
Optimal Extragradient-Based Algorithms for Stochastic Variational Inequalities with Separable Structure

Huizhuo Yuan[◇] Chris Junchi Li[†] Gauthier Gidel[‡]

Michael I. Jordan^{†,□} Quanquan Gu[◇] Simon S. Du^{*}

[◇] Department of Computer Science, University of California, Los Angeles
{hzyuan, qgu}@cs.ucla.edu

[†] Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
{junchili, jordan}@cs.berkeley.edu

[‡] DIRO, Université de Montréal and Mila
gauthier.gidel@umontreal.ca

[□] Department of Statistics, University of California, Berkeley

^{*} Paul G. Allen School of Computer Science and Engineering, University of Washington
ssdu@cs.washington.edu

Abstract

We consider the problem of solving stochastic monotone variational inequalities with a separable structure using a stochastic first-order oracle. Building on standard extragradient for variational inequalities we propose a novel algorithm—stochastic *accelerated gradient-extragradient* (AG-EG)—for strongly monotone variational inequalities (VIs). Our approach combines the strengths of extragradient and Nesterov acceleration. By showing that its iterates remain in a bounded domain and applying scheduled restarting, we prove that AG-EG has an optimal convergence rate for strongly monotone VIs. Furthermore, when specializing to the particular case of bilinearly coupled strongly-convex-strongly-concave saddle-point problems, including bilinear games, our algorithm achieves fine-grained convergence rates that match the respective lower bounds, with the stochasticity being characterized by an additive statistical error term that is optimal up to a constant prefactor.

1 Introduction

The variational inequality (VI) problem plays a central role in a wide range of optimization problems with convex structure, including convex minimization, saddle-point problems, and games [Facchinei and Pang, 2003, Nemirovski, 2004, Nemirovski et al., 2009, Juditsky et al., 2011, Jordan et al., 2023]. A general VI problem aims to find a solution $z^* \in \mathcal{Z}$ that satisfies:

$$\langle \mathcal{W}(z^*), z^* - z \rangle \leq 0, \quad \forall z \in \mathcal{Z}, \quad (1)$$

where \mathcal{Z} is a finite-dimensional closed and convex feasible set and $\mathcal{W}(\cdot)$ is a monotone operator in the following form:

$$\mathcal{W}(z) = \nabla \mathcal{F}(z) + \mathcal{H}(z) + J'(z) \equiv \mathbb{E}_\xi[\nabla \tilde{\mathcal{F}}(z; \xi)] + \mathbb{E}_\zeta[\tilde{\mathcal{H}}(z; \zeta)] + J'(z), \quad (2)$$

where \mathcal{F} is continuously differentiable with L -Lipschitz continuous gradient and is μ -strongly convex, \mathcal{H} is an M -Lipschitz monotone operator, $J' \in \partial J$ is the subgradient of a simple and convex function, ξ and ζ are drawn from distributions \mathcal{D}_ξ and \mathcal{D}_ζ , respectively. This formulation captures a separable structure in which \mathcal{H} usually models the competing forces in a system, and J models

a nonsmooth factor. In addition, we consider the stochastic setting where we can only access $\nabla\mathcal{F}$ and \mathcal{H} through their unbiased estimators $\nabla\tilde{\mathcal{F}}(\mathbf{z}; \xi)$ and $\tilde{\mathcal{H}}(\mathbf{z}; \zeta)$ respectively.

A notable instance of the VI problem (1) with separable structure (2) is the widely studied *bilinearly coupled strongly-convex-strongly-concave saddle-point problem*:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \max_{\mathbf{y} \in \mathbb{R}^m} \mathcal{F}(\mathbf{x}, \mathbf{y}) = F(\mathbf{x}) + H(\mathbf{x}, \mathbf{y}) - G(\mathbf{y}) \equiv \mathbb{E}_\xi [f(\mathbf{x}; \xi)] + \mathbb{E}_\zeta [h(\mathbf{x}, \mathbf{y}; \zeta)] - \mathbb{E}_\xi [g(\mathbf{y}; \xi)], \quad (3)$$

where $H(\mathbf{x}, \mathbf{y}) \equiv \mathbf{x}^\top \mathbf{B} \mathbf{y} - \mathbf{x}^\top \mathbf{u}_x + \mathbf{u}_y^\top \mathbf{y}$ is the bilinear coupling function with the coupling matrix $\mathbf{B} \in \mathbb{R}^{n \times m}$. Note that (3) is a special instance of (1) when taking $\mathcal{F}(\mathbf{z}) = F(\mathbf{x}) + G(\mathbf{y})$, $\mathcal{H}(\mathbf{z}) = [\nabla_x H(\mathbf{x}, \mathbf{y}); -\nabla_y H(\mathbf{x}, \mathbf{y})]$ and $J = 0$. In addition to a wide range of applications in economics, problems of form (3) are becoming increasingly important in machine learning. For instance, (3) appears in reinforcement learning, differentiable games, regularized empirical risk minimization, and robust optimization formulations. It can also be seen as a local approximation of the nonconvex-nonconcave minimax games—e.g., the generative adversarial network (GAN) [Goodfellow et al., 2020]—around a local Nash equilibrium [Mescheder et al., 2017; Nagarajan and Kolter, 2017].

In this paper, we aim to improve the efficiency of solving (1) by utilizing the structural information of the monotone operator in (2). More specifically, we consider the case when \mathcal{F} is strongly monotone, or zero. Although optimal convergence results have been obtained for the monotone VI problem (1) [Chen et al., 2017] as well as the special case of convex-concave saddle-point problem with bilinear coupling (3) [Chen et al., 2014], it remains open how to design an optimal algorithm for the strongly monotone VI problem. Notably, for the special case (3) when F and/or G are strongly convex, several concurrent works have independently obtained the optimal convergence rates [Kovalev et al., 2022; Thekumparampil et al., 2022; Jin et al., 2022; Metelev et al., 2022; Li et al., 2022b]. On the other hand, when both F and G are zero, optimal convergence results have been obtained by [Li et al., 2022a] and the accelerated-gradient optimistic gradient approach [Li et al., 2022b]. We defer a more complete overview of related work to the appendix.

1.1 Main Contributions

We start with the strongly monotone VI problem in an unbounded feasible set, extending the scope of recent work such as [Jin et al., 2022] and going beyond earlier studies that focus on nonstrongly monotone VIs in a bounded feasible set [Juditsky et al., 2011; Chen et al., 2017]. We propose a class of algorithms named stochastic *accelerated gradient-extragradient* (AG-EG), which combine Nesterov acceleration with the extragradient method. By employing either a strong-convexity shifting technique or a scheduled restarting scheme, our algorithm achieves convergence rates that match the lower bounds for the general strongly monotone VI problem (1), the special SC-SC saddle-point problem (3), and bilinear games, in both deterministic and stochastic settings, thus providing a unified optimal solution. In sharp contrast to the accelerated mirror-prox (AMP) algorithm proposed by [Chen et al., 2017; Jordan et al., 2023], our analysis does not rely on the boundedness of the feasible set \mathcal{Z} , which makes our algorithm projection-free. We also extend our algorithm to VIs with bounded feasible set and/or nondifferentiable convex regularization through proximal mapping. We summarize our contributions as follows:

- (1) We present a direct approach for separable strongly monotone VIs, where the iteration complexity lower bound due to [Zhang et al., 2022] is matched as $\mathcal{O}\left(\sqrt{\frac{L}{\mu}} + \frac{M}{\mu} + \frac{\sigma^2}{\mu^2 \varepsilon^2}\right) \log\left(\frac{L}{\mu} \frac{1}{\varepsilon}\right)$, which admits a sharp near-unity coefficient [§2.3, Theorem 2.3]. Here σ^2 is the weighted, uniform variance bound on the stochastic gradient and stochastic operator.
- (2) We also present a stochastic AG-EG algorithm equipped with scheduled restarting, which achieves the sharpest possible iteration complexity of $\mathcal{O}\left(\left(\sqrt{\frac{L}{\mu}} + \frac{M}{\mu}\right) \log\left(\frac{1}{\varepsilon}\right) + \frac{\sigma^2}{\mu^2 \varepsilon^2}\right)$ for finding an ε -optimal point. The deterministic part matches the complexity lower bound in [Zhang et al., 2022], while the stochastic part matches the optimal statistical error.

¹VIs in an unbounded feasible set is more difficult to solve because existing algorithms and analyses crucially rely on the boundedness of the feasible set.

When specializing the VI problem to bilinearly coupled SC-SC saddle-point problems, our results have the following implications:

Strongly-convex-strongly-concave (SC-SC) Saddle-Point Problem. For the class of SC-SC saddle-point problems, the stochastic AG-EG descent-ascent Algorithm [1](#), equipped with scaling reduction, achieves an iteration complexity of

$$\mathcal{O}\left(\left(\sqrt{\frac{L_F}{\mu_F}} \vee \frac{L_G}{\mu_G} + \sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\mu_F \mu_G}}\right) \log\left(\frac{1}{\varepsilon}\right) + \frac{\sigma^2}{\mu_F^2 \varepsilon^2}\right), \quad (4)$$

where $F : \mathbb{R}^n \rightarrow \mathbb{R}$ is L_F -smooth and μ_F -strongly convex, $G : \mathbb{R}^m \rightarrow \mathbb{R}$ is L_G -smooth and μ_G -strongly convex. When the optimization problem is deterministic, the complexity upper bound matches the lower bound of [Zhang et al. \[2022\]](#) [[§3.1](#), [Corollary 2.8](#)].

Bilinear Games. For bilinear games ($\nabla f(\mathbf{x}; \xi) = \mathbf{0}$ and $\nabla g(\mathbf{y}; \xi) = \mathbf{0}$ almost surely), Algorithm [1](#), equipped with scheduled restarting achieves an iteration complexity of

$$\mathcal{O}\left(\sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top)}} \log\left(\frac{\sqrt[4]{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top)\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}}{\sigma_{\text{Bil}}}\right) + \frac{\sigma_{\text{Bil}}^2}{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top)\varepsilon^2}\right), \quad (5)$$

where σ_{Bil}^2 is the variance of the stochastic gradient on the bilinear coupling term. When there is no randomness, this complexity result reduces to $\mathcal{O}\left(\sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top)}} \log\left(\frac{1}{\varepsilon}\right)\right)$ for bilinear games, matching the lower bound of [Ibrahim et al. \[2020\]](#) [[§3.2](#), [Corollary 3.3](#)]².

Organization. The rest of this paper is organized as follows. Section [2](#) proposes the Accelerated Gradient-Extragradient Descent-Ascent algorithm for strongly monotone VIs, showing that it achieves an accelerated convergence rate, and extending to VIs with bounded domains with proximal operator. Section [3](#) discusses two specific instances of saddle-point problems, where our proposed AG-EG algorithm has a convergence rate that matches the corresponding lower bounds. Finally, Section [4](#) summarizes our results and suggests future directions.

Notation. Let $\lambda_{\max}(\mathbf{M})$ (resp. $\lambda_{\min}(\mathbf{M})$) be the largest (resp. smallest) eigenvalue of a real symmetric matrix \mathbf{M} . Let $a \vee b \equiv \max(a, b)$ (resp. $a \wedge b \equiv \min(a, b)$) denote the maximum (resp. minimum) value of two reals a, b . For two nonnegative real sequences (a_n) and (b_n) , we write $a_n = \mathcal{O}(b_n)$ or $a_n \lesssim b_n$ (resp. $a_n = \Omega(b_n)$ or $a_n \gtrsim b_n$) to denote $a_n \leq Cb_n$ (resp. $a_n \geq Cb_n$) for all $n \geq 1$ for a positive, numerical constant C , and let $a_n \asymp b_n$ if both $a_n \lesssim b_n$ and $a_n \gtrsim b_n$ hold. We also let $a_n = \tilde{\mathcal{O}}(b_n)$ denote $a_n \leq Cb_n$ where C hides a polylogarithmic factor in problem-dependent constants. We let $[\mathbf{x}; \mathbf{y}] \in \mathbb{R}^{n+m}$ concatenate two vectors $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{y} \in \mathbb{R}^m$. Finally for two real symmetric matrices \mathbf{A} and \mathbf{B} , we denote $\mathbf{A} \preceq \mathbf{B}$ (resp. $\mathbf{A} \succeq \mathbf{B}$) when $\mathbf{v}^\top(\mathbf{A} - \mathbf{B})\mathbf{v} \leq 0$ (resp. $\mathbf{v}^\top(\mathbf{A} - \mathbf{B})\mathbf{v} \geq 0$) holds for all vectors \mathbf{v} .

2 Accelerated Gradient-Extragradient Descent-Ascent Algorithm

In this section, we focus on accelerating the extragradient algorithm for the strongly monotone VI problem in [\(1\)](#) with separable structure [\(2\)](#). Our algorithm design draws inspiration from the work of [Chen et al. \[2017\]](#) on the stochastic Accelerated MirrorProx (AMP) algorithm for nonstrongly monotone VIs. The AMP algorithm applies Nesterov-type acceleration on top of the mirror-prox method [[Korpelevich, 1976](#), [Nemirovski, 2004](#)] and attains the optimal iteration complexity of $\mathcal{O}\left(\sqrt{\frac{L}{\varepsilon}} + \frac{M}{\varepsilon}\right)$. However, the big-O notation hides the diameter of the feasible set, and the existing theory for the AMP algorithm can only deal with VIs with bounded domain. Our algorithm not only achieves the optimal convergence rates for the strongly monotone VI problem with separable structure but we also remove the dependency on the diameter of the feasible set. Therefore, our algorithm can deal with VIs with unbounded domains.

²For the function class of bilinear games, we assume that $n = m$ where \mathbf{B} is a nonsingular square matrix, so that $\lambda_{\min}(\mathbf{B}\mathbf{B}^\top) > 0$ and the complexity makes sense. See [§3.2](#) for more on this.

Throughout §2, we maintain conceptual simplicity by presenting all our algorithm designs in the deterministic setting, while presenting the convergence results in the more general stochastic setting. These results can be easily reduced to the deterministic setting when the stochastic noise vanishes.

2.1 Setting and Assumptions

In this section, we formally introduce our assumptions. We first state the smoothness and monotonicity assumptions that we impose on \mathcal{F} and \mathcal{H} .

Assumption 2.1 (Monotonicity, strong convexity and smoothness) *We assume that function $\mathcal{F}(\cdot)$ is continuously differentiable with L -Lipschitz continuous gradient and is μ -strongly convex. That is, for any $\mathbf{z}, \mathbf{z}' \in \mathcal{Z}$,*

$$\frac{\mu}{2} \|\mathbf{z} - \mathbf{z}'\|^2 \leq \mathcal{F}(\mathbf{z}) - \mathcal{F}(\mathbf{z}') - \nabla \mathcal{F}(\mathbf{z}')^\top (\mathbf{z} - \mathbf{z}') \leq \frac{L}{2} \|\mathbf{z} - \mathbf{z}'\|^2.$$

Furthermore, operator $\mathcal{H}(\cdot)$ is monotone and M -Lipschitz in the sense that for any $\mathbf{z}, \mathbf{z}' \in \mathcal{Z}$,

$$\langle \mathcal{H}(\mathbf{z}) - \mathcal{H}(\mathbf{z}'), \mathbf{z} - \mathbf{z}' \rangle \geq 0, \quad \|\mathcal{H}(\mathbf{z}) - \mathcal{H}(\mathbf{z}')\| \leq M \|\mathbf{z} - \mathbf{z}'\|.$$

Second, we impose assumptions on the noise variance.

Assumption 2.2 (Unbiased gradients and variance bounds) *We assume that $\mathbf{z} \in \mathcal{Z}$, samples $\xi \sim \mathcal{D}_\xi$ and $\zeta \sim \mathcal{D}_\zeta$ are drawn from given distributions such that the following conditions hold: $\mathbb{E}_\xi[\nabla \tilde{\mathcal{F}}(\mathbf{z}; \xi)] = \nabla \mathcal{F}(\mathbf{z})$, $\mathbb{E}_\zeta[\tilde{\mathcal{H}}(\mathbf{z}; \zeta)] = \mathcal{H}(\mathbf{z})$, and*

$$\mathbb{E}_\xi \left[\|\nabla \tilde{\mathcal{F}}(\mathbf{z}; \xi) - \nabla \mathcal{F}(\mathbf{z})\|^2 \right] \leq \sigma_{\text{Str}}^2, \quad \mathbb{E}_\zeta \left[\|\tilde{\mathcal{H}}(\mathbf{z}; \zeta) - \mathcal{H}(\mathbf{z})\|^2 \right] \leq \sigma_{\text{Bil}}^2. \quad (6)$$

For all results in this work, we suppose that Assumptions 2.1 and 2.2 hold with appropriate parameter settings. Given a desired accuracy $\varepsilon > 0$, our goal is to find an ε -optimal point defined as:

Definition 2.1 (ε -Optimal point) *A point $\mathbf{z} \in \mathcal{Z}$ is called an ε -optimal point for the VI problem in (1) if $\|\mathbf{z} - \mathbf{z}^*\| \leq \varepsilon$.*

2.2 The ExtraGradient (EG) Algorithm

We first consider the case where \mathcal{Z} is the entire space \mathbb{R}^n and the objective is smooth ($J = 0$). The extragradient (EG) algorithm, introduced by Korpelevich [1976], is designed to address cyclic behavior in saddle-point problems by introducing an extrapolated point for gradient evaluation. In the context of VI problems (1), let \mathbf{z}_t represents the t -th iterate of the EG algorithm. The update rule of EG is as follows:

$$\mathbf{z}_{t+1} = \mathbf{z}_t - \eta \mathcal{W}(\mathbf{z}_t - \eta \mathcal{W}(\mathbf{z}_t)), \quad (7)$$

where $\eta > 0$ is the step size. For a L -smooth and μ -strongly monotone operator \mathcal{W} , Tseng [1995], Mokhtari et al. [2020], Gidel et al. [2019a] have shown that the EG algorithm achieves an iteration complexity of $\mathcal{O}(\kappa \log(1/\varepsilon))$, where $\kappa = L/\mu$ denotes the condition number of the problem.

2.3 Accelerating the ExtraGradient Algorithm, Direct Approach

The convergence rate of the EG algorithm is far from optimal for the strongly monotone VI problem in (1) with separable structure (2). Firstly, the update rule in (7) takes \mathcal{W} as a whole without utilizing the separable structure. This prevents us from exploiting the properties of $\nabla \mathcal{F}$. Secondly, in the case of bilinear games, the established lower bound for EG is $\Omega(\sqrt{\kappa} \log(1/\varepsilon))$ rather than $\Omega(\kappa \log(1/\varepsilon))$. This discrepancy highlights the potential for accelerating the EG algorithm in various directions. We first rewrite the EG update rule in (7) as follows:

$$\begin{aligned} \mathbf{z}_{t-\frac{1}{2}} &= \mathbf{z}_{t-1} - \eta \mathcal{W}(\mathbf{z}_{t-1}) = \mathbf{z}_{t-1} - \eta (\mathcal{H}(\mathbf{z}_{t-1}) + \nabla \mathcal{F}(\mathbf{z}_{t-1})), \\ \mathbf{z}_t &= \mathbf{z}_{t-1} - \eta \mathcal{W}(\mathbf{z}_{t-\frac{1}{2}}) = \mathbf{z}_{t-1} - \eta (\mathcal{H}(\mathbf{z}_{t-\frac{1}{2}}) + \nabla \mathcal{F}(\mathbf{z}_{t-\frac{1}{2}})). \end{aligned} \quad (8)$$

To accelerate the process based on $\nabla \mathcal{F}$, we consider Nesterov's second acceleration scheme on minimizing a single convex function \mathcal{F} [Tseng, 2008; Lan and Zhou, 2018; Lin et al., 2020c]:

$$\mathbf{z}_{t-1}^{\text{md}} = (1 - \alpha_t) \mathbf{z}_{t-1}^{\text{ag}} + \alpha_t \mathbf{z}_{t-1}, \quad \mathbf{z}_t = \mathbf{z}_{t-1} - \frac{\eta}{\alpha_t} \nabla \mathcal{F}(\mathbf{z}_{t-1}^{\text{md}}), \quad \mathbf{z}_t^{\text{ag}} = (1 - \alpha_t) \mathbf{z}_{t-1}^{\text{ag}} + \alpha_t \mathbf{z}_t, \quad (9)$$

where α_t is the extrapolation step size in the standard three-line Nesterov scheme. Here we adopt the notation z^{md} and z^{ag} to indicate the middle point and the aggregated point [Chen et al. 2017], respectively. Next, to achieve acceleration, we replace the gradient of $\nabla \mathcal{F}$ evaluated at both z_{t-1} and $z_{t-\frac{1}{2}}$ in (8) by the gradient evaluated at the extrapolated point z_{t-1}^{md} in (9). Furthermore, we shift the index of z^{ag} by $\frac{1}{2}$ to indicate the use of $z_{t-\frac{1}{2}}$ instead of z_t in the z^{ag} update in (9). In addition, we take into account the μ -strong convexity of \mathcal{F} and shift the gradient of the strongly convex part $\nabla_z \left[\frac{\mu}{2} \|z - z_0\|^2 \right] = \mu(z - z_0)$ from $\nabla \mathcal{F}(z)$ to $\mathcal{H}(z)$ as $\mathcal{W}(z) = (\nabla \mathcal{F}(z) - \mu(z - z_0)) + (\mathcal{H}(z) + \mu(z - z_0))$, we obtain the following update rule for a direct version of an accelerated EG algorithm (different step size schemes for η_t are required for different algorithmic designs):

$$\begin{cases} z_{t-1}^{\text{md}} = (1 - \alpha_t)z_{t-\frac{3}{2}}^{\text{ag}} + \alpha_t z_{t-1}, \\ z_{t-\frac{1}{2}} = z_{t-1} - \eta_t (\mathcal{H}(z_{t-1}) + \nabla \mathcal{F}(z_{t-1}^{\text{md}}) - \mu(z_{t-1}^{\text{md}} - z_{t-1})), \\ z_t = z_{t-1} - \eta_t (\mathcal{H}(z_{t-\frac{1}{2}}) + \nabla \mathcal{F}(z_{t-1}^{\text{md}}) - \mu(z_{t-1}^{\text{md}} - z_{t-\frac{1}{2}})), \\ z_{t-\frac{1}{2}}^{\text{ag}} = (1 - \alpha_t)z_{t-\frac{3}{2}}^{\text{ag}} + \alpha_t z_{t-\frac{1}{2}}. \end{cases} \quad (10)$$

We call the algorithm in (10) the *accelerated gradient-extragradient, direct approach* (AG-EG-Direct), and postpone its full description to Algorithm 2 in §C.1. The final output of the direct approach is z_T after T iterates. The following theorem records the convergence rate and iteration complexity of AG-EG (direct approach).

Theorem 2.3 (Convergence of stochastic AG-EG, direct approach) *Suppose Assumptions 2.1 and 2.2 hold. Fix any $r \in (0, 1)$, $\beta \in (0, \infty)$, let $\kappa_\beta = \frac{L}{\mu} + \frac{(1+\beta)M^2}{\mu^2}$ and set the step size upper bound $\bar{\alpha} \equiv \frac{r}{1+\sqrt{1+r\kappa_\beta}}$. For any sequence of step sizes $\alpha_t \in (0, \bar{\alpha}]$ and $\eta_t = \frac{\alpha_t}{\mu}$, the iterates of stochastic AG-EG (direct approach) satisfy that for all $t = 1, \dots, T$, we have*

$$\mathbb{E} \|z_t - z^*\|^2 \leq \|z_0 - z^*\|^2 \left(\frac{L}{\mu} + 1 \right) \prod_{s=1}^t (1 - \alpha_s) + \frac{3\sigma^2}{\mu^2} \sum_{s=1}^t \alpha_s^2 \prod_{\tau=s+1}^t (1 - \alpha_\tau), \quad (11)$$

where we define $\sigma = \frac{1}{\sqrt{3}} \sqrt{\frac{1}{1-r} \sigma_{\text{Str}}^2 + (2 + \frac{1}{\beta}) \sigma_{\text{Bil}}^2}$.

In the rest of the paper, we use the same definition σ as in Theorem 2.3. The proof of Theorem 2.3 is provided in §D.4. We further note that one possible choice of step size is to let $\alpha_t \equiv \alpha$, such that (11) reduces to

$$\mathbb{E} \|z_t - z^*\|^2 \leq \|z_0 - z^*\|^2 \left(\frac{L}{\mu} + 1 \right) e^{-\alpha t} + \frac{3\sigma^2}{\mu^2} \alpha.$$

For any given $T \geq 1$, by choosing the optimal $\alpha = \frac{1}{T} \left(1 + \log \left(\frac{\mu^2 T}{3\sigma^2} \left(\frac{L}{\mu} + 1 \right) \|z_0 - z^*\|^2 \right) \right) \wedge \bar{\alpha}$, (11) implies

$$\mathbb{E} \|z_T - z^*\|^2 \leq \|z_0 - z^*\|^2 \left(\frac{L}{\mu} + 1 \right) e^{-\bar{\alpha} T} + \frac{3\sigma^2}{\mu^2 T} \left(1 + \log \left(\frac{\mu^2 T}{3\sigma^2} \left(\frac{L}{\mu} + 1 \right) \|z_0 - z^*\|^2 \right) \right).$$

Prescribing the desired accuracy $\varepsilon > 0$, the iteration complexity to output an ε -optimal minimax point is⁵

$$\mathcal{O} \left(\left(\sqrt{\frac{L}{\mu}} + \frac{M}{\mu} + \frac{\sigma^2}{\mu^2 \varepsilon^2} \right) \log \left(\left(\frac{L}{\mu} + 1 \right) \|z_0 - z^*\|^2 / \varepsilon^2 \right) \right).$$

We conjecture that the logarithmic factor in the optimal statistical rate $\frac{\sigma^2}{\mu^2 \varepsilon^2}$ is removable using a proper diminishing step size, a possibility that we reserve for future study. In the setting of deterministic optimization, setting $\sigma = 0$ and $r \rightarrow 1^-$, $\beta \rightarrow 0^+$ in Theorem 2.3 we obtain the optimal iteration complexity bound as follows:

$$\left(1 + \sqrt{1 + \frac{L}{\mu} + \frac{M^2}{\mu^2}} \right) \log \left(\left(\frac{L}{\mu} + 1 \right) / \varepsilon^2 \right). \quad (12)$$

⁵Throughout this work, we focus on the iteration complexity whereas the required number of queries to the stochastic gradient oracle is three times the iteration complexity.

Algorithm 1 Stochastic Accelerated Gradient-Extra Gradient (AG-EG) Descent-Ascent Algorithm, with Scheduled Restarting

Require: Initialization $z_0^{[0]}$, total number of epochs $\mathcal{S} \geq 1$, total number of per-epoch iterates $(T_s : s = 1, \dots, \mathcal{S})$, stepsizes $(\alpha_t, \eta_t : t = 1, 2, \dots)$.

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

 Set $z_{-\frac{1}{2}}^{\text{ag}} \leftarrow z_0^{[s-1]}, z_0 \leftarrow z_0^{[s-1]}, z_0^{\text{md}} \leftarrow z_0^{[s-1]}$

for $t = 1, 2, \dots, T_s$ **do**

 Draw samples $\xi_{t-\frac{1}{2}} \sim \mathcal{D}_\xi$ from oracle, and also $\zeta_{t-\frac{1}{2}}, \zeta_t \sim \mathcal{D}_\zeta$ independently from oracle

$z_{t-\frac{1}{2}} \leftarrow z_{t-1} - \eta_t \left(\tilde{\mathcal{H}}(z_{t-1}; \zeta_{t-\frac{1}{2}}) + \nabla \tilde{\mathcal{F}}(z_{t-1}^{\text{md}}; \xi_{t-\frac{1}{2}}) \right)$

$z_{t-\frac{1}{2}}^{\text{ag}} \leftarrow (1 - \alpha_t) z_{t-\frac{3}{2}}^{\text{ag}} + \alpha_t z_{t-\frac{1}{2}}$

$z_t \leftarrow z_{t-1} - \eta_t \left(\tilde{\mathcal{H}}(z_{t-\frac{1}{2}}; \zeta_t) + \nabla \tilde{\mathcal{F}}(z_{t-1}^{\text{md}}; \xi_{t-\frac{1}{2}}) \right)$

$z_t^{\text{md}} \leftarrow (1 - \alpha_{t+1}) z_{t-\frac{1}{2}}^{\text{ag}} + \alpha_{t+1} z_t$

end for

 Set $z_0^{[s]} \leftarrow z_{T_s-\frac{1}{2}}^{\text{ag}}$ {//Warm-start using the output of the previous epoch}

end for

Output: $z_0^{[\mathcal{S}]}$

Remark 2.4 Our complexity bounds fundamentally differs from the previous analysis [Chen et al. 2017, Jordan et al. 2023] for separable smooth (strongly) monotone VIs. The convergence results in previous studies are dependent on the diameter of the domain, whereas our convergence rate is independent of the domain parameters and eliminates the need for projection onto a bounded domain. Moreover, our contributions go beyond those of Chen et al. [2017] by extending the analysis to the strongly monotone case. In comparison with Jordan et al. [2023], we design an algorithm where $\nabla \mathcal{F}$ is strongly monotone and resolve the open problem of extending the analysis to the stochastic case. Additionally, our complexity bound in (12) indicates a near-unity coefficient on the condition-number exponent, improving the corresponding coefficient in Chen et al. [2017 Theorem 15] by an asymptotic factor of 4.

The direct approach, which reduces to EG when $\nabla \mathcal{F} = 0$ and $\mu = 0$, falls short of attaining optimality within the specific regime of bilinear games. In the next subsection, we will introduce a new algorithm that can overcome this limitation.

2.4 Accelerating the Extra Gradient Algorithm with Scheduled Restarting

In this subsection, we solve problem (1) by further accelerating the stochastic EG algorithm. Rather than directly relying on the strong monotonicity of $\nabla \mathcal{F}$, the inner updates of our new algorithm are identical to the updates in (10) with $\mu = 0$. Due to the domain-independent nature of our analysis, we can apply the scheduled restarting technique [O’Donoghue and Candes, 2015, Roulet and d’Aspremont, 2017, Renegar and Grimmer, 2022] to the outer loop, accelerating the algorithm from sublinear convergence to linear convergence. In addition, the output of our algorithm is the aggregated point $z_{T-\frac{1}{2}}^{\text{ag}}$ after T iterates. We present the full algorithm in Algorithm 1.

We first present the convergence rate of a single epoch (i.e., the inner loop) of Algorithm 1 in Theorem 2.5. To accommodate more flexibility in the choice of parameters, we introduce three constants r, β , and C in the theorem statement.

Theorem 2.5 (Convergence of stochastic AG-EG, one epoch) Suppose Assumptions 2.1 and 2.2 hold. For any fixed epoch length $T \geq 1$, any constant $r \in (0, 1)$, $\beta \in (0, \infty)$, $C \in (0, \infty)$, choose step sizes $\alpha_t = \frac{2}{t+1}$ and η_t such that

$$\frac{t}{\eta_t} = \frac{2}{r} L \vee B + \sqrt{\frac{1+\beta}{r}} Mt, \quad (13)$$

where $B = \frac{\sigma\sqrt{T(T+1)}}{C\sqrt{\mathbb{E}\|z_0 - z^*\|^2}}$. The output $z_{T-\frac{1}{2}}^{ag}$ of a single epoch of Algorithm 1 satisfies

$$\mathbb{E}\left\|z_{T-\frac{1}{2}}^{ag} - z^*\right\|^2 \leq \frac{2}{\mu(T+1)} \left(\frac{2L}{rT} + A\sqrt{\frac{1+\beta}{r}}M \right) \mathbb{E}\|z_0 - z^*\|^2 + \frac{2(\frac{1}{C}+C)\sigma}{\mu\sqrt{T}} \sqrt{\mathbb{E}\|z_0 - z^*\|^2}, \quad (14)$$

where the prefactor $A \equiv 1 + C^2 B \eta_1 \leq 1 + C^2$ reduces to 1 when $\sigma = 0$.

The proof of Theorem 2.5 is provided in §D.3. We make a few remarks on Theorem 2.5 as follows:

Remark 2.6 In the setting of deterministic optimization, by taking $\sigma = 0$, $r \rightarrow 1^-$, $\beta \rightarrow 0^+$ in our analysis, with step size choice $\eta_t = \frac{t}{2L+Mt}$, we obtain that

$$\|z_{T-\frac{1}{2}}^{ag} - z^*\|^2 \leq \frac{2}{\mu(T+1)} \left(\frac{2L}{rT} + M \right) \|z_0 - z^*\|^2, \quad (15)$$

In this setting, the algorithm is independent of B and requires no knowledge of $\|z_0 - z^*\|^2$. In the face of stochasticity, we choose $C = 1$ when the initial distance to the optimal point is known. Alternatively, when only an over-estimate Γ_0 of $\sqrt{\mathbb{E}\|z_0 - z^*\|^2}$ is available, we can set (large enough) $C = \frac{\Gamma_0}{\sqrt{\mathbb{E}\|z_0 - z^*\|^2}} \geq 1$ to obtain

$$\mathbb{E}\|z_{T-\frac{1}{2}}^{ag} - z^*\|^2 \leq \frac{2}{\mu(T+1)} \left(\frac{2L}{rT} + 2\sqrt{\frac{1+\beta}{r}}M \right) \Gamma_0^2 + \frac{4\sigma}{\mu\sqrt{T}}\Gamma_0. \quad (16)$$

Remark 2.7 When the constants are not a concern, the coarse-grained choices of $r = \frac{1}{2}$ and $\beta = 1$ would suffice. Nevertheless, to optimize the constants, the tradeoff between the deviation of r from 1 and β from 0 is crucial, as it determines a balance between the stochastic gradient noise variance and the convergence rate coefficients.

To prepare for our multi-epoch result with the help of scheduled restarting, we perform an induction based on (16) as follows. Supposing that $\mathbb{E}\|z_0^{[s-1]} - z^*\|^2 \leq \Gamma_0^2 e^{1-s}$ hold, by taking $r = \frac{1}{2}$ and $\beta = 1$, we have

$$\mathbb{E}\|z_0^{[s]} - z^*\|^2 \lesssim \frac{L}{\mu T_s^2} \Gamma_0^2 e^{1-s} + \frac{M}{\mu T_s} \Gamma_0^2 e^{1-s} + \frac{\sigma}{\mu\sqrt{T_s}} \Gamma_0 e^{\frac{1-s}{2}}.$$

Setting the right-hand side of the above inequality to satisfy $\leq \Gamma_0^2 e^{-s}$, and solving for T_s , we need the epoch length satisfies $T_s \asymp \sqrt{\frac{L}{\mu}} + \frac{M}{\mu} + \frac{\sigma^2}{\mu^2 \Gamma_0^2 e^{1-s}}$. Thus, we can obtain the total iteration complexity as

$$\sum_{s=1}^S \left[\sqrt{\frac{L}{\mu}} + \frac{M}{\mu} + \frac{\sigma^2}{\mu^2 \Gamma_0^2 e^{1-s}} \right] = \left(\sqrt{\frac{L}{\mu}} + \frac{M}{\mu} \right) S + \frac{\sigma^2}{\mu^2 \Gamma_0^2} \cdot \frac{e^S - 1}{e - 1},$$

where $S \equiv \left\lceil \log \frac{\Gamma_0^2}{\varepsilon^2} \right\rceil$. This yields the following multi-epoch iteration complexity bound:

Corollary 2.8 (Iteration complexity of stochastic AG-EG with scheduled restarting) Under the same condition of Theorem 2.5, the stochastic AG-EG with scheduled restarting in Algorithm 1 with epoch length $T_s \asymp \sqrt{\frac{L}{\mu}} + \frac{M}{\mu} + \frac{\sigma^2}{\mu^2 \Gamma_0^2 e^{1-s}}$ has a total iteration complexity of

$$\mathcal{O} \left(\left(\sqrt{\frac{L}{\mu}} + \frac{M}{\mu} \right) \log \left(\frac{1}{\varepsilon} \right) + \frac{\sigma^2}{\mu^2 \varepsilon^2} \right). \quad (17)$$

Note that the hard instance constructed by Zhang et al. [2022] can be modified in a straightforward way to establish a lower bound of $\Omega \left(\left(\sqrt{\frac{L}{\mu}} + \frac{M}{\mu} \right) \log \left(\frac{1}{\varepsilon} \right) \right)$ for our monotone VI (1), demonstrating the optimality of Corollary 2.8 in the deterministic separable setting. An alternative optimality argument proceeds as follows: the first term $\sqrt{\frac{L}{\mu}}$ matches the lower bound for the minimization of a strongly convex function \mathcal{F} [Nesterov, 2004], and the second term $\frac{M}{\mu}$ matches the lower bound for VI for non-strongly monotone operator when $\nabla \mathcal{F} = 0$ [Ouyang and Xu, 2021]. This together gives a lower bound for solving monotone VI (1) via a similar argument by Thekumparampil et al. [2022].

It is worth noting that while both complexity bounds in Corollary 2.8 and Theorem 2.3 match the lower bound in Zhang et al. (2022) for strongly monotone VIs with separable structure, the direct approach in §2.3 reduces to the *last-iterate* independent-sample stochastic extragradient (SEG) algorithm in *bilinear games*. Consequently, the deterministic part ($\sigma = 0$) fails to match the lower bound in Ibrahim et al. (2020). In the stochastic case with noise variance bounded away from zero, the direct approach in §2.3 can exhibit *nonconvergence behavior* (Hsieh et al. 2020, §3). The AG-EG algorithm in §2.4 resolves this issue by restarting the *average-iterate* SEG, matching the lower bound results (see §3.2 for more details). In addition, the complexity bound in (17) also eliminates the log prefactor of the statistical error term $\frac{\sigma^2}{\mu^2 \varepsilon^2}$ compared to Theorem 2.3. The optimality of our algorithm lies in not only the optimization complexity but also the statistical error rate $\frac{\sigma^2}{\mu^2 \varepsilon^2}$. Here the ε -optimal point \mathbf{z} is defined as $\|\mathbf{z} - \mathbf{z}^*\| \leq \varepsilon$.⁴

2.5 Extension of AG-EG to Proximal Algorithms

In previous subsections, we have focused on the case where the feasible set \mathcal{Z} represents the entire space and the nondifferentiable convex function J is dropped. We now extend the AG-EG algorithm and its analysis to the more general setting that has a bounded feasible set (via Euclidean projection onto the feasible set) as well as a nondifferentiable convex regularization term (via a proximal operator). These settings are useful in various applications, such as the variational inequality on the Lorentz cone where projection onto $\mathcal{Z} = \{(\mathbf{x}, t) \in \mathbb{R}^{(n+1)} : \|\mathbf{x}\| \leq t\}$ is required (Chen et al. 2017), and the two-player game that involves projection onto the probability simplex, among others. To deal with bounded feasible set \mathcal{Z} , we adopt a variant of the EG algorithm, where we project the extrapolated point and the main iterates back onto the feasible set \mathcal{Z} of \mathcal{W} :

$$\begin{aligned} \mathbf{z}_{t-\frac{1}{2}} &= P_{\mathcal{Z}}[\mathbf{z}_{t-1} - \eta \mathcal{W}(\mathbf{z}_{t-1})] = \arg \min_{\mathbf{z} \in \mathcal{Z}} \langle \mathbf{z} - \mathbf{z}_{t-1}, \eta \mathcal{W}(\mathbf{z}_{t-1}) \rangle + \frac{1}{2} \|\mathbf{z} - \mathbf{z}_{t-1}\|^2, \\ \mathbf{z}_t &= P_{\mathcal{Z}}\left[\mathbf{z}_{t-1} - \eta \mathcal{W}(\mathbf{z}_{t-\frac{1}{2}})\right] = \arg \min_{\mathbf{z} \in \mathcal{Z}} \left\langle \mathbf{z} - \mathbf{z}_{t-1}, \eta \mathcal{W}(\mathbf{z}_{t-\frac{1}{2}}) \right\rangle + \frac{1}{2} \|\mathbf{z} - \mathbf{z}_{t-1}\|^2, \end{aligned} \quad (18)$$

where $P_{\mathcal{Z}}(\mathbf{z}) = \arg \min_{\mathbf{z}' \in \mathcal{Z}} \|\mathbf{z} - \mathbf{z}'\|^2$ is the Euclidean projection operator. To handle the nondifferentiable simple convex function J , we replace the projection operator in (18) by the following proximal mapping defined via a Bregman divergence $\mathcal{B}(\cdot, \cdot)$:

$$\text{prox}_{\mathcal{Z}}^J(\mathbf{v}) \equiv \arg \min_{\mathbf{u} \in \mathcal{Z}} \langle \mathbf{v}, \mathbf{u} - \mathbf{z} \rangle + \mathcal{B}(\mathbf{z}, \mathbf{u}) + J(\mathbf{u}). \quad (19)$$

In fact, (18) can be seen as a special case of (19) when choosing the Bregman divergence $\mathcal{B}(\mathbf{z}, \mathbf{u}) = \frac{1}{2} \|\mathbf{z} - \mathbf{u}\|^2$ and $J(\mathbf{u})$ as the set indicator function of the feasible set \mathcal{Z} . Therefore, by substituting the prox-mapping (19) into the AG-EG updates introduced in §2.4, we obtain the more general proximal AG-EG algorithm in Algorithm 3 (See in §C.2), which reduces to Algorithm 1 when $J = 0$, $\mathcal{B}(\mathbf{z}, \mathbf{u}) = \frac{1}{2} \|\mathbf{z} - \mathbf{u}\|^2$ and $\mathcal{Z} = \mathbb{R}^n$. Moreover, we assume that $\mathcal{B}(\cdot, \cdot)$ is $\mu_{\mathcal{B}}$ -strongly convex. Without loss of generality, in contrast to the previous assumption of μ -strong convexity for \mathcal{F} , we instead assume that \mathcal{F} is μ -strongly convex with respect to the Bregman divergence $\mathcal{B}(\cdot, \cdot)$ (See, for example, Hazan and Kale (2014), Xu et al. (2018)). Similar to Corollary 2.8, we have the following iteration complexity result, whose proof is deferred to §D.5:

Corollary 2.9 (Iteration complexity of stochastic proximal AG-EG with scheduled restarting)
Under the same condition of Theorem 2.5, the stochastic proximal AG-EG with scheduled restarting in Algorithm 3 with epoch length $T_s \asymp \sqrt{\frac{L}{\mu \mu_{\mathcal{B}}}} + \frac{M}{\mu \mu_{\mathcal{B}}} + \frac{\sigma^2 \mathcal{B}(\mathbf{z}_0, \mathbf{z}^*)}{\mu^2 \mu_{\mathcal{B}} \Gamma_0^2 e^{1-s}}$, has a total iteration complexity of

$$\mathcal{O}\left(\left(\sqrt{\frac{L}{\mu \mu_{\mathcal{B}}}} + \frac{M}{\mu \mu_{\mathcal{B}}}\right) \log\left(\frac{1}{\varepsilon}\right) + \frac{\sigma^2 \mathcal{B}(\mathbf{z}_0, \mathbf{z}^*)}{\mu^2 \mu_{\mathcal{B}} \varepsilon^2}\right).$$

For the deterministic case, proximal AG-EG with scheduled restarting has a total iteration complexity of $\mathcal{O}\left(\left(\sqrt{\frac{L}{\mu \mu_{\mathcal{B}}}} + \frac{M}{\mu \mu_{\mathcal{B}}}\right) \log\left(\frac{1}{\varepsilon}\right)\right)$ to output an ε -optimal point of (1).

⁴The optimal statistical error rate $\frac{\sigma^2}{\mu^2 T}$ has been achieved by a multistage algorithm in Fallah et al. (2020), where the ε -optimal point is defined by $\|\mathbf{z} - \mathbf{z}^*\|^2 \leq \varepsilon$. In our paper, the ε -optimal point is defined by $\|\mathbf{z} - \mathbf{z}^*\| \leq \varepsilon$. Therefore, our statistical error rate can be translated into $\frac{\sigma^2}{\mu^2 T}$ using their definition, which matches their result.

3 Implications for Specific Instances

In this section, we discuss the implications of our AG-EG algorithm and its convergence rates when applying to two instances of saddle-point problems.

3.1 Strongly-Convex-Strongly-Concave Saddle-Point Problem

For the stochastic bilinearly-coupled SC-SC saddle-point problem (3), we note that the smoothness and strong convexity parameters L_F , L_G , μ_F , and μ_G of F and G may differ. To accommodate these variations in curvature information, we employ a scaling reduction technique. This technique enables us to convert the SC-SC with equal strong convexity parameters for F and G by reparametrizing the objective function. The same argument is also applicable to the direct approach.

In lieu of (3), we consider

$$\min_{\hat{\mathbf{x}}} \max_{\hat{\mathbf{y}}} \hat{\mathcal{F}}(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = F(\hat{\mathbf{x}}) + \hat{H}(\hat{\mathbf{x}}, \hat{\mathbf{y}}) - \hat{G}(\hat{\mathbf{y}}),$$

where $\hat{\mathcal{F}}(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = \mathcal{F}(\mathbf{x}, \mathbf{y})$ with the symbolic reparametrization $\hat{\mathbf{x}} = \mathbf{x}$, $\hat{\mathbf{y}} = \sqrt{\frac{\mu_G}{\mu_F}} \mathbf{y}$, $\hat{H}(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = H(\mathbf{x}, \mathbf{y})$, $\hat{G}(\hat{\mathbf{y}}) = G(\mathbf{y})$ and also their derivatives $\nabla_{\hat{\mathbf{y}}} \hat{H}(\hat{\mathbf{x}}, \hat{\mathbf{y}}) = \sqrt{\frac{\mu_F}{\mu_G}} \nabla_{\mathbf{y}} H(\mathbf{x}, \mathbf{y})$, $\nabla \hat{G}(\hat{\mathbf{y}}) = \sqrt{\frac{\mu_F}{\mu_G}} \nabla G(\mathbf{y})$ (the stochastic oracles \hat{h}, \hat{g} follow the same rule). It is straightforward to verify that $\hat{\mathcal{F}}(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is μ -strongly-convex- μ -strongly-concave. The essence of our update rules can be summarized by the rescaled updates on \mathbf{y} :

$$\begin{aligned} \hat{\mathbf{y}}_t &= \hat{\mathbf{y}}_{t-1} - \eta_t \left(-\nabla_{\hat{\mathbf{y}}} h(\hat{\mathbf{x}}_{t-\frac{1}{2}}, \hat{\mathbf{y}}_{t-\frac{1}{2}}; \zeta_t) + \nabla g(\hat{\mathbf{y}}_{t-1}^{\text{md}}; \xi_{t-\frac{1}{2}}) \right) \\ &\Leftrightarrow \mathbf{y}_t = \mathbf{y}_{t-1} - \eta_t \cdot \frac{\mu_F}{\mu_G} \left(-\nabla_{\mathbf{y}} h(\mathbf{x}_{t-\frac{1}{2}}, \mathbf{y}_{t-\frac{1}{2}}; \zeta_t) + \nabla g(\mathbf{y}_{t-1}^{\text{md}}; \xi_{t-\frac{1}{2}}) \right). \end{aligned}$$

Therefore, it suffices to analyze Algorithm 3 for $\hat{\mathcal{F}}(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ and due to this scaling reduction, we only need to prove all results for the case of $\mu_F = \mu_G = \mu$. It is also straightforward to justify corresponding scaling changes as: $L = L_F \vee \frac{\mu_F}{\mu_G} L_G$, $M = \sqrt{\frac{\mu_F}{\mu_G}} \lambda_{\max}(\mathbf{B}^\top \mathbf{B})$, and $\mu = \mu_F$. The following corollary is recovered by reverting the scaling reduction from $\hat{\mathcal{F}}(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ to $\mathcal{F}(\mathbf{x}, \mathbf{y})$.

Corollary 3.1 (Iteration complexity of stochastic AG-EG on SC-SC saddle-point problem)

For solving (3), Algorithm 1 with an epoch length $T_s \asymp \sqrt{\frac{L_F}{\mu_F} \vee \frac{L_G}{\mu_G}} + \sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\mu_F \mu_G}} + \frac{\sigma^2}{\mu_F^2 \Gamma_0^2 e^{1-s}}$ has a total iteration complexity of

$$\mathcal{O} \left(\left(\sqrt{\frac{L_F}{\mu_F} \vee \frac{L_G}{\mu_G}} + \sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\mu_F \mu_G}} \right) \log \left(\frac{1}{\varepsilon} \right) + \frac{\sigma^2}{\mu_F^2 \varepsilon^2} \right).$$

In the deterministic case, the iteration complexity in Theorem 2.8 matches the lower bound established by Zhang et al. [2022], i.e., $\Omega \left(\left(\sqrt{\frac{L_F}{\mu_F} \vee \frac{L_G}{\mu_G}} + \sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\mu_F \mu_G}} \right) \log \left(\frac{1}{\varepsilon} \right) \right)$. Moreover, our algorithm achieves the optimal statistical rate of $\frac{\sigma^2}{\mu_F^2 \varepsilon^2}$ up to a constant prefactor.

Remark 3.2 A well-known finding regarding the second scheme of Nesterov acceleration is its connection to the primal-dual method [Lan and Zhou, 2018; Lin et al., 2020c]. This finding has been incorporated into the design of the LPD algorithm [Thekumparampil et al., 2022], where a Chambolle-Pock-style primal-dual method is utilized as an approximation of proximal point methods, instead of the extragradient used in this paper. The LPD algorithm [Thekumparampil et al., 2022] also achieves the optimal complexity for the deterministic bilinearly-coupled saddle-point problem.

3.2 Bilinear Games

In this subsection, we consider the particular case of bilinear games. We assume $n = m$ such that \mathbf{B} is a nonsingular square matrix, $\nabla f(\mathbf{x}; \xi) = \mathbf{0}$ and $\nabla g(\mathbf{y}; \xi) = \mathbf{0}$ a.s., so (3) reduces to

$$\min_{\mathbf{x}} \max_{\mathbf{y}} \mathcal{F}(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{\zeta} [h(\mathbf{x}, \mathbf{y}; \zeta)] = H(\mathbf{x}, \mathbf{y}) = \mathbf{x}^\top \mathbf{B} \mathbf{y} - \mathbf{x}^\top \mathbf{u}_x + \mathbf{u}_y^\top \mathbf{y}, \quad (20)$$

and Algorithm 3 reduces to the independent-sample extragradient descent-ascent algorithm for (20). The saddle point $[z^*; \omega_y^*]$ in this case is the unique solution to the linear equation

$$\begin{bmatrix} \mathbf{0} & \mathbf{B} \\ -\mathbf{B}^\top & \mathbf{0} \end{bmatrix} \begin{bmatrix} z^* \\ \omega_y^* \end{bmatrix} = \begin{bmatrix} \mathbf{u}_x \\ \mathbf{u}_y \end{bmatrix}, \quad \text{which has a closed-form solution } \begin{bmatrix} z^* \\ \omega_y^* \end{bmatrix} = \begin{bmatrix} -(\mathbf{B}^\top)^{-1} \mathbf{u}_y \\ \mathbf{B}^{-1} \mathbf{u}_x \end{bmatrix}.$$

Our results imply the following iteration complexity for solving stochastic bilinear games.

Corollary 3.3 (Iteration complexity of stochastic AG-EG, bilinear games) For solving (20), choose the step sizes $\alpha_t = \frac{2}{t+1}$ and $\eta_t \equiv \frac{1}{\sqrt{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}}$, in which case Algorithm 1 with an epoch length $T_s \asymp \sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top)}}$ has the total iteration complexity of

$$\mathcal{O} \left(\sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top)}} \log \left(\frac{\sqrt[4]{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top) \lambda_{\max}(\mathbf{B}^\top \mathbf{B})}}{\sigma_{\text{Bil}}} \right) + \frac{\sigma_{\text{Bil}}^2}{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top) \varepsilon^2} \right). \quad (21)$$

Note that our choice of the step size is maximal and is independent of the noise. In the deterministic setting, letting $\sigma_{\text{Bil}} \asymp \varepsilon \sqrt{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top) \lambda_{\max}(\mathbf{B}^\top \mathbf{B})}$, the complexity bound in Corollary 3.3 reduces to $\mathcal{O} \left(\sqrt{\frac{\lambda_{\max}(\mathbf{B}^\top \mathbf{B})}{\lambda_{\min}(\mathbf{B}\mathbf{B}^\top)}} \log \left(\frac{1}{\varepsilon} \right) \right)$, which matches the lower bound in Ibrahim et al. [2020]. Notably, Azizian et al. [2020b] proposed an algorithm achieving an upper bound that matches the lower bound in Ibrahim et al. [2020]. Li et al. [2022a] also proposed a lower-bound matching SEG algorithm that uses a shared sample in both steps under an unbounded noise assumption. In contrast, our algorithm is in the independent-sample setting with bounded noise variance.

Remark 3.4 Standard acceleration techniques do not attain the optimal nonasymptotic convergence rate for bilinear games [Gidel et al. 2019b]. This limitation applies to various algorithms, including the direct approach [§2.3], as well as several other acceleration techniques [Thekumparampil et al. 2022, Kovalev et al. 2022, Jin et al. 2022], all of which fall short of achieving optimal acceleration for bilinear games. Therefore, matching both lower bounds in a single algorithm in the general stochastic setting has been an open problem. While Li et al. [2022b] present an algorithm that achieves both lower bounds in a single algorithm, it relies on the use of optimistic gradients rather than extragredients on the bilinear coupling function. Furthermore, our algorithm and analysis is more general than those in Li et al. [2022b] as we can handle the general variational inequality with proximal operators.

4 Conclusions

We have presented a stochastic extragradient-based acceleration algorithm, AG-EG, for solving stochastic monotone variational inequalities with separable structure. The iteration complexity of our algorithm matches the lower bound and is independent of the size of the feasible set. When specialized to solving the bilinearly coupled saddle-point problem (3), our AG-EG algorithm simultaneously matches lower bounds due to Zhang et al. [2022] and Ibrahim et al. [2020] for strongly-convex-strongly-concave and bilinear games, respectively. To the best of our knowledge, this is the first time that all three lower bounds have been met by a single algorithm. There are some remaining issues to be addressed, however, including the case of one-sided nonstrong convexity, the setting of unbounded noise variance, and the characterization of the full parameter regime dependency on $\lambda_{\min}(\mathbf{B}\mathbf{B}^\top)$. These are left as important directions for future research.

Acknowledgments and Disclosure of Funding

This work is supported in part by Canada CIFAR AI Chair to GG, by the Mathematical Data Science program of the Office of Naval Research under grant number N00014-18-1-2764 and also the Vannevar Bush Faculty Fellowship program under grant number N00014-21-1-2941 and National Science Foundation (NSF) grant IIS-1901252 to MIJ. This work is also supported in part by the NSF CAREER Award 1906169 to QG and by NSF Award’s IIS-2110170 and DMS-2134106 to SSD.

References

- Waïss Azizian, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. A tight and unified analysis of gradient-based methods for a whole spectrum of differentiable games. In *International Conference on Artificial Intelligence and Statistics*, pages 2863–2873. PMLR, 2020a.
- Waïss Azizian, Damien Scieur, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. Accelerating smooth games by manipulating spectral shapes. In *International Conference on Artificial Intelligence and Statistics*, pages 1705–1715. PMLR, 2020b.
- Yunmei Chen, Guanghui Lan, and Yuyuan Ouyang. Optimal primal-dual methods for a class of saddle point problems. *SIAM Journal on Optimization*, 24(4):1779–1814, 2014.
- Yunmei Chen, Guanghui Lan, and Yuyuan Ouyang. Accelerated schemes for a class of variational inequalities. *Mathematical Programming*, 165(1):113–149, 2017.
- Ziyi Chen, Qunwei Li, and Yi Zhou. Finding local minimax points via (stochastic) cubic-regularized gda: Global convergence and complexity. *arXiv preprint arXiv:2110.07098*, 2021.
- Michael B Cohen, Aaron Sidford, and Kevin Tian. Relative Lipschitzness in extragradient methods and a direct recipe for acceleration. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, volume 185, page 62, 2021.
- Christoph Dann, Gerhard Neumann, Jan Peters, et al. Policy evaluation with temporal differences: A survey and comparison. *Journal of Machine Learning Research*, 15:809–883, 2014.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training GANs with optimism. In *International Conference on Learning Representations*, 2018.
- Simon S Du, Jianshu Chen, Lihong Li, Lin Xiao, and Dengyong Zhou. Stochastic variance reduction methods for policy evaluation. In *International Conference on Machine Learning*, pages 1049–1058. PMLR, 2017.
- Francisco Facchinei and Jong-Shi Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer, 2003.
- Alireza Fallah, Asuman Ozdaglar, and Sarath Pattathil. An optimal multistage stochastic gradient method for minimax problems. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 3573–3579. IEEE, 2020.
- Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization i: A generic algorithmic framework. *SIAM Journal on Optimization*, 22(4):1469–1492, 2012.
- Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. In *International Conference on Learning Representations*, 2019a.
- Gauthier Gidel, Reyhane Askari Hemmat, Mohammad Pezeshki, Rémi Le Priol, Gabriel Huang, Simon Lacoste-Julien, and Ioannis Mitliagkas. Negative momentum for improved game dynamics. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1802–1811. PMLR, 2019b.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Zhishuai Guo, Zhuoning Yuan, Yan Yan, and Tianbao Yang. Fast objective & duality gap convergence for nonconvex-strongly-concave min-max problems. *arXiv preprint arXiv:2006.06889*, 2020.
- Elad Hazan and Satyen Kale. Beyond the regret minimization barrier: optimal algorithms for stochastic strongly-convex optimization. *The Journal of Machine Learning Research*, 15(1): 2489–2512, 2014.

- Yu-Guan Hsieh, Franck Iutzeler, Jérôme Malick, and Panayotis Mertikopoulos. Explore aggressively, update conservatively: Stochastic extragradient methods with variable stepsize scaling. In *Advances in Neural Information Processing Systems*, volume 33, pages 16223–16234, 2020.
- Adam Ibrahim, Waiss Azizian, Gauthier Gidel, and Ioannis Mitliagkas. Linear lower bounds and conditioning of differentiable games. In *International Conference on Machine Learning*, pages 4583–4593. PMLR, 2020.
- Alfredo N Iusem, Alejandro Jofré, Roberto Imbuzeiro Oliveira, and Philip Thompson. Extragradient method with variance reduction for stochastic variational inequalities. *SIAM Journal on Optimization*, 27(2):686–724, 2017.
- Samy Jelassi, Carles Domingo-Enrich, Damien Scieur, Arthur Mensch, and Joan Bruna. Extragradient with player sampling for faster convergence in n-player games. In *International Conference on Machine Learning*, pages 4736–4745. PMLR, 2020.
- Yujia Jin, Aaron Sidford, and Kevin Tian. Sharper rates for separable minimax and finite sum optimization via primal-dual extragradient methods. In *Conference on Learning Theory*, pages 4362–4415. PMLR, 2022.
- Michael I Jordan, Tianyi Lin, and Manolis Zampetakis. First-order algorithms for nonlinear generalized Nash equilibrium problems. *Journal of Machine Learning Research*, 24(38):1–46, 2023.
- Anatoli Juditsky, Arkadi Nemirovski, and Claire Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 1(1):17–58, 2011.
- Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Ekonomika i Matematicheskie Metody*, 12:747–756, 1976.
- Georgios Kotsalis, Guanghui Lan, and Tianjiao Li. Simple and optimal methods for stochastic variational inequalities, I: Operator extrapolation. *arXiv preprint arXiv:2011.02987*, 2020.
- Dmitry Kovalev, Alexander Gasnikov, and Peter Richtárik. Accelerated primal-dual gradient method for smooth and convex-concave saddle-point problems with bilinear coupling. *Advances in Neural Information Processing Systems*, 35:21725–21737, 2022.
- Guanghui Lan and Yuyuan Ouyang. Mirror-prox sliding methods for solving a class of monotone variational inequalities. *arXiv preprint arXiv:2111.00996*, 2021.
- Guanghui Lan and Yi Zhou. An optimal randomized incremental gradient method. *Mathematical Programming*, 171:167–215, 2018.
- Chris Junchi Li, Yaodong Yu, Nicolas Loizou, Gauthier Gidel, Yi Ma, Nicolas Le Roux, and Michael Jordan. On the convergence of stochastic extragradient for bilinear games using restarted iteration averaging. In *International Conference on Artificial Intelligence and Statistics*, pages 9793–9826. PMLR, 2022a.
- Chris Junchi Li, Angela Yuan, Gauthier Gidel, and Michael I Jordan. Nesterov meets optimism: Rate-optimal optimistic-gradient-based method for stochastic bilinearly-coupled minimax optimization. *arXiv preprint arXiv:2210.17550*, 2022b.
- Tengyuan Liang and James Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *International Conference on Artificial Intelligence and Statistics*, pages 907–915. PMLR, 2019.
- Tianyi Lin, Chi Jin, and Michael Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pages 6083–6093. PMLR, 2020a.
- Tianyi Lin, Chi Jin, and Michael I Jordan. Near-optimal algorithms for minimax optimization. In *Conference on Learning Theory*, pages 2738–2779. PMLR, 2020b.
- Zhouchen Lin, Huan Li, and Cong Fang. Accelerated optimization for machine learning. *Nature Singapore: Springer*, 2020c.

- Nicolas Loizou, Hugo Berard, Alexia Jolicoeur-Martineau, Pascal Vincent, Simon Lacoste-Julien, and Ioannis Mitliagkas. Stochastic hamiltonian gradient methods for smooth games. In *International Conference on Machine Learning*, pages 6370–6381. PMLR, 2020.
- Luo Luo, Yujun Li, and Cheng Chen. Finding second-order stationary points in nonconvex-strongly-concave minimax optimization. *Advances in Neural Information Processing Systems*, 35:36667–36679, 2022.
- Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. In *International Conference on Learning Representations*, 2018.
- Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. The numerics of GANs. *Advances in Neural Information Processing Systems*, 30, 2017.
- Dmitriy Metev, Alexander Rogozin, Alexander Gasnikov, and Dmitry Kovalev. Decentralized saddle-point problems with different constants of strong convexity and strong concavity. *arXiv preprint arXiv:2206.00090*, 2022.
- Konstantin Mishchenko, Dmitry Kovalev, Egor Shulgin, Peter Richtárik, and Yura Malitsky. Revisiting stochastic extragradient. In *International Conference on Artificial Intelligence and Statistics*, pages 4573–4582. PMLR, 2020.
- Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *International Conference on Artificial Intelligence and Statistics*, pages 1497–1507. PMLR, 2020.
- Vaishnavh Nagarajan and J Zico Kolter. Gradient descent GAN optimization is locally stable. *Advances in Neural Information Processing Systems*, 30, 2017.
- Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19:1574–1609, 2009.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2004.
- Yurii Nesterov and Laura Scrimali. Solving strongly monotone variational and quasi-variational inequalities. *Discrete & Continuous Dynamical Systems*, 31(4):1383, 2011.
- Yuyuan Ouyang and Yangyang Xu. Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems. *Mathematical Programming*, 185(1-2):1–35, 2021.
- Brendan O’donoghue and Emmanuel Candes. Adaptive restart for accelerated gradient schemes. *Foundations of computational mathematics*, 15(3):715–732, 2015.
- Hassan Rafique, Mingrui Liu, Qihang Lin, and Tianbao Yang. Weakly-convex–concave min–max optimization: provable algorithms and applications in machine learning. *Optimization Methods and Software*, pages 1–35, 2021.
- James Renegar and Benjamin Grimmer. A simple nearly optimal restart scheme for speeding up first-order methods. *Foundations of Computational Mathematics*, 22(1):211–256, 2022.
- Vincent Roulet and Alexandre d’Aspremont. Sharpness, restart and acceleration. *Advances in Neural Information Processing Systems*, 30, 2017.
- Ernest K Ryu, Kun Yuan, and Wotao Yin. Ode analysis of stochastic gradient methods with optimism and anchoring for minimax problems. *arXiv preprint arXiv:1905.10899*, 2019.

- Othmane Sebbouh, Marco Cuturi, and Gabriel Peyré. Randomized stochastic gradient descent ascent. In *International Conference on Artificial Intelligence and Statistics*, pages 2941–2969. PMLR, 2022.
- Kiran K Thekumparampil, Niao He, and Sewoong Oh. Lifted primal-dual method for bilinearly coupled smooth minimax optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 4281–4308. PMLR, 2022.
- Paul Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1-2):237–252, 1995.
- Paul Tseng. On accelerated proximal gradient methods for convex-concave optimization. *submitted to SIAM Journal on Optimization*, 2(3), 2008.
- Jialei Wang and Lin Xiao. Exploiting strong convexity from data with primal-dual first-order algorithms. In *International Conference on Machine Learning*, pages 3694–3702, 2017.
- Yuanhao Wang and Jian Li. Improved algorithms for convex-concave minimax optimization. *Advances in Neural Information Processing Systems*, 33:4800–4810, 2020.
- Lin Xiao, Adams Wei Yu, Qihang Lin, and Weizhu Chen. Dscovr: Randomized primal-dual block coordinate algorithms for asynchronous distributed optimization. *Journal of Machine Learning Research*, 20(1):1634–1691, 2019.
- Guangzeng Xie, Yuze Han, and Zhihua Zhang. Dippa: An improved method for bilinear saddle point problems. *arXiv preprint arXiv:2103.08270*, 2021.
- Pan Xu, Tianhao Wang, and Quanquan Gu. Continuous and discrete-time accelerated stochastic mirror descent for strongly convex functions. In *International Conference on Machine Learning*, pages 5492–5501. PMLR, 2018.
- Yan Yan, Yi Xu, Qihang Lin, Lijun Zhang, and Tianbao Yang. Stochastic primal-dual algorithms with faster convergence than $o(1/\sqrt{T})$ for problems without bilinear structure. *arXiv preprint arXiv:1904.10112*, 2019.
- Yan Yan, Yi Xu, Qihang Lin, Wei Liu, and Tianbao Yang. Optimal epoch stochastic gradient descent ascent methods for min-max optimization. *Advances in Neural Information Processing Systems*, 33:5789–5800, 2020.
- Junchi Yang, Antonio Orvieto, Aurelien Lucchi, and Niao He. Faster single-loop algorithms for minimax optimization without strong concavity. In *International Conference on Artificial Intelligence and Statistics*, pages 5485–5517. PMLR, 2022.
- Junyu Zhang, Mingyi Hong, and Shuzhong Zhang. On lower iteration complexity bounds for the convex concave saddle point problems. *Mathematical Programming*, 194(1-2):901–935, 2022.
- Yuchen Zhang and Lin Xiao. Stochastic primal-dual coordinate method for regularized empirical risk minimization. *Journal of Machine Learning Research*, 18:1–42, 2017.