$  belongs to. In other words,  $i \in C_j$  iff  $B(i) = j$ . Now, for each  $j \in [b]$  pick  $\xi_j \in C_j$  uniformly at random, and define  $S = \cup_{j \in [b]} \{\xi_j\}$ . We call the resulting method SPPM-SS. Clearly,  $p_i := \mathbf{Prob}(i \in S) = \frac{1}{|C_{B(i)}}$ . The expressions (32) take on the form

$$\mu_{AS} = \mu_{SS} := \min_{(i_1, \dots, i_b) \in C_1 \times \dots \times C_b} \sum_{j=1}^b \frac{\mu_{i_j} |C_j|}{n}$$

and

$$\sigma_{\star,AS}^2 = \sigma_{\star,SS}^2 := \left( \frac{1}{\prod_{j=1}^b |C_j|} \right) \sum_{(i_1, \dots, i_b) \in C_1 \times \dots \times C_b} \left\| \sum_{j=1}^b \frac{|C_j|}{n} \nabla f_{i_j}(x_\star) \right\|^2.$$

We now consider two extreme cases:

- If  $b = 1$ , then SPPM-SS = SPPM-US. Let's see, as a sanity check, whether we recover the right rate as well. We have  $C_1 = [n]$ ,  $|C_1| = n$ ,  $\left( \prod_{j=1}^b \frac{1}{|C_j|} \right) = \frac{1}{n}$  and hence

$$\mu_{SS} = \mu_{US} := \min_i \mu_i, \quad \sigma_{\star,SS}^2 = \sigma_{\star,US}^2 := \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(x_\star)\|^2.$$

So, indeed, we recover the same rate as SPPM-US.

- If  $b = n$ , then SPPM-SS = SPPM-FS. Let’s see, as a sanity check, whether we recover the right rate as well. We have  $C_i = \{i\}$  for all  $i \in [n]$ ,  $\left(\prod_{j=1}^b \frac{1}{|C_j|}\right) = 1$ , and hence

$$\mu_{\text{SS}} = \mu_{\text{FS}} := \frac{1}{n} \sum_{i=1}^n \mu_i, \quad \sigma_{\star, \text{SS}}^2 = \sigma_{\star, \text{FS}}^2 := 0.$$

So, indeed, we recover the same rate as SPPM-FS.

*Proof of Theorem 4.* From Lemma 4 we have that Assumption 7 holds for the iterates of SPPM-AS (Algorithm 4) with

$$A_1 = 0, B_1 = 0, C_1 = \sigma_{\star, \text{AS}}^2, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

From Theorem 1, choosing  $\alpha > 0$ ,  $\theta = \frac{1}{(1+\gamma\mu)^2}$ ,  $\zeta = \frac{\gamma^2 \sigma_{\star, \text{AS}}^2}{(1+\gamma\mu)^2}$  (see (15) and (16)), we obtain that, for any  $k \geq 0$  and any  $\gamma > 0$ , the iterates of SPPM-AS (Algorithm 4) satisfy

$$\mathbb{E} \left[ \|x_k - x_\star\|^2 \right] \leq \left( \frac{1}{1 + \gamma\mu_{\text{AS}}} \right)^{2k} \|x_0 - x_\star\|^2 + \frac{\gamma\sigma_{\star, \text{AS}}^2}{\gamma\mu_{\text{AS}}^2 + 2\mu_{\text{AS}}}.$$

Alternatively, we can derive Theorem 4 from Theorem 2. Let  $C$  be any (necessarily nonempty) subset of  $[n]$  such that  $p_C > 0$ . Recall that in view of (29) we have

$$f_C(x) = \sum_{i \in C} \frac{1}{np_i} f_i(x);$$

i.e.,  $f_C$  is a conic combination of the functions  $\{f_i : i \in C\}$  with weights  $w_i = \frac{1}{np_i}$ . Since each  $f_i$  is  $\mu_i$ -strongly convex, Lemma 10 says that  $f_C$  is  $\mu_C$ -strongly convex with

$$\mu_C := \sum_{i \in C} \frac{\mu_i}{np_i}.$$

So, every such  $f_C$  is  $\mu$ -strongly convex with

$$\mu = \mu_{\text{AS}} := \min_{C \subseteq [n], p_C > 0} \sum_{i \in C} \frac{\mu_i}{np_i}.$$

Further, the quantity  $\sigma_\star^2$  from (20) is equal to

$$\begin{aligned} \sigma_\star^2 &:= \mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|\nabla f_\xi(x_\star)\|^2 \right] \\ &\stackrel{(31)}{=} \sum_{C \subseteq [n], p_C > 0} p_C \|\nabla f_C(x_\star)\|^2 \\ &\stackrel{(29)}{=} \sum_{C \subseteq [n], p_C > 0} p_C \left\| \sum_{i \in C} \frac{1}{np_i} \nabla f_i(x_\star) \right\|^2 \\ &:= \sigma_{\star, \text{AS}}^2. \end{aligned}$$

It now only remains to apply Theorem 2. □

## B.5 STOCHASTIC PROXIMAL POINT METHOD WITH OPTIMAL GRADIENT CORRECTION (SPPM\*)

We showed that SPPM converges up to a neighborhood of size  $\frac{\gamma\sigma_\star^2}{\gamma\mu^2 + 2\mu}$ , where

$$\sigma_\star^2 := \mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|\nabla f_\xi(x_\star)\|^2 \right].$$

We now describe a simple trick which gets rid of the neighborhood when  $\sigma_\star^2 > 0$ . The trick is of a conceptual nature: as is, it is practically useless. However, it will serve as an inspiration for a trick that can be implemented.

**Algorithm 5** Stochastic Proximal Point Method with Optimal Gradient Correction (SPPM\*)

- 
- 1: **Parameters:** learning rate  $\gamma > 0$ , starting point  $x_0 \in \mathbb{R}^d$
  - 2: **for**  $k = 0, 1, 2, \dots$  **do**
  - 3:   Sample  $\xi_k \sim \mathcal{D}$
  - 4:    $h_k = \nabla f_{\xi_k}(x_*)$
  - 5:    $x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k + \gamma h_k)$   $= \text{prox}_{\gamma \tilde{f}_{\xi_k}}(x_k)$
  - 6: **end for**
- 

Let us reformulate the problem by adding a smart zero.

For each  $\xi \sim \mathcal{D}$ , define

$$\tilde{f}_\xi(x) := f_\xi(x) - \langle \nabla f_\xi(x_*), x \rangle, \quad (36)$$

and instead of solving (1), consider solving the problem

$$\min_{x \in \mathbb{R}^d} \left\{ \tilde{f}(x) := \mathbb{E}_{\xi \sim \mathcal{D}} [\tilde{f}_\xi(x)] \right\}. \quad (37)$$

Observations:

- We do not know  $\nabla f_\xi(x_*)$ , and hence formulation (37) is not of practical interest.
- It is easy to see that  $f = \tilde{f}$ , and hence problems (1) and (37) are equivalent. Indeed,

$$\begin{aligned} \tilde{f}(x) &\stackrel{(37)}{=} \mathbb{E}_{\xi \sim \mathcal{D}} [\tilde{f}_\xi(x)] \stackrel{(36)}{=} \mathbb{E}_{\xi \sim \mathcal{D}} [f_\xi(x) - \langle \nabla f_\xi(x_*), x \rangle] \\ &= \mathbb{E}_{\xi \sim \mathcal{D}} [f_\xi(x)] - \mathbb{E}_{\xi \sim \mathcal{D}} [\langle \nabla f_\xi(x_*), x \rangle] \\ &= \mathbb{E}_{\xi \sim \mathcal{D}} [f_\xi(x)] - \left\langle \underbrace{\mathbb{E}_{\xi \sim \mathcal{D}} [\nabla f_\xi(x_*)]}_{=\nabla f(x_*)=0}, x \right\rangle \stackrel{(1)}{=} f(x). \end{aligned}$$

- All stochastic gradients of  $\tilde{f}$  at  $x_*$  are zero. Indeed,

$$\nabla \tilde{f}(x) \stackrel{(36)}{=} \nabla f_\xi(x) - \nabla f_\xi(x_*),$$

and hence  $\nabla \tilde{f}_\xi(x_*) = 0$ .

- It is easy to see that since  $f_\xi$  is differentiable and  $\mu$ -strongly convex, then so is  $\tilde{f}_\xi$ .

**Hiding the prox.** Further, recall that if for some  $x \in \mathbb{R}^d$  and differentiable and convex  $\phi$  we let  $x_+ := \text{prox}_\phi(x)$ , then  $x_+ = x - \nabla \phi(x_+)$ . Therefore, steps 4 and 5 of the method can be written in the equivalent form

$$x_{k+1} = x_k + \gamma h_k - \gamma \nabla f_{\xi_k}(x_{k+1}) = x_k - \gamma (\nabla f_{\xi_k}(x_{k+1}) - \nabla f_{\xi_k}(x_*)).$$

**Lemma 5** (SPPM\*). *Assumption 7 holds for the iterates of SPPM\* (Algorithm 5) with*

$$A_1 = 0, B_1 = 0, C_1 = 0, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

*Proof.* Recall that the iterates of SPPM\* have the form

$$x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k + \gamma \nabla f_{\xi_k}(x_*)).$$

Thus,  $h_k = \nabla f_{\xi_k}(x_*)$ . Let  $\sigma^2 \equiv 0$ . Clearly, (7) holds. Furthermore,

$$\mathbb{E} \left[ \|h_k - \nabla f_{\xi_k}(x_*)\|^2 \mid x_k, \sigma_k^2 \right] = 0.$$

Hence, Assumption 7 holds with

$$A_1 = 0, B_1 = 0, C_1 = 0, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

□



1296 The convergence of SPPM\* is captured by the following theorem.

1297 **Theorem 5.** *Let Assumption 1 (differentiability), Assumption 2 ( $\mu$ -strong convexity) hold. Let  $x_0 \in \mathbb{R}^d$  be an arbitrary starting point. Then for any  $k \geq 0$  and any  $\gamma > 0$ , the iterates of SPPM\* (Algorithm 5) satisfy*

$$1301 \quad \mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1 + \gamma\mu} \right)^{2k} \|x_0 - x_*\|^2. \quad (38)$$

1302  
1303  
1304 Commentary:

- 1305 • The convergence neighborhood is fully removed; the method converges to the exact solution! The result is identical to (22); i.e., to the rate of SPPM in the interpolation regime.
  - 1306 – The method converges to  $x_*$  for any fixed  $\gamma > 0$  as long as  $k \rightarrow \infty$ .
  - 1307 – The method converges to  $x_*$  for any fixed  $k \geq 1$  as long as  $\gamma \rightarrow \infty$ .
- 1308 • The method is practically useless since it relies on the knowledge of the optimal stochastic gradients  $\nabla f_\xi(x_*)$  for all  $\xi$  as hyper-parameters of the method. Needless to say, these vectors are rarely known.
- 1309 • If we follow (36), add “smart zero” to (30), and apply SPPM to the resulting formulation, we automatically get a “star” variant of SPPM-AS, and Theorem 5 captures the complexity of this method.

1310  
1311  
1312 *Proof of Theorem 5.* From Lemma 5 we know that Assumption 7 holds for the iterates of SPPM\* (Algorithm 5) with

$$1313 \quad A_1 = 0, B_1 = 0, C_1 = 0, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

1314 Therefore, from Theorem 1, by choosing any  $\alpha > 0$ ,  $\theta = \left( \frac{1}{1 + \gamma\mu} \right)^2$ ,  $\zeta = 0$  (see (15) and (16)), we have

$$1315 \quad \mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1 + \gamma\mu} \right)^{2k} \|x_0 - x_*\|^2.$$

1316 □

## 1317 B.6 STOCHASTIC PROXIMAL POINT METHOD WITH GRADIENT CORRECTION (SPPM-GC)

1318 We consider the stochastic optimization problem (1), i.e.,

$$1319 \quad \min_{x \in \mathbb{R}^d} \{f(x) := \mathbb{E}_{\xi \sim \mathcal{D}} [f_\xi(x)]\},$$

1320 and rely on Assumption 1 (differentiability of  $f_\xi$ ) and Assumption 2 ( $\mu$ -strong convexity of  $f_\xi$ ). Recall that this implies strong convexity of  $f$ . Hence  $f$  has a unique minimizer, which we denote  $x_*$ .

1321 We have already described the SPPM\* method – this a variant of SPPM without the neighborhood term in the convergence bound. In order to run it, we need to know  $\nabla f_\xi(x_*)$  for all  $\xi$ , which is of course something we almost never know; one exception to this is the *interpolation regime*, defined by the assumption that  $\nabla f_\xi(x_*) = 0$  for all  $\xi$ .

1322 In this section we describe a practical method inspired by SPPM\*. The method is based on the following ideas:

- 1323 • While we do not know  $\nabla f_\xi(x_*)$ , what if this quantity could be in some appropriate/useful sense approximated by some vector  $g_\xi$  we *do* know?
- 1324 • One option is to require a quantity such as

$$1325 \quad \|g_\xi - \nabla f_\xi(x_*)\|^2 \quad \text{or} \quad \mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|g_\xi - \nabla f_\xi(x_*)\|^2 \right]$$

1326 to be “small” in some sense. However, where can such vectors come from? And what should “small” mean?

- One idea is to make an extra assumption on the functions  $f_\xi$  that would somehow automatically guarantee the existence and availability of such vectors. So, we trade off easy availability of vectors  $g_\xi$  for a limitation on the class of problems we solve this way.

Recall the Similarity Assumption 5. There exists  $\delta \geq 0$  such that

$$\mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|\nabla f_\xi(x) - \nabla f(x) - \nabla f_\xi(x_\star)\|^2 \right] \leq \delta^2 \|x - x_\star\|^2, \quad \forall x \in \mathbb{R}^d.$$

The above assumption says that we can consider the vectors  $g_\xi = \nabla f_\xi(x) - \nabla f(x)$  for any  $x$ , and that by “small enough” we require a bound by  $\delta^2 \|x - x_\star\|^2$ . Why these particular choices make sense will become clear from the convergence proof. Let us now make some observations about the class of functions satisfying Assumption 5:

1. The approximation of  $\nabla f_\xi(x_\star)$  by  $\nabla f_\xi(x) - \nabla f(x)$  gets better as  $x$  gets closer to  $x_\star$ .
2. The smaller  $\delta$  is, the better the approximation.
3. If  $\delta = 0$ , we get perfect approximation.
4. Since

$$\mathbb{E}_{\xi \sim \mathcal{D}} [\nabla f_\xi(x) - \nabla f_\xi(x_\star)] = \nabla f(x),$$

the left-hand side of (5) is the variance of the random vector  $\nabla f_\xi(x) - \nabla f_\xi(x_\star)$  as an estimator of  $\nabla f(x) = \nabla f(x) - \nabla f(x_\star)$ .

5. It follows that (5) can be equivalently written in the form

$$\mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|\nabla f_\xi(x) - \nabla f_\xi(x_\star)\|^2 \right] - \|\nabla f(x) - \nabla f(x_\star)\|^2 \leq \delta^2 \|x - x_\star\|^2, \quad \forall x \in \mathbb{R}^d. \quad (39)$$

6. Note that (39) holds if the following stronger condition holds:

$$\mathbb{E}_{\xi \sim \mathcal{D}} \left[ \|\nabla f_\xi(x) - \nabla f_\xi(x_\star)\|^2 \right] \leq \delta^2 \|x - x_\star\|^2, \quad \forall x \in \mathbb{R}^d. \quad (40)$$

7. Furthermore, (40) holds if the following even stronger condition holds: there exists  $\delta \geq 0$  such that

$$\|\nabla f_\xi(x) - \nabla f_\xi(x_\star)\| \leq \delta \|x - x_\star\|, \quad \forall x \in \mathbb{R}^d \quad (41)$$

for all  $\xi$ .

8. Finally, (41) holds if there exists  $\delta \geq 0$  such that  $\nabla f_\xi$  is  $\delta$ -Lipschitz for all  $\xi$ :

$$\|\nabla f_\xi(x) - \nabla f_\xi(y)\| \leq \delta \|x - y\|, \quad \forall x, y \in \mathbb{R}^d. \quad (42)$$

For each  $\xi \in [n]$  and any “parameter”  $v \in \mathbb{R}^d$ , define

$$\tilde{f}_\xi^v(x) := f_\xi(x) - \langle \nabla f_\xi(v) - \nabla f(v), x \rangle, \quad (43)$$

and instead of solving (1), consider solving the problem

$$\min_{x \in \mathbb{R}^d} \left\{ \tilde{f}^v(x) := \mathbb{E}_{\xi \sim \mathcal{D}} \left[ \tilde{f}_\xi^v(x) \right] \right\}. \quad (44)$$

---

#### Algorithm 6 Stochastic Proximal Point Method with Gradient Correction (SPPM-GC)

---

- 1: **Parameters:** learning rate  $\gamma > 0$ , starting point  $x_0 \in \mathbb{R}^d$
  - 2: **for**  $k = 0, 1, 2, \dots$  **do**
  - 3:   Sample  $\xi_k \sim \mathcal{D}$
  - 4:    $h_k = \nabla f_{\xi_k}(x_k) - \nabla f(x_k)$
  - 5:    $x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k + \gamma h_k) = \text{prox}_{\gamma \tilde{f}_{\xi_k}^{x_k}}(x_k)$
  - 6: **end for**
- 

The algorithm can be interpreted as:

- a practical variant of SPPM\* in which we use the computable correction  $h_k = \nabla f_{\xi_k}(x_k) - \nabla f(x_k)$  instead of incomputable correction  $\nabla f_{\xi_k}(x_*)$ ,
- SPPM applied to the reformulated problem (44), with the “control” vector  $v = x_k$  at iteration  $k$ :

$$v = x_k, \quad x_{k+1} = \text{prox}_{\gamma \tilde{f}_{\xi_k}^v}(x_k)$$

**Hiding the prox.** Further, recall that if for some  $x \in \mathbb{R}^d$  and differentiable and convex  $\phi$  we let  $x_+ := \text{prox}_\phi(x)$ , then  $x_+ = x - \nabla \phi(x_+)$ . Therefore, steps 4 and 5 of the method can be written in the equivalent form

$$x_{k+1} = x_k + \gamma h_k - \gamma \nabla f_{\xi_k}(x_{k+1}) = x_k - \gamma (\nabla f_{\xi_k}(x_{k+1}) - \nabla f_{\xi_k}(x_k) + \nabla f(x_k)).$$

**Lemma 6 (SPPM-GC).** *Suppose Assumption 5 holds with  $\delta > 0$ . Assumption 7 holds for the iterates of SPPM-GC (Algorithm 6) with*

$$A_1 = \delta^2, B_1 = 0, C_1 = 0, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

*Proof.* Recall that the iterates of SPPM-GC have the form

$$x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k + \gamma h_k),$$

where  $h_k = \nabla f_{\xi_k}(x_k) - \nabla f(x_k)$ . Let  $\sigma^2 \equiv 0$ . Since Assumption 5 holds, we get

$$\mathbb{E} \left[ \|h_k - \nabla f_{\xi_k}(x_*)\|^2 \middle| x_k, \sigma_k^2 \right] = \mathbb{E} \left[ \|\nabla f_{\xi_k}(x_k) - \nabla f(x_k) - \nabla f_{\xi_k}(x_*)\|^2 \middle| x_k \right] \leq \delta^2 \|x_k - x_*\|^2.$$

Hence, Assumption 7 holds with

$$A_1 = \delta^2, B_1 = 0, C_1 = 0, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

□

**Theorem 6.** *Let Assumption 1 (differentiability), Assumption 2 ( $\mu$ -strong convexity), and Assumption 5 ( $\delta$ -similarity) hold. Choose any  $x_0 \in \mathbb{R}^d$ . Then for any  $\gamma > 0$ , and all  $k \geq 0$ , we have*

$$\mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1 + \gamma^2 \delta^2}{(1 + \gamma \mu)^2} \right)^k \|x_0 - x_*\|^2. \quad (45)$$

Commentary:

1. **Perfect similarity.** If  $\delta = 0$ , for all  $\gamma > 0$  we get

$$\mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{1}{1 + \gamma \mu} \right)^{2k} \|x_0 - x_*\|^2. \quad (46)$$

We get convergence even with  $k = 1$  provided that  $\gamma$  is chosen large enough. This rate is identical to what Theorem 2 predicts in the interpolation regime. However, the methods are different, and we do not need to assume interpolation regime here. Instead, we assume perfect similarity ( $\delta = 0$ ), and the availability of the gradient of  $f$ . So, while the rates are exactly the same, both the methods and the assumptions are different.

2. **General case.** It can be shown that the expression  $\frac{1 + \gamma^2 \delta^2}{(1 + \gamma \mu)^2}$  is minimized for  $\gamma = \frac{\mu}{\delta^2}$ . With this choice of the stepsize we get

$$\mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \left( \frac{\delta^2}{\delta^2 + \mu^2} \right)^k \|x_0 - x_*\|^2 = \left( 1 - \frac{\mu^2}{\delta^2 + \mu^2} \right)^k \|x_0 - x_*\|^2. \quad (47)$$

This means that

$$k \geq \left( 1 + \frac{\delta^2}{\mu^2} \right) \log \left( \frac{\|x_0 - x_*\|^2}{\varepsilon} \right) \quad (48)$$

iterations suffice to guarantee  $\mathbb{E} \left[ \|x_k - x_*\|^2 \right] \leq \varepsilon$ .

1458 *Proof of Theorem 6.* From Lemma 6 we have that Assumption 7 holds for the iterates of SPPM-GC  
 1459 (Algorithm 6) with

$$1461 \quad A_1 = \delta^2, B_1 = 0, C_1 = 0, \quad \text{and} \quad A_2 = 0, B_2 = 0, C_2 = 0.$$

1463 From Theorem 1, choosing any  $\alpha > 0$ ,  $\theta = \frac{1+\gamma^2\delta^2}{(1+\gamma\mu)^2}$ ,  $\zeta = 0$  (see (15) and (16)), we have

$$1466 \quad \mathbb{E} \left[ \|x_k - x_\star\|^2 \right] \leq \left( \frac{1 + \gamma^2\delta^2}{(1 + \gamma\mu)^2} \right)^k \|x_0 - x_\star\|^2.$$

□

## 1472 B.7 LOOPLESS STOCHASTIC VARIANCE REDUCED PROXIMAL POINT METHOD (L-SVRP / 1473 SPPM-LGC)

1474 We consider the stochastic optimization problem (1), i.e.,

$$1476 \quad \min_{x \in \mathbb{R}^d} \{f(x) := \mathbb{E}_{\xi \sim \mathcal{D}} [f_\xi(x)]\},$$

1479 and rely on Assumption 1 (differentiability of  $f_\xi$ ) and Assumption 2 ( $\mu$ -strong convexity of  $f_\xi$ ).  
 1480 Recall that this implies strong convexity of  $f$ . Hence  $f$  has a unique minimizer, which we denote  
 1481  $x_\star$ .

1482 Note that SPPM-GC needs to compute  $\nabla f(x_k)$  in iteration  $k$ . This can be very costly or even  
 1483 impossible to do in practice. To make this more clear, consider the problem

$$1485 \quad \min_{x \in \mathbb{R}^d} \left\{ f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x) \right\}$$

1488 as a special case of (1).

- 1491 • **One worker.** Assume we have a single machine solving this problem. Moreover, assume  
 1492 it takes one unit of time to this machine to compute  $\nabla f_i$  for any  $i$ , and  $n$  units of time  
 1493 to compute  $\nabla f$ . If the computation of  $\nabla f$  is the bottleneck (i.e., if it is more expensive  
 1494 than the evaluation of the proximity operator of  $f_i$ ), then an attempt to design a method  
 1495 addressing this bottleneck would be justified.
- 1496 • **Parallel workers.** Assume we have  $n$  workers able to work in parallel. Then  $\nabla f$  can be  
 1497 computed in 1 unit of time if communication among the workers is instantaneous. However,  
 1498 it may still be desirable to avoid having to compute the gradient:
  - 1499 – The server aggregating the  $n$  gradients computed by the workers may have limited  
 1500 capacity, and it may take more time for it to be able to compute the average of a very  
 1501 large number of vectors.
  - 1502 – Some workers may be not available at all times.

1504 These considerations justify the desire to reduce the reliance of SPPM-GC on the computation of  
 1505  $\nabla f$ . The key idea is to compute the gradient only periodically, i.e., to be “lazy” about computing  
 1506 the gradient. In particular, we flip a biased coin, and compute a new gradient if the coin lands the  
 1507 right way. Otherwise, we use the previously computed gradient instead. We shall formalize this in  
 1508 the next section.

1509 We are now ready to present the stochastic proximal point method with lazy gradient correction  
 1510 (SPPM-LGC). In the literature, the method is known under the name loopless stochastic variance  
 1511 reduced proximal (L-SVRP) point method.

**Algorithm 7** Loopless Stochastic Variance Reduced Proximal Point Method (L-SVRP / SPPM-LGC)

---

```

1: Parameters: learning rate  $\gamma > 0$ , starting point  $x_0 \in \mathbb{R}^d$ , starting control vector  $w_0 \in \mathbb{R}^d$ ,
   probability  $p \in (0, 1]$ 
2: for  $k = 0, 1, 2, \dots$  do
3:   Sample  $\xi_k \sim \mathcal{D}$ 
4:   Set  $h_k = \nabla f_{\xi_k}(w_k) - \nabla f(w_k)$ 
5:    $x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k + \gamma h_k)$ 
6:   Set  $w_{k+1} = \begin{cases} x_{k+1} & \text{with probability } p \\ w_k & \text{with probability } 1 - p \end{cases}$ 
7: end for

```

---

Note that as intended, L-SVRP indeed reduces to SPPM-GC when  $w_0 = x_0$  and  $p = 1$ .

**Hiding the prox.** Further, recall that if for some  $x \in \mathbb{R}^d$  and differentiable and convex  $\phi$  we let  $x_+ := \text{prox}_\phi(x)$ , then  $x_+ = x - \nabla\phi(x_+)$ . Therefore, steps 4 and 5 of the method can be written in the equivalent form

$$x_{k+1} = x_k + \gamma h_k - \gamma \nabla f_{\xi_k}(x_{k+1}) = x_k - \gamma (\nabla f_{\xi_k}(x_{k+1}) - \nabla f_{\xi_k}(w_k) + \nabla f(w_k)).$$

**L-SVRP vs L-SVRG.** The name L-SVRP was intentionally coined to resemble the name L-SVRG, which is a method proposed in and studied by Kovalev et al. (2020). This method has the form

$$x_{k+1} = x_k - \gamma (\nabla f_{\xi_k}(x_k) - \nabla f_{\xi_k}(w_k) + \nabla f(w_k)),$$

with Step 6 being identical. That is, the only difference here is that L-SVRP involves  $\nabla f_{\xi_k}(x_{k+1})$  while L-SVRG uses  $\nabla f_{\xi_k}(x_k)$  in the same place. So, ‘‘P’’ in L-SVRP refers to the proximal nature of the term  $\nabla f_{\xi_k}(x_{k+1})$ , while ‘‘G’’ in L-SVRG refers to the gradient nature of the corresponding term  $\nabla f_{\xi_k}(x_k)$ .

**Loopless vs loopy structure.** The word ‘‘loopless’’ refers to the way the control vector  $w_{k+1}$  is updated in Step 6. The alternative to this, used in the famous SVRG method of Johnson and Zhang (2013), is to update  $w_{k+1}$  once every  $m$  iterations, where  $m$  is an appropriately chosen parameter. This change introduces an outer loop into the method, and makes it look a bit more cumbersome. More importantly, the loopless nature of L-SVRG is useful in three ways:

- (i) leads to a somewhat sharper analysis,
- (ii) makes the method easier to analyze, and
- (iii) allows for easier to extensions / modifications.

The last two points are more important than the first one.

**Lemma 7** (L-SVRP). *Suppose Assumption 5 holds with  $\delta > 0$ . Assumption 7 holds for the iterates of L-SVRP (Algorithm 7) with*

$$A_1 = 0, B_1 = \delta^2, C_1 = 0, \quad \text{and} \quad A_2 = p, B_2 = 1 - p, C_2 = 0.$$

*Proof of Lemma 7.* Recall that the iterates of L-SVRP have the form

$$x_{k+1} = \text{prox}_{\gamma f_{\xi_k}}(x_k + \gamma h_k),$$

where  $h_k$  is defined as  $h_k = \nabla f_{\xi_k}(w_k) - \nabla f(w_k)$ , and  $w_k$  is updated in a loopless fashion. Let  $\phi_k = w_k$ . Then

$$\mathbb{E}[h_k | x_k, \phi_k] = \mathbb{E}[\nabla f_{\xi_k}(w_k) - \nabla f(w_k) | x_k, w_k] = 0,$$

and hence (7) holds. If, moreover, Assumption 5 holds, then

$$\mathbb{E} \left[ \|h_k - \nabla f_{\xi_k}(x_\star)\|^2 \middle| x_k, w_k \right] \stackrel{(80)}{\leq} \delta^2 \|w_k - x_\star\|^2,$$

which means that (8) holds with  $A_1 = 0$ ,  $B_1 = \delta^2$  and  $C_1 = 0$  if we let  $\sigma^2(z) := \|z - x_\star\|^2$  (since then  $\sigma_k^2 = \sigma^2(w_k) = \|w_k - x_\star\|^2$ ). On the other hand, from the proof of auxiliary Lemma 12 we know that

$$\mathbb{E} \left[ \|w_{k+1} - x_\star\|^2 \middle| x_{k+1}, w_k \right] \stackrel{(81)}{=} p \|x_{k+1} - x_\star\|^2 + (1-p) \|w_k - x_\star\|^2,$$

which means that (9) holds with  $A_2 = p$ ,  $B_2 = 1 - p$ , and  $C_2 = 0$ . In summary, Assumption 7 holds with

$$A_1 = 0, B_1 = \delta^2, C_1 = 0, \quad \text{and} \quad A_2 = p, B_2 = 1 - p, C_2 = 0.$$

□

**Theorem 7.** *Let Assumption 1 (differentiability), Assumption 2 ( $\mu$ -strong convexity), and Assumption 5 ( $\delta$ -similarity) hold. Choose any  $x_0, w_0 \in \mathbb{R}^d$ . Then for any  $p \in (0, 1]$ ,  $\gamma > 0$ ,  $\alpha > 0$  and all  $k \geq 0$ , we have*

$$\mathbb{E} [\Psi_k] \leq \max \left\{ \frac{1 + \alpha p}{(1 + \gamma\mu)^2}, \frac{1 + \alpha p}{(1 + \gamma\mu)^2} \frac{\gamma^2 \delta^2}{\alpha} + 1 - p \right\}^k \Psi_0, \quad (49)$$

where

$$\Psi_k := \|x_k - x_\star\|^2 + \alpha \|w_k - x_\star\|^2. \quad (50)$$

If  $\delta = 0$ , we have the more precise result

$$\mathbb{E} [\|x_{k+1} - x_\star\|^2] \leq \frac{1}{(1 + \gamma\mu)^2} \mathbb{E} [\|x_k - x_\star\|^2]. \quad (51)$$

Commentary:

1. **Convergence vs divergence.** Clearly, it is possible for the maximum in (49) to *not* be smaller than 1. In this case, the theorem gives a meaningless result. Whether or not the value is smaller than 1 depends on the choice of the parameters  $\alpha$ ,  $p$  and  $\gamma$  in relation to the strong convexity constant  $\mu$ . For example, it's clear that if  $\gamma$  and  $p$  are fixed and  $\alpha$  is too large, the expression  $\frac{1 + \alpha p}{(1 + \gamma\mu)^2}$  might exceed 1, rendering the rate vacuous.
2. **Optimal choice of  $\alpha$ .** Note that  $\alpha \mapsto \frac{1 + \alpha p}{(1 + \gamma\mu)^2}$  is linear and increasing, and  $\alpha \mapsto \frac{1 + \alpha p}{(1 + \gamma\mu)^2} \frac{\gamma^2 \delta^2}{\alpha} + 1 - p$  is convex and decreasing (make sure you understand why!). Moreover, while the first function has a finite value at  $\alpha = 0$ , the second function blows up as  $\alpha$  approaches 0 from the right. This means that the maximum of these two functions will be minimized at the point where the graphs of the two functions intersect, i.e., at  $\alpha$  satisfying

$$\frac{1 + \alpha p}{(1 + \gamma\mu)^2} = \frac{1 + \alpha p}{(1 + \gamma\mu)^2} \frac{\gamma^2 \delta^2}{\alpha} + 1 - p. \quad (52)$$

In the  $p = 1$  case (L-SVRP reduces to SPPM-GC in this regime), the equation simplifies to  $1 = \frac{\gamma^2 \delta^2}{\alpha}$ , i.e., the optimal solution is  $\alpha = \gamma^2 \delta^2$ , and (49) reduces to

$$\mathbb{E} [\Psi_k] \leq \left( \frac{1 + \gamma^2 \delta^2}{(1 + \gamma\mu)^2} \right)^k \Psi_0. \quad (53)$$

This is the same result we obtained in (47) for the SPPM-GC method, up to the choice of the Lyapunov function. However, if we initialize with  $w_0 = x_0$ , then

$$\Psi_k = (1 + \gamma\mu) \|x_k - x_\star\|^2, \quad (54)$$

and plugging this into (53) gives

$$\mathbb{E} [\|x_k - x_\star\|^2] \leq \left( \frac{1 + \gamma^2 \delta^2}{(1 + \gamma\mu)^2} \right)^k \|x_0 - x_\star\|^2, \quad (55)$$

which is exactly (47).

The  $0 < p < 1$  case turns out to be a more cumbersome. After a bit of algebra, one obtains that equation (52) is equivalent to the quadratic equation

$$p\alpha^2 + b\alpha + c = 0,$$

where  $b = 1 - (1 + \gamma\mu)^2(1 - p) - p\gamma^2\delta^2$  and  $c = \gamma^2\delta^2$ . The roots of the quadratic are

$$\alpha = \frac{-b \pm \sqrt{b^2 - 4pc}}{2p}.$$

It seems it is awkward to work with this expression since the resulting rate will become hard to parse and interpret. So, we'll give up on working with perfectly optimal  $\alpha$ . Nevertheless, we will show how to choose some (slightly suboptimal)  $\alpha$  in Corollary 8 which also gives the right complexity result.

Admittedly, it's not easy to understand how good is the rate provided by (49). The following corollary sheds light on what is achievable.

**Corollary 8.** *If we choose  $\alpha = \frac{\gamma\mu}{p}$  and  $\gamma = \frac{p}{p\frac{\delta^2}{\mu} + (1-p)\mu}$ , then for any  $\varepsilon > 0$  we have*

$$k \geq \left(\frac{1}{p} + \frac{\delta^2}{\mu^2}\right) \log\left(\frac{\Psi_0}{\varepsilon}\right) \quad \Rightarrow \quad \mathbb{E}[\Psi_k] \leq \varepsilon. \quad (56)$$

*Proof.* Since  $\alpha = \frac{\gamma\mu}{p}$ , we have

$$A(\gamma) := \frac{1 + \alpha p}{(1 + \gamma\mu)^2} = \frac{1}{1 + \gamma\mu}$$

and

$$B(\gamma) := \frac{1 + \alpha p}{(1 + \gamma\mu)^2} \frac{\gamma^2\delta^2}{\alpha} + 1 - p = \frac{1}{1 + \gamma\mu} \frac{p\gamma\delta^2}{\mu} + 1 - p.$$

Plugging this into (50) leads to

$$\mathbb{E}[\Psi_k] \leq \max\{A(\gamma), B(\gamma)\}^k \Psi_0, \quad (57)$$

We will now select stepsize  $\gamma$  which minimizes

$$\gamma \mapsto \max\{A(\gamma), B(\gamma)\}.$$

Notice that  $\gamma \mapsto A(\gamma)$  is decreasing to zero on  $(0, \infty)$ , with  $A(\gamma)$  blowing up to  $\infty$  as  $\gamma$  approaches zero from the right. Further,  $\gamma \mapsto B(\gamma)$  is increasing in  $(0, \infty)$ . This means that  $\max\{A(\gamma), B(\gamma)\}$  is minimized at the point where the graphs of the two functions intersect, i.e., at  $\gamma$  satisfying  $A(\gamma) = B(\gamma)$ . Direct calculation shows that the solution of this is

$$\gamma = \gamma_\star := \frac{p}{p\frac{\delta^2}{\mu} + (1-p)\mu}, \quad (58)$$

and hence

$$\mathbb{E}[\Psi_k] \leq \max\{A(\gamma_\star), B(\gamma_\star)\}^k \Psi_0 = A(\gamma_\star)^k \Psi_0 = \left(\frac{1}{1 + \gamma_\star\mu}\right)^k \Psi_0 = \left(1 - \frac{\gamma_\star\mu}{1 + \gamma_\star\mu}\right)^k \Psi_0.$$

This implies that

$$k \geq \left(1 + \frac{1}{\gamma_\star\mu}\right) \log\left(\frac{\Psi_0}{\varepsilon}\right) \quad \Rightarrow \quad \mathbb{E}[\Psi_k] \leq \varepsilon. \quad (59)$$

Plugging  $\gamma_\star$  into this iteration complexity result gives

$$1 + \frac{1}{\gamma_\star\mu} \stackrel{(58)}{=} 1 + \frac{p\frac{\delta^2}{\mu} + (1-p)\mu}{p\mu} = \frac{1}{p} + \frac{\delta^2}{\mu^2}.$$

Plugging this back into (59) gives the final result

$$k \geq \left(\frac{1}{p} + \frac{\delta^2}{\mu^2}\right) \log\left(\frac{\Psi_0}{\varepsilon}\right) \quad \Rightarrow \quad \mathbb{E}[\Psi_k] \leq \varepsilon. \quad (60)$$

□

1674 Commentary:

- 1675  
1676 1. **Comparing to SPPM with Gradient Correction (i.e.,  $p = 1$ ).** Recall that L-SVRP re-  
1677 duces to SPPM-GC when  $w_0 = x_0$  and  $p = 1$ . Therefore, one would expect the rates to be  
1678 the same. First, notice that  $w_k = x_k$  for all  $k$ , and as a result, the Lyapunov function (50)  
1679 reduces to

$$1680 \Psi_k = (1 + \gamma\mu)\|x_k - x_\star\|^2, \quad (61)$$

1681 and hence (57) reduces to

$$1682 (1 + \gamma\mu)\mathbb{E} \left[ \|x_k - x_\star\|^2 \right] \stackrel{(61)}{=} \mathbb{E} [\Psi_k] \leq \max \left\{ \left( \frac{1}{1 + \gamma\mu} \right)^k, \left( \frac{1}{1 + \gamma\mu} \frac{\gamma\delta^2}{\mu} \right)^k \right\} \Psi_0.$$

1683 If we choose  $\gamma \leq \frac{\mu}{\delta^2}$ , then the first term in the max dominates, and we get

$$1684 \max \left\{ \left( \frac{1}{1 + \gamma\mu} \right)^k, \left( \frac{1}{1 + \gamma\mu} \frac{\gamma\delta^2}{\mu} \right)^k \right\} \Psi_0 = \left( \frac{1}{1 + \gamma\mu} \right)^k \Psi_0$$

$$1685 \stackrel{(61)}{=} \left( \frac{1}{1 + \gamma\mu} \right)^k (1 + \gamma\mu)\|x_0 - x_\star\|^2.$$

1686 Combining the above observations, we get

$$1687 \mathbb{E} \left[ \|x_k - x_\star\|^2 \right] \leq \left( \frac{1}{1 + \gamma\mu} \right)^k \|x_0 - x_\star\|^2 = \left( 1 - \frac{1}{1 + \frac{1}{\gamma\mu}} \right)^k \|x_0 - x_\star\|^2.$$

1688 It is easy to see that this means that if we choose  $k \geq \left( 1 + \frac{1}{\gamma\mu} \right) \log \left( \frac{\|x_0 - x_\star\|^2}{\varepsilon} \right)$ , then  
1689  $\mathbb{E} \left[ \|x_k - x_\star\|^2 \right] \leq \varepsilon$ . The best rate is obtained for the largest allowed stepsize, i.e., for  
1690  $\gamma = \frac{\mu}{\delta^2}$  (recall that this was the optimal stepsize choice for SPPM-GC established in Ap-  
1691 pendix B.6), in which case we conclude that

$$1692 k \geq \left( 1 + \frac{\delta^2}{\mu^2} \right) \log \left( \frac{\|x_0 - x_\star\|^2}{\varepsilon} \right) \Rightarrow \mathbb{E} \left[ \|x_k - x_\star\|^2 \right] \leq \varepsilon.$$

1693 If more similarity (i.e., smaller  $\delta$ ) or more strong convexity (i.e., larger  $\mu$ ) is present, fewer  
1694 iterations are needed to solve the problem. Note that we get the same result as in (48);  
1695 so, we do not lose anything by doing the analysis in the  $\delta > 0$  case using the Lyapunov  
1696 approach.

- 1697 2. **Comparison to the result of Khaled and Jin (2023).** Choosing  $\alpha = \frac{\gamma\mu}{p}$ ,  $\gamma = \frac{\mu}{2\delta^2}$ ,  $p = \frac{1}{n}$ ,  
1698 we retrieve the convergence guarantees of Khaled and Jin (2023). Indeed, from Corollary 8  
1699 we have that

$$1700 A(\gamma) = \frac{1}{1 + \gamma\mu}, \quad B(\gamma) = \frac{1}{1 + \gamma\mu} \frac{\gamma\delta^2 p}{\mu} + 1 - p.$$

1701 The condition on the stepsize states that  $\frac{\gamma\delta^2}{\mu} \leq \frac{1}{2}$ . It implies that  $B(\gamma) \leq 1 - \frac{p}{2}$ . Let  
1702  $\rho = \min \left\{ \frac{\gamma\mu}{1 + \gamma\mu}, \frac{p}{2} \right\}$ . Clearly,  $\mathbb{E} \left[ \|x_k - x_\star\|^2 \right] \leq \mathbb{E} [\Psi_k]$ . Further, as  $w_0 = x_0$ ,

$$1703 \mathbb{E} [\Psi_0] = \|x_0 - x_\star\|^2 + \frac{\gamma\mu}{p} \|w_0 - x_\star\|^2 = \left( 1 + \frac{\gamma\mu}{p} \right) \|x_0 - x_\star\|^2.$$

1704 For any  $k \geq 0$ , we obtain

$$1705 \mathbb{E} \left[ \|x_k - x_\star\|^2 \right] \leq \left( 1 + \frac{\gamma\mu}{p} \right) (1 - \rho)^k \|x_0 - x_\star\|^2,$$

1706 therefore,

$$1707 \mathbb{E} \left[ \|x_K - x_\star\|^2 \right] \leq \left( 1 + \frac{\gamma\mu}{p} \right) \exp \{-\rho K\} \|x_0 - x_\star\|^2.$$



If we run L-SVRP for

$$K \geq \frac{1}{\rho} \log \left( \frac{\|x_0 - x_\star\|^2 \left(1 + \frac{\gamma\mu}{p}\right)}{\varepsilon} \right).$$

Making the substitutions  $\rho = \min \left\{ \frac{\gamma\mu}{1+\gamma\mu}, \frac{p}{2} \right\}$ ,  $\gamma = \frac{\mu}{2\delta^2}$ ,  $p = \frac{1}{n}$ , we arrive at

$$K \geq 2 \max \left\{ \frac{1}{2} + \frac{\delta^2}{\mu^2}, n \right\} \log \left( \frac{\|x_0 - x_\star\|^2 \left(1 + \frac{\mu^2 n}{2\delta^2}\right)}{\varepsilon} \right),$$

which is even slightly better than the result by Khaled and Jin (2023).

*Proof of Theorem 7.* From Lemma 7 we know that Assumption 7 holds for the iterates of L-SVRP (Algorithm 7) with

$$A_1 = 0, B_1 = \delta^2, C_1 = 0, \quad \text{and} \quad A_2 = p, B_2 = 1 - p, C_2 = 0.$$

From Theorem 1 choosing any  $\alpha > 0$ ,  $\theta = \max \left\{ \frac{1+\alpha p}{(1+\gamma\mu)^2}, \frac{1+\alpha p}{(1+\gamma\mu)^2} \frac{\gamma^2 \delta^2}{\alpha} + 1 - p \right\}$ ,  $\zeta = 0$  (see (15) and (16)), we get

$$\mathbb{E} [\Psi_k] \leq \max \left\{ \frac{1 + \alpha p}{(1 + \gamma\mu)^2}, \frac{1 + \alpha p}{(1 + \gamma\mu)^2} \frac{\gamma^2 \delta^2}{\alpha} + 1 - p \right\}^k \Psi_0,$$

where

$$\Psi_k := \|x_k - x_\star\|^2 + \alpha \|w_k - x_\star\|^2.$$

□

## B.8 POINT SAGA (Point SAGA)

The main motivation is to give one more example of a SPPM method based on the idea of gradient correction which does not need to compute the full/exact gradient of  $f$  in each iteration. We consider another well-known variance-reduced SPPM method called Point SAGA. However, we will revert back to the finite-sum optimization problem

$$\min_{x \in \mathbb{R}^d} \left\{ f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x) \right\}.$$

---

### Algorithm 8 Point SAGA (Point SAGA)

---

- 1: **Parameters:** learning rate  $\gamma > 0$ , starting point  $x_0 \in \mathbb{R}^d$ , starting control vectors  $w_0^i \in \mathbb{R}^d$  for  $i \in [n]$
  - 2: **for**  $k = 0, 1, 2, \dots$  **do**
  - 3:   Sample  $i_k \in \{1, \dots, n\}$  uniformly at random
  - 4:   Set  $h_k = \nabla f_{i_k}(w_k^{i_k}) - \frac{1}{n} \sum_{j=1}^n \nabla f_j(w_k^j)$
  - 5:    $x_{k+1} = \text{prox}_{\gamma f_{i_k}}(x_k + \gamma h_k)$
  - 6:   Set  $w_{k+1}^j = \begin{cases} x_{k+1} & \text{for } j = i_k \\ w_k^j & \text{for } j \neq i_k \end{cases}$
  - 7: **end for**
- 

Commentary:

1. Compared to Algorithm 7, Algorithm 8 uses additional memory to store the table of control vectors  $w_k^i$  or computed gradients  $\nabla f_i(w_k^i)$ .

2. In each iteration, the method update only one ‘‘column’’ in a memory table replacing the old control vector/gradients with the corresponding new one.

The structure of Point SAGA will not allow us perform the analysis under the similarity assumption (Assumption 5). Instead, we will rely on the stronger Assumption 6.

We assume that there exists  $\nu > 0$  such that the inequality

$$\frac{1}{n} \sum_{j=1}^n \left\| \nabla f_j(x^j) - \frac{1}{n} \sum_{i=1}^n \nabla f_i(x^i) - \nabla f_j(x_*) \right\|^2 \leq \nu^2 \frac{1}{n} \sum_{j=1}^n \|x^j - x_*\|^2$$

holds for all  $x^1, \dots, x^n \in \mathbb{R}^d$ .

This inequality can be written in the form

$$\frac{1}{n} \sum_{j=1}^n \|\nabla f_j(x^j) - \nabla f_j(x_*)\|^2 - \left\| \frac{1}{n} \sum_{i=1}^n \nabla f_i(x^i) \right\|^2 \leq \nu^2 \frac{1}{n} \sum_{j=1}^n \|x^j - x_*\|^2, \quad \forall x^j \in \mathbb{R}^d, j \in [n]. \quad (62)$$

Thus, (62) holds, if the following condition is assumed

$$\frac{1}{n} \sum_{j=1}^n \|\nabla f_j(x^j) - \nabla f_j(x_*)\|^2 \leq \nu^2 \frac{1}{n} \sum_{j=1}^n \|x^j - x_*\|^2, \quad \forall x^j \in \mathbb{R}^d, j \in [n]. \quad (63)$$

Moreover, (63) is equivalent to the following condition: for all  $j \in [n]$ , we have

$$\|\nabla f_j(x) - \nabla f_j(x_*)\|^2 \leq \nu^2 \|x - x_*\|^2, \quad \forall x \in \mathbb{R}^d. \quad (64)$$

Finally, (64) holds, if each  $f_j$  is  $\nu$ -smooth, i.e.

$$\|\nabla f_j(x) - \nabla f_j(y)\|^2 \leq \nu^2 \|x - y\|^2, \quad \forall x, y \in \mathbb{R}^d. \quad (65)$$

In summary, we have the following relations between the above conditions:

$$(65) \Rightarrow (64) \equiv (63) \Rightarrow (62) \equiv (6).$$

**Lemma 8.** *Suppose Assumption 6 holds with  $\nu > 0$ . Then Assumption 7 holds for the iterates of Point SAGA (Algorithm 8) with*

$$A_1 = 0, B_1 = \nu^2, C_1 = 0, \quad \text{and} \quad A_2 = \frac{1}{n}, B_2 = \frac{n-1}{n}, C_2 = 0.$$

*Proof of Lemma 8.* Let  $\sigma_k = \frac{1}{n} \sum_{i=1}^n \|w_k^i - x_*\|^2$ ,  $\phi_k = (w_k^1, \dots, w_k^n)$ . Recalling that

$$h_k := \nabla f_{i_k}(w_k^{i_k}) - \frac{1}{n} \sum_{j=1}^n \nabla f_j(w_k^j), \quad (66)$$

we have

$$\mathbb{E}[h_k | x_k, \phi_k] = \mathbb{E} \left[ \frac{1}{n} \sum_{j=1}^n \nabla f_j(w_k^j) - \frac{1}{n} \sum_{j=1}^n \nabla f_j(w_k^j) \middle| x_k, \phi_k \right] = 0. \quad (67)$$

Further,

$$\begin{aligned} \mathbb{E} \left[ \|h_k - \nabla f_{i_k}(x_*)\|^2 \middle| x_k, \phi_k \right] &\stackrel{(66)}{=} \mathbb{E} \left[ \left\| \nabla f_{i_k}(w_k^{i_k}) - \frac{1}{n} \sum_{j=1}^n \nabla f_j(w_k^j) - \nabla f_{i_k}(x_*) \right\|^2 \middle| x_k, \phi_k \right] \\ &= \frac{1}{n} \sum_{i=1}^n \left\| \nabla f_i(w_k^i) - \frac{1}{n} \sum_{j=1}^n \nabla f_j(w_k^j) - \nabla f_i(x_*) \right\|^2 \\ &\stackrel{(6)}{\leq} \frac{\nu^2}{n} \sum_{i=1}^n \|w_k^i - x_*\|^2. \end{aligned}$$

Therefore, we have that  $A_1 = 0$ ,  $B_1 = \nu^2$ ,  $C_1 = 0$ .

$$\begin{aligned} \mathbb{E} [\sigma_{k+1}^2 | x_{k+1}, \phi_k] &= \mathbb{E} \left[ \frac{1}{n} \sum_{i=1}^n \|w_{k+1}^i - x_\star\|^2 \middle| x_{k+1}, \phi_k \right] \\ &= \frac{1}{n} \sum_{i_k=1}^n \left[ \frac{1}{n} \|x_{k+1} - x_\star\|^2 + \frac{1}{n} \sum_{j \neq i_k} \|w_k^j - x_\star\|^2 \right] \\ &= \frac{1}{n} \|x_{k+1} - x_\star\|^2 + \frac{n-1}{n} \sigma_k^2. \end{aligned}$$

Therefore, we have that  $A_2 = \frac{1}{n}$ ,  $B_2 = \frac{n-1}{n}$ ,  $C_2 = 0$ .  $\square$

The convergence of Point SAGA is captured by the following theorem.

**Theorem 9.** *Let Assumption 1, Assumption 2 and Assumption 6 hold. Chose any  $x_0, w_0^1, \dots, w_0^n \in \mathbb{R}^d$ . Then for any  $\gamma > 0$ , and all  $k \geq 0$ , we have*

$$\mathbb{E} [\Psi_k] \leq \max \left\{ \left( \frac{1}{1 + \gamma\mu} \right)^k, \left( \frac{1}{1 + \gamma\mu} \frac{\gamma\nu^2}{\mu n} + 1 - \frac{1}{n} \right)^k \right\} \Psi_0, \quad (68)$$

where

$$\Psi_k := \|x_k - x_\star\|^2 + \gamma\mu \sum_{i=1}^n \|w_k^i - x_\star\|^2. \quad (69)$$

Clearly, it is possible for the maximum in (68) to not be smaller than 1. In this case, the theorem produces the meaningless result. Whether or not the value is smaller than 1 depends on the choice of  $\gamma$  with respect to the strong convexity constant  $\mu$ , the number of individual functions  $n$ , the similarity constant  $\nu$ .

*Proof of Theorem 9.* From Lemma 8 we have that Assumption 7 holds for the iterates of Point SAGA (Algorithm 8) with

$$A_1 = 0, B_1 = \nu^2, C_1 = 0, \quad \text{and} \quad A_2 = \frac{1}{n}, B_2 = \frac{n-1}{n}, C_2 = 0.$$

From Theorem 1, choosing  $\alpha = \gamma\mu n$ ,  $\theta = \max \left\{ \frac{1}{1 + \gamma\mu}, \frac{1}{1 + \gamma\mu} \frac{\gamma\nu^2}{\mu n} + 1 - \frac{1}{n} \right\}$ ,  $\zeta = 0$  (see (15) and (16)), we obtain

$$\mathbb{E} [\Psi_k] \leq \max \left\{ \left( \frac{1}{1 + \gamma\mu} \right)^k, \left( \frac{1}{1 + \gamma\mu} \frac{\gamma\nu^2}{\mu n} + 1 - \frac{1}{n} \right)^k \right\} \Psi_0,$$

where

$$\Psi_k := \|x_k - x_\star\|^2 + \gamma\mu \sum_{i=1}^n \|w_k^i - x_\star\|^2. \quad \square$$

**Corollary 10.** *If we choose  $\gamma = \frac{1}{\frac{\nu^2}{\mu} + (n-1)\mu}$ , then, for any  $\varepsilon > 0$ , we have*

$$k \geq \left( n + \frac{\nu^2}{\mu^2} \right) \log \left( \frac{\Psi_0}{\varepsilon} \right) \quad \Rightarrow \quad \mathbb{E} [\Psi_k] \leq \varepsilon.$$

*Proof of Corollary 10.* Notice that  $A(\gamma) := \frac{1}{1 + \gamma\mu}$  is decreasing for  $\gamma > 0$ , and  $B(\gamma) := \frac{\gamma\nu^2}{1 + \gamma\mu} + 1 - \frac{1}{n}$  is increasing for  $\gamma > 0$ . This means that  $\max \{ \}$  is minimized at  $\gamma := \gamma_\star$  where  $A(\gamma) = B(\gamma)$ . Direct calculation shows that the solution of this is

$$\gamma = \gamma_\star := \frac{1}{\frac{\nu^2}{\mu} + (n-1)\mu},$$

1890 and hence

$$1891 \mathbb{E} [\Psi_k] \leq \max \{A(\gamma_\star), B(\gamma_\star)\}^k \Psi_0 = A(\gamma_\star)^k \Psi_0 = \left(\frac{1}{1 + \gamma_\star \mu}\right)^k \Psi_0 = \left(1 - \frac{\gamma_\star \mu}{1 + \gamma_\star \mu}\right)^k \Psi_0.$$

1894 This implies that

$$1895 k \geq \left(1 + \frac{1}{\gamma_\star \mu}\right) \log \left(\frac{\Psi_0}{\varepsilon}\right) \Rightarrow \mathbb{E} [\Psi_k] \leq \varepsilon.$$

1897 Plugging  $\gamma_\star$  into this iteration complexity result gives

$$1898 1 + \frac{1}{\gamma_\star \mu} = 1 + \frac{\frac{\nu^2}{\mu n} + (1 - \frac{1}{n})\mu}{\frac{\mu}{n}} = n + \frac{\nu^2}{\mu^2}.$$

1901 We obtain the final result

$$1902 k \geq \left(n + \frac{\delta^2}{\mu^2}\right) \log \left(\frac{\Psi_0}{\varepsilon}\right) \Rightarrow \mathbb{E} [\Psi_k] \leq \varepsilon.$$

1905 □

## 1907 C AUXILIARY LEMMA FOR SPPM-LC

1909 Our main result relies on a single lemma only.

1910 **Lemma 9.** *Let Assumption 1, Assumption 2 and Assumption 7 hold. Then*

$$1911 \mathbb{E} [\|x_{k+1} - x_\star\|^2] \leq \frac{(1 + \gamma^2 A_1)}{(1 + \gamma\mu)^2} \mathbb{E} [\|x_k - x_\star\|^2] + \frac{\gamma^2 B_1}{(1 + \gamma\mu)^2} \mathbb{E} [\sigma_k^2] + \frac{\gamma^2 C_1}{(1 + \gamma\mu)^2}. \quad (70)$$

1915 *Proof.* By combining Fact 1 and Fact 2, we get

$$1916 \begin{aligned} \|x_{k+1} - x_\star\|^2 &\stackrel{\text{Fact 1}}{=} \left\| \text{prox}_{\gamma f_{\xi_k}}(x_k + \gamma h_k) - \text{prox}_{\gamma f_{\xi_k}}(x_\star + \gamma \nabla f_{\xi_k}(x_\star)) \right\|^2 \\ &\stackrel{\text{Fact 2}}{\leq} \frac{1}{(1 + \gamma\mu)^2} \|x_k + \gamma h_k - (x_\star + \gamma \nabla f_{\xi_k}(x_\star))\|^2 \\ &= \frac{1}{(1 + \gamma\mu)^2} \|x_k - x_\star + \gamma (h_k - \nabla f_{\xi_k}(x_\star))\|^2. \end{aligned}$$

1923 Thus, we have that

$$1924 \|x_{k+1} - x_\star\|^2 \leq \frac{1}{(1 + \gamma\mu)^2} \left( \|x_k - x_\star\|^2 + 2\gamma \langle h_k - \nabla f_{\xi_k}(x_\star), x_k - x_\star \rangle + \gamma^2 \|h_k - \nabla f_{\xi_k}(x_\star)\|^2 \right).$$

1927 We can use it since it holds irrespective of the choice of  $h_k$ . Taking conditional expectation on both sides, we get

$$1928 \begin{aligned} \mathbb{E} [\|x_{k+1} - x_\star\|^2 | x_k, \phi_k] &\leq \frac{1}{(1 + \gamma\mu)^2} \mathbb{E} [\|x_k - x_\star\|^2 | x_k, \phi_k] \\ &\quad + \frac{2\gamma}{(1 + \gamma\mu)^2} \langle \mathbb{E} [h_k - \nabla f_{\xi_k}(x_\star) | x_k, \phi_k], x_k - x_\star \rangle \\ &\quad + \frac{\gamma^2}{(1 + \gamma\mu)^2} \mathbb{E} [\|h_k - \nabla f_{\xi_k}(x_\star)\|^2 | x_k, \phi_k]. \end{aligned} \quad (71)$$

1936 Note that

$$1937 \mathbb{E} [\|x_k - x_\star\|^2 | x_k, \phi_k] = \|x_k - x_\star\|^2. \quad (72)$$

1939 Further,

$$1940 \begin{aligned} \mathbb{E} [h_k - \nabla f_{\xi_k}(x_\star) | x_k, \phi_k] &= \mathbb{E} [h_k | x_k, \phi_k] - \mathbb{E} [\nabla f_{\xi_k}(x_\star) | x_k, \phi_k] \\ 1941 &= \mathbb{E} [h_k | x_k, \phi_k] - \nabla f(x_\star) \\ 1942 &= \mathbb{E} [h_k | x_k, \phi_k] \\ 1943 &= 0, \end{aligned} \quad (73)$$

where the last equality follows from Assumption 7, relation (7). Relation (8) of Assumption 7 says that

$$\mathbb{E} \left[ \|h_k - \nabla f_{\xi_k}(x_*)\|^2 \middle| x_k, \phi_k \right] \leq A_1 \|x_k - x_*\|^2 + B_1 \sigma_k^2 + C_1. \quad (74)$$

Plugging (74), (73) and (72) into (71), we obtain

$$\mathbb{E} \left[ \|x_{k+1} - x_*\|^2 \middle| x_k, \phi_k \right] \leq \frac{1}{(1 + \gamma\mu)^2} \left( (1 + \gamma^2 A_1) \|x_k - x_*\|^2 + \gamma^2 B_1 \sigma_k^2 + \gamma^2 C_1 \right). \quad (75)$$

It only remains to take expectation on both sides and apply the tower property.  $\square$

## D AUXILIARY LEMMAS

### D.1 SPPM-AS

**Lemma 10.** *Let  $\phi_1, \dots, \phi_m : \mathbb{R}^d \rightarrow \mathbb{R}$  be differentiable functions, with  $\phi_i$  being  $\mu_i$ -strongly convex for all  $i \in [m]$ . Further, let  $w_1, \dots, w_m$  be positive scalars. Then the function  $\phi := \sum_{i=1}^m w_i \phi_i$  is  $\mu$ -strongly convex with  $\mu = \sum_{i=1}^m w_i \mu_i$ .*

*Proof.* By assumption,

$$\phi_i(y) + \langle \nabla \phi_i(y), x - y \rangle + \frac{\mu_i}{2} \|x - y\|^2 \leq \phi_i(x), \quad \forall x, y \in \mathbb{R}^d. \quad (76)$$

This means that

$$\sum_{i=1}^m w_i \left( \phi_i(y) + \langle \nabla \phi_i(y), x - y \rangle + \frac{\mu_i}{2} \|x - y\|^2 \right) \leq \sum_{i=1}^m w_i \phi_i(x), \quad \forall x, y \in \mathbb{R}^d,$$

which is equivalent to

$$\phi(y) + \langle \nabla \phi(y), x - y \rangle + \frac{\sum_{i=1}^m w_i \mu_i}{2} \|x - y\|^2 \leq \phi(x), \quad \forall x, y \in \mathbb{R}^d, \quad (77)$$

So,  $\phi$  is  $\mu$ -strongly convex.  $\square$

### D.2 L-SVRP

**Lemma 11.** *Recalling that*

$$h_k := \nabla f_{\xi_k}(w_k) - \nabla f(w_k), \quad (78)$$

*we can write*

$$\begin{aligned} \mathbb{E} \left[ \langle h_k - \nabla f_{\xi_k}(x_*), x_k - x_* \rangle \middle| x_k, w_k \right] &= \langle \mathbb{E} [h_k - \nabla f_{\xi_k}(x_*) \middle| x_k, w_k], x_k - x_* \rangle \\ &\stackrel{(78)}{=} \langle \mathbb{E} [\nabla f_{\xi_k}(w_k) - \nabla f(w_k) - \nabla f_{\xi_k}(x_*) \middle| x_k, w_k], x_k - x_* \rangle \\ &= \left\langle \underbrace{\nabla f(w_k) - \nabla f(w_k) - \nabla f(x_*)}_{=0}, x_k - x_* \right\rangle \\ &= 0, \end{aligned} \quad (79)$$

and

$$\begin{aligned} \mathbb{E} \left[ \|h_k - \nabla f_{\xi_k}(x_*)\|^2 \middle| x_k, w_k \right] &\stackrel{(78)}{=} \mathbb{E} \left[ \|\nabla f_{\xi_k}(w_k) - \nabla f(w_k) - \nabla f_{\xi_k}(x_*)\|^2 \middle| x_k, w_k \right] \\ &\stackrel{(5)}{\leq} \delta^2 \|w_k - x_*\|^2. \end{aligned} \quad (80)$$

**Lemma 12.** *Observe that by the way  $w_{k+1}$  is defined, we have*

$$\mathbb{E} \left[ \|w_{k+1} - x_*\|^2 \middle| x_{k+1}, w_k \right] = p \|x_{k+1} - x_*\|^2 + (1 - p) \|w_k - x_*\|^2. \quad (81)$$

*Taking expectation again, and applying the tower property of expectation, we get*

$$\begin{aligned} \mathbb{E} \left[ \|w_{k+1} - x_*\|^2 \right] &= \mathbb{E} \left[ \mathbb{E} \left[ \|w_{k+1} - x_*\|^2 \middle| x_{k+1}, w_k \right] \right] \\ &\stackrel{(81)}{=} p \mathbb{E} \left[ \|x_{k+1} - x_*\|^2 \right] + (1 - p) \mathbb{E} \left[ \|w_k - x_*\|^2 \right], \end{aligned} \quad (82)$$

*which is what we set out to prove.*

## E AUXILIARY FACTS

**Fact 1** (Every point is a fixed point). *Let  $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$  be a differentiable convex function. Then*

$$\text{prox}_{\gamma\phi}(x + \gamma\nabla\phi(x)) = x, \quad \forall\gamma > 0, \quad \forall x \in \mathbb{R}^d.$$

*In particular, if  $x_*$  is a minimizer of  $\phi$ , then  $\text{prox}_{\gamma\phi}(x_*) = x_*$ .*

*Proof.* Pick any  $x \in \mathbb{R}^d$  and  $\gamma > 0$ . Evaluating the proximity operator at

$$y := x + \gamma\nabla\phi(x) \tag{83}$$

gives

$$\text{prox}_{\gamma\phi}(y) = \arg \min_{x' \in \mathbb{R}^d} \left( \phi(x') + \frac{1}{2\gamma} \|x' - y\|^2 \right).$$

This is a strongly convex minimization problem, and hence the (necessarily unique) minimizer

$$x' := \text{prox}_{\gamma\phi}(y) \tag{84}$$

of this problem satisfies the first-order optimality condition

$$\nabla\phi(x') + \frac{1}{\gamma}(x' - y) = 0.$$

Note that  $x' = x$  satisfies this equation, and hence

$$x \stackrel{(84)}{=} \text{prox}_{\gamma\phi}(y) \stackrel{(83)}{=} \text{prox}_{\gamma\phi}(x + \gamma\nabla\phi(x)).$$

□

The next statement is (Khaled and Jin, 2023, Fact 4)

**Fact 2** (Contractivity of the prox). *If  $\phi$  is differentiable and  $\mu$ -strongly convex, then for all  $\gamma > 0$  and for any  $x, y \in \mathbb{R}^d$  we have*

$$\|\text{prox}_{\gamma\phi}(x) - \text{prox}_{\gamma\phi}(y)\|^2 \leq \frac{1}{(1 + \gamma\mu)^2} \|x - y\|^2.$$

*Proof.* This lemma can be seen as a tighter version of (Mishchenko et al., 2022, Lemma 5) though our proof technique is different. Note that  $p(x) = \text{prox}_{\gamma h}(x)$  satisfies  $\gamma\nabla h(p(x)) + [p(x) - x] = 0$ , or equivalently  $p(x) = x - \gamma\nabla h(p(x))$ . Using this we have

$$\begin{aligned} \|p(x) - p(y)\|^2 &= \|[x - \gamma\nabla h(p(x))] - [y - \gamma\nabla h(p(y))]\|^2 \\ &= \|[x - y] - \gamma[\nabla h(p(x)) - \nabla h(p(y))]\|^2 \\ &= \|x - y\|^2 + \gamma^2 \|\nabla h(p(x)) - \nabla h(p(y))\|^2 - 2\gamma \langle x - y, \nabla h(p(x)) - \nabla h(p(y)) \rangle. \end{aligned} \tag{85}$$

Now note that

$$\begin{aligned} \langle x - y, \nabla h(p(x)) - \nabla h(p(y)) \rangle &= \langle p(x) + \gamma\nabla h(p(x)) - [p(y) + \gamma\nabla h(p(y))], \nabla h(p(x)) - \nabla h(p(y)) \rangle \\ &= \langle p(x) - p(y), \nabla h(p(x)) - \nabla h(p(y)) \rangle + \gamma \|\nabla h(p(x)) - \nabla h(p(y))\|^2. \end{aligned} \tag{86}$$

Combining Equations (85) and (86) we get

$$\begin{aligned} \|p(x) - p(y)\|^2 &= \|x - y\|^2 + \gamma^2 \|\nabla h(p(x)) - \nabla h(p(y))\|^2 - 2\gamma \langle p(x) - p(y), \nabla h(p(x)) - \nabla h(p(y)) \rangle \\ &\quad - 2\gamma^2 \|\nabla h(p(x)) - \nabla h(p(y))\|^2 \\ &= \|x - y\|^2 - \gamma^2 \|\nabla h(p(x)) - \nabla h(p(y))\|^2 - 2\gamma \langle p(x) - p(y), \nabla h(p(x)) - \nabla h(p(y)) \rangle. \end{aligned} \tag{87}$$

2052 Let  $D_h(u, v) = h(u) - h(v) - \langle \nabla h(v), u - v \rangle$  be the Bregman divergence associated with  $h$  at  $u, v$ .  
 2053 It is easy to show that

$$2054 \quad \langle u - v, \nabla h(u) - \nabla h(v) \rangle = D_h(u, v) + D_h(v, u).$$

2056 This is a special case of the three-point identity (Chen and Teboulle, 1993, Lemma 3.1). Using this  
 2057 with  $u = p(x)$  and  $v = p(y)$  and plugging back into (87) we get

$$2058 \quad \|p(x) - p(y)\|^2 = \|x - y\|^2 - \gamma^2 \|\nabla h(p(x)) - \nabla h(p(y))\|^2 - 2\gamma [D_h(p(x), p(y)) + D_h(p(y), p(x))].$$

2060 Note that because  $h$  is strongly convex, we have that  $D_h(p(y), p(x)) \geq \frac{\mu}{2} \|p(y) - p(x)\|^2$  and  
 2061  $D_h(p(x), p(y)) \geq \frac{\mu}{2} \|p(y) - p(x)\|^2$ , hence

$$2062 \quad \|p(x) - p(y)\|^2 \leq \|x - y\|^2 - \gamma^2 \|\nabla h(p(x)) - \nabla h(p(y))\|^2 - 2\gamma\mu \|p(x) - p(y)\|^2. \quad (88)$$

2065 Strong convexity implies that for any two points  $u, v$

$$2066 \quad \|\nabla h(u) - \nabla h(v)\|^2 \geq \mu^2 \|u - v\|^2,$$

2068 see (Nesterov, 2018, Theorem 2.1.10) for a proof. Using this in Equation (88) with  $u = p(x)$  and  
 2069  $v = p(y)$  yields

$$2070 \quad \|p(x) - p(y)\|^2 \leq \|x - y\|^2 - \gamma^2 \mu^2 \|p(x) - p(y)\|^2 - 2\gamma\mu \|p(x) - p(y)\|^2.$$

2072 Rearranging gives

$$2073 \quad [1 + \gamma^2 \mu^2 + 2\gamma\mu] \|p(x) - p(y)\|^2 \leq \|x - y\|^2.$$

2075 It remains to notice that  $(1 + \gamma\mu)^2 = 1 + \gamma^2 \mu^2 + 2\gamma\mu$ . □

2077 **Fact 3** (Recurrence). Assume that a sequence  $\{s_k\}_{k \geq 0}$  of positive real numbers for all  $k \geq 0$   
 2078 satisfies

$$2079 \quad s_{k+1} \leq as_k + b,$$

2081 where  $0 < a < 1$  and  $b \geq 0$ . Then the sequence for all  $k \geq 0$  satisfies

$$2082 \quad s_k \leq a^k s_0 + b \min \left\{ k, \frac{1}{1-a} \right\}. \quad (89)$$

2085 *Proof.* Unrolling the recurrence, we get

$$2086 \quad s_k \leq as_{k-1} + b \leq a(as_{k-2} + b) + b \leq \dots \leq a^k s_0 + b \sum_{i=0}^{k-1} a^i. \quad (90)$$

2090 We can now bound the sum  $\sum_{i=0}^{k-1} a^i$  in two different ways. First, since  $a < 1$ , we get the estimate

$$2091 \quad \sum_{i=0}^{k-1} a^i \leq \sum_{i=0}^{k-1} 1 = k. \quad (91)$$

2095 Second, we sum a geometric series

$$2096 \quad \sum_{i=0}^{k-1} a^i \leq \sum_{i=0}^{\infty} a^i = \frac{1}{1-a}. \quad (92)$$

2100 Note that either of these bounds can be better. So, we apply the best of these bounds. Substituting  
 2101 Equations (91) and (92) into (90) gives (89). □

2102  
2103  
2104  
2105