

Accelerated Saddle Flow Dynamics for Bilinearly Coupled Minimax Problems

Yingzhu Liu, Enrique Mallada, Zhongkui Li, and Pengcheng You

Abstract—Minimax problems have attracted much attention due to various applications in constrained optimization problems and zero-sum games. Identifying saddle points within these problems is crucial, and saddle flow dynamics offer a straightforward yet useful approach. This study focuses on a class of bilinearly coupled minimax problems with strongly convex-linear objective functions. We design an accelerated algorithm based on saddle flow dynamics, achieving a convergence rate beyond the stereotype limit (the strong convexity constant). The algorithm is derived from a sequential two-step transformation of a given objective function. First, a change of variables is applied to render the objective function better-conditioned, introducing strong concavity (from linearity) while preserving strong convexity. Second, proximal regularization, when staggered with the first step, further enhances the strong convexity of the objective function by shifting some of the obtained strong concavity. After these transformations, saddle flow dynamics based on the new objective function can be tuned for accelerated exponential convergence. Besides, such an approach can be extended to weakly convex-weakly concave functions and still guarantees exponential convergence to one stationary point. The theory is verified by a numerical test on an affine equality-constrained convex optimization problem.

I. INTRODUCTION

The study of saddle flow dynamics traces back to the seminal work [1] and contributes to the fundamentals of mathematical optimization from a dynamical system perspective. Saddle flow dynamics aim to search for minimax saddle points by combining gradient descent with gradient ascent on two respective subsets of variables (that form a partition). This approach is the basis of primal-dual methods used for solving constrained optimization problems [2] and best-response dynamics used for locating Nash equilibria in zero-sum games [3], [4]. As a result, saddle flow dynamics have been widely used for resources allocation and stabilizing controller design in a variety of areas, including power systems [5], [6], communication networks [7], [8], and cloud computing [9].

To fully understand the dynamic behavior of saddle flows (and its variants), the convergence properties have been extensively studied. Preliminary results [10]–[13] are centered

on asymptotic stability, primarily using advanced analytical tools to provide insights into the important special case of primal-dual dynamics. More recent efforts [14]–[16] have been made to study the exponential convergence of saddle flow dynamics for not only theory development but also practical uses. For example, [14], [15] have shown particular exponential convergence in the absence of strong convexity. [2] has further extended the study to discrete-time problems, showing linear convergence – the counterpart of exponential convergence in continuous time.

A fundamental challenge in studying these exponentially convergent dynamics is to estimate their decay rates using several algorithmic variants [17]–[23]. A vast amount of literature resorts to proximal regularization, especially when handling non-smooth functions. In particular, [17] only shows the existence of exponential convergence rates related to strong convexity but fails to provide an estimate. Lower bounds on the convergence rate have been developed in [18], [19] and [20], respectively, using saddle flow variants and frequency-domain Integral Quadratic Constraint (IQC) approaches. [21] further establishes a best-case upper-bound rate with a Lyapunov method in the time domain. Another set of results [22], [23] develop a novel projection on the standard Lagrangian and derive lower bounds on the decay rates when proving (semi-)global exponential stability for the augmented primal-dual dynamics. By and large, the rate estimates provided in the literature depend on strong convexity, regularization coefficients, and singular values of any coupling matrices. Notably, the constant of strong convexity seems to be a bottleneck inherent in the convergence rate of saddle flow dynamics.

In this paper, we focus particularly on a class of bilinearly coupled minimax problems with strongly convex-linear objective functions and design an algorithm based on carefully designed saddle flow dynamics that exploits the problem structure and achieves an exponential convergence rate beyond the strong convexity constant. The algorithm involves two sequential steps - change-of-variable conditioning and proximal regularization - that jointly enhance the strong convexity-strong concavity of a given objective function. Building upon our recent results [24], this property immediately results in a lower-bound estimate for the convergence rate that breaks the stereotype limit. In addition, such a performance guarantee makes the algorithm suited for a class of weakly convex-weakly concave functions with exponential convergence to a stationary point. Extensive simulations are run to compare our algorithm with existing methods and the actual rate with its lower-bound rate estimate.

This work was supported in part by the National Science and Technology Major Project under grant 2022ZD0116401, and in part by the National Natural Science Foundation of China under grants 72201007, 72431001, T2121002, 72131001, U2241214, 62373008, and 723B1001.

Y. Liu, Z. Li, and P. You are with the College of Engineering, Peking University, Beijing, China. Y. Liu and Z. Li are also with the State Key Laboratory for Turbulence and Complex Systems, Peking University, Beijing, China. {yzliucoe, zhongkli, pcyou}@pku.edu.cn

E. Mallada is with the Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD, USA. mallada@jhu.edu

Our contributions are summarized as follows.

- (i) For bilinearly coupled minimax problems with strongly convex-linear objective functions, we design a saddle flow-based algorithm with an exponential convergence rate beyond the given strong convexity constant.
- (ii) We provide an explicit lower-bound estimate for the exponential convergence rate of the algorithm, which offers a guideline for choosing parameters to further optimize the rate.
- (iii) The accelerated rate guarantee accommodates weakly convex-weakly concave functions, thus showing exponential convergence in nonconvex-nonconcave scenarios.
- (iv) Numerical results validate the proposed algorithm's superiority to many existing methods and further show that the lower-bound rate estimate is almost tight.

Notation: Let $\mathbb{R}$ be the set of real numbers. $I_n \in \mathbb{R}^{n \times n}$ denotes the identity matrix of size n . Given a twice differentiable function $L(x, y) \in \mathcal{C}^2$ with $L : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}$, we use $\frac{\partial}{\partial x} L(x, y) \in \mathbb{R}^{1 \times n}$ and $\frac{\partial}{\partial y} L(x, y) \in \mathbb{R}^{1 \times m}$ to denote the partial derivatives with respect to x and y , respectively. We further define $\nabla_x L(x, y) := [\frac{\partial}{\partial x} L(x, y)]^T$. Meanwhile, $\frac{\partial^2}{\partial x^2} L(x, y) \in \mathbb{R}^{n \times n}$ and $\frac{\partial^2}{\partial y^2} L(x, y) \in \mathbb{R}^{m \times m}$ represent the second-order partial derivatives of $L(x, y)$ with respect to x and y , respectively.

II. PROBLEM AND RESULT

A. Problem Statement

In this paper, we aim to solve bilinearly coupled minimax problems of the form:

$$\min_{x \in \mathbb{R}^n} \max_{y \in \mathbb{R}^m} L(x, y) := f(x) + \eta y^T A x - \eta y^T b \quad (1)$$

where $\eta > 0$ denotes a constant, $b \in \mathbb{R}^m$ is a constant vector, and $A \in \mathbb{R}^{m \times n}$ is the coupling matrix. $f(x) : \mathbb{R}^n \mapsto \mathbb{R}$ is assumed to be twice differentiable. We assume a saddle point (x_*, y_*) exists. The minimax problem (1) is equivalent to an affine equality-constrained optimization problem:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & f(x) \\ \text{s.t.} \quad & A x - b = 0 : y \in \mathbb{R}^m \end{aligned} \quad (2)$$

where $y \in \mathbb{R}^m$ represents the dual variable. Note that $L(x, y)$ is the (weighted) Lagrangian of the problem (2) and $(x_*, \eta y_*)$ is an optimal primal-dual solution.

The primary goal is to locate a saddle (min-max) point of $L(x, y)$, as indicated in (1). We will particularly focus on saddle flow dynamics that run a continuous-time version of gradient descent ascent on $L(x, y)$:

Definition 1 (Saddle Flow Dynamics). *The saddle flow dynamics on $L(x, y)$ are defined as the following dynamic law:*

$$\dot{x} = -\nabla_x L(x, y), \quad (3a)$$

$$\dot{y} = +\nabla_y L(x, y). \quad (3b)$$

In the following subsections, we will develop an algorithm based on saddle flow dynamics that exponentially converges to a saddle point of $L(x, y)$.

B. Algorithm

Our algorithm is defined on a set of new auxiliary variables. Three auxiliary variables $u \in \mathbb{R}^n$, $p, v \in \mathbb{R}^m$ are introduced, and their connection with the original variables x and y is briefly illustrated in Figure 1.




Fig. 1: Algorithm development.

First, we propose a change of variables to transform (x, y) into (u, p) dictated by:

$$\begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} I_n & \alpha A^T \\ 0 & B \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}, \quad (4)$$

where $\alpha > 0$ is a positive constant, and $B \in \mathbb{R}^{m \times m}$ is a given full-rank transformation matrix, with $\bar{\sigma}_B I \succeq B B^T \succeq \underline{\sigma}_B I$ for some $\bar{\sigma}_B > 0$ and $\underline{\sigma}_B > 0$. We obtain $\bar{L}(u, p)$ as the counterpart of the Lagrangian $L(x, y)$ in the (u, p) space:

$$\begin{aligned} \bar{L}(u, p) := & f(u - \alpha A^T B^{-1} p) + \eta p^T B^{-T} (A u - b) \\ & - \eta \alpha \|A^T B^{-1} p\|^2. \end{aligned} \quad (5)$$

This conditioning step exploits the linearity in y . Second, we introduce a proximal regularizer on the variable p to obtain a Moreau envelope function defined on (u, v) in the following form:

$$\tilde{L}(u, v) := \max_{p \in \mathbb{R}^m} \left\{ \bar{L}(u, p) - \frac{\rho}{2} \|p - v\|^2 \right\}, \quad (6)$$

where $\rho > 0$ denotes a constant regularization coefficient. The saddle flow dynamics on the Moreau envelope are expected to reach an equilibrium point (u_*, v_*) , which immediately implies a saddle point (x_*, y_*) of the original Lagrangian $L(x, y)$ via the inverse change of variables.

More specifically, we define our algorithm as the following dynamic law on (u, v) :

$$\dot{u} = -\nabla_u \tilde{L}(u, v) = -\nabla f(u - \alpha A^T B^{-1} \tilde{p}) - \eta A^T B^{-1} \tilde{p}, \quad (7a)$$

$$\dot{v} = +\nabla_v \tilde{L}(u, v) = \rho(\tilde{p} - v), \quad (7b)$$

where $\tilde{p}$ is a shorthand for the mapping $\tilde{p}(u, v) : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}^m$ that represents the unique maximizer in (6) (uniqueness to be shown later) subject to

$$\begin{aligned} -\alpha B^{-T} A \nabla f(u - \alpha A^T B^{-1} \tilde{p}(u, v)) + \eta B^{-T} (A u - b) \\ - 2\eta \alpha B^{-T} A A^T B^{-1} \tilde{p}(u, v) - \rho(\tilde{p}(u, v) - v) = 0. \end{aligned} \quad (8)$$

We will provide the conditions under which the algorithm (7) is guaranteed to achieve the goal and analyze its performance in the next subsection.

C. Main Result

Before stating the main result, we outline the assumptions required for the algorithm (7) that basically characterize the objective function $f(x)$ and the matrix A of the affine equality constraint in (2).

Assumption 1. $f(x) \in \mathcal{C}^2$ is m_f -strongly convex and l_f -smooth, i.e., $l_f I \preceq \nabla^2 f(x) \preceq m_f I$ with $l_f > 0$ and $m_f > 0$.

Given Assumption 1, $L(x, y)$ is a strongly convex-linear function.

Assumption 2. The coupling matrix A is full row rank with $\bar{\sigma}_A I \preceq AA^T \preceq \underline{\sigma}_A I$ for some $\bar{\sigma}_A > 0$ and $\underline{\sigma}_A > 0$.

Given Assumptions 1 and 2, the saddle point (x_*, y_*) of $L(x, y)$ is unique since the problem (2) meets the linear independence constraint qualification and admits a unique primal-dual solution [25], [26]. Note that the saddle point (x_*, y_*) of $L(x, y)$ is also the equilibrium point of the corresponding saddle flow dynamics (3).

Our first argument is that an equilibrium point (u_*, v_*) of the algorithm (7) is immediately related to the target saddle point (x_*, y_*) , as the following theorem suggests.

Theorem 1 (Characterization of Saddle Point). *Let Assumptions 1 and 2 hold. Given an equilibrium point (u_*, v_*) of the dynamic law (7), the unique saddle point (x_*, y_*) of $L(x, y)$ can be attained by*

$$x_* = u_* - \alpha A^T B^{-1} v_*, \quad (9a)$$

$$y_* = B^{-1} v_*. \quad (9b)$$

The proof of Theorem 1 is provided in Appendix A, which is an immediate consequence of the algorithm design in Section III.

Given Theorem 1, it remains to show the convergence of the algorithm (7) to an equilibrium point, which we establish along with a lower-bound rate estimate in the following theorem.

Theorem 2 (Convergence). *Let Assumptions 1 and 2 hold. Given any positive constants η and α that satisfy $2\eta > \alpha l_f$, the dynamic law (7) globally exponentially converges to a unique equilibrium point (u_*, v_*) . More precisely, given $w := (u, v)$,*

$$\|w(t) - w_*\| \leq \|w(0) - w_*\| e^{-ct}$$

holds with the lower-bound rate

$$c := \min \left\{ m_f + \frac{\underline{\sigma}_B}{\bar{\sigma}_B} \frac{\underline{\sigma}_A \gamma_1}{(2\eta\alpha - \alpha^2 m_f) \bar{\sigma}_A + \rho \underline{\sigma}_B}, \frac{\rho \underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f)}{\underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f) + \bar{\sigma}_B \rho} \right\},$$

where γ_1 is a constant with

$$0 \leq \gamma_1 \leq \min\{(\eta - \alpha m_f)^2, (\eta - \alpha l_f)^2\}.$$

The proof will become clear as we develop the algorithm (7) in detail in Section III.

Theorem 2 not only guarantees the exponential convergence of the algorithm (7), but also implies the following corollary that suggests an accelerated rate beyond the constant of the strong convexity of $f(x)$.

Corollary 3. *Given any positive constants ρ , η , $\bar{\sigma}_B$ and $\underline{\sigma}_B$ that satisfy*

$$1) \bar{\sigma}_B m_f < \underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f),$$

$$2) \rho > \frac{\underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f)}{\underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f) - \bar{\sigma}_B m_f} m_f,$$

the dynamical law (7) achieves a convergence rate lower bounded by

$$c > m_f.$$

In Corollary 3, the first condition ensures that the right-hand side of the inequality in the second condition is positive. Technically, the second condition can be derived from the expression for the rate c . The intuition will be provided later in Section III.

Remark 1. *The parameters that satisfy all the conditions always exist. A straightforward way is to pick a sufficiently small $\bar{\sigma}_B$ and a sufficiently large ρ . For example, the following set of parameters*

$$\eta = 8, \alpha = \frac{1}{l_f}, \bar{\sigma}_B = \frac{\underline{\sigma}_A}{m_f l_f}, \rho = \frac{3}{2} m_f \quad (10)$$

satisfy all the conditions.

III. ALGORITHM DEVELOPMENT AND ANALYSIS

In this section, we explain the development of the algorithm, which naturally provides insights into the accelerated convergence performance of the algorithm (7). Inspired by [27] and [28], we aim to enhance the convexity-concavity of the original Lagrangian in (1) by exploiting the bilinear couplings.

A. Change-of-Variable Conditioning

Recall the change of variables (4) in the first step. We obtain the expression for the Lagrangian $L(x, y)$ in the (u, p) space as $\bar{L}(u, p)$ in (5). The role of the change-of-variable conditioning (4) is captured below.

Lemma 4. *Let Assumptions 1 and 2 hold. Given any positive constants η and α that satisfy $2\eta > \alpha l_f$, $\bar{L}(u, p)$ is m_f -strongly convex in u and $\frac{\underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f)}{\bar{\sigma}_B}$ -strongly concave in p .*

Proof. Lemma 4 follows immediately from the calculation of the respective second-order partial derivatives:

$$\frac{\partial^2}{\partial u^2} \bar{L}(u, p) = \nabla^2 f(u - \alpha A^T B^{-1} p) \succeq m_f I_n,$$

$$\begin{aligned} \frac{\partial^2}{\partial p^2} \bar{L}(u, p) &= \alpha^2 B^{-T} A \nabla^2 f(u - \alpha A^T B^{-1} p) A^T B^{-1} \\ &\quad - 2\eta \alpha B^{-T} A A^T B^{-1} \\ &= -B^{-T} A \left(2\eta \alpha I_n - \alpha^2 \nabla^2 f(u - \alpha A^T B^{-1} p) \right) A^T B^{-1} \\ &\preceq -B^{-T} A (2\eta \alpha - \alpha^2 l_f) I_n A^T B^{-1} \preceq -\frac{\underline{\sigma}_A}{\bar{\sigma}_B} (2\eta \alpha - \alpha^2 l_f) I_m. \end{aligned}$$

■

Lemma 4 suggests that the change-of-variable conditioning enhances the concavity (from linearity) while maintaining the strong convexity for the objective function. In fact, the constant of strong concavity of $\bar{L}(u, p)$ in p could be further optimized with proper choices of parameters. See an example below.

Remark 2. When first fixing η and $\bar{\sigma}_B$, setting $\alpha = \eta/l_f$ gives the largest possible constant of strong concavity $\frac{\eta^2 \underline{\sigma}_A}{\bar{\sigma}_B l_f}$. In this case, $\frac{\eta^2}{\bar{\sigma}_B} \geq \frac{m_f l_f}{\underline{\sigma}_A}$ suffices to guarantee

$$\frac{\partial^2}{\partial u^2} \bar{L}(u, p) \succeq m_f I_n, \quad \frac{\partial^2}{\partial p^2} \bar{L}(u, p) \preceq -m_f I_m.$$

Furthermore, $\bar{L}(u, p)$ is relevant since its unique saddle point (u_*, p_*) is directly connected with the saddle point (x_*, y_*) of the original Lagrangian $L(x, y)$, as the corollary states.

Corollary 5. (u_*, p_*) is the unique saddle point of $\bar{L}(u, p)$ if and only if (x_*, y_*) is the unique saddle point of $L(x, y)$ with

$$x_* = u_* - \alpha A^T B^{-1} p_*, \quad (11a)$$

$$y_* = B^{-1} p_*. \quad (11b)$$

B. Proximal Regularization

This subsection builds on $\bar{L}(u, p)$ and aims to further improve its strong convexity-strong concavity. As the second step, we apply proximal regularization on p to obtain the Moreau envelope of $\bar{L}(u, p)$:

$$\begin{aligned} \tilde{L}(u, v) &= \max_{p \in \mathbb{R}^m} \left\{ \bar{L}(u, p) - \frac{\rho}{2} \|p - v\|^2 \right\} \\ &= f(u - \alpha A^T B^{-1} \tilde{p}) + \eta \tilde{p}^T B^{-T} (Au - b) \\ &\quad - \eta \alpha \|A^T B^{-1} \tilde{p}\|^2 - \frac{\rho}{2} \|\tilde{p} - v\|^2, \end{aligned} \quad (12)$$

where $\tilde{p}$ is the solution to the optimality condition (8). Note that $\tilde{p}$ is unique due to the strong concavity of $\bar{L}(u, p)$ in p . It can be efficiently computed, either using its closed form when it is available, i.e., the inverse function of $\nabla f(\cdot)$ has an explicit expression, or using numerical methods since $\nabla f(\cdot)$ is strictly monotone.

The following lemma characterizes the impact of proximal regularization on the convexity-concavity property.

Lemma 6. Let Assumptions 1 and 2 hold. Given any positive constants η and α that satisfy $2\eta > \alpha l_f$, $\tilde{L}(u, v)$ is c_1 -strongly convex in u and c_2 -strongly concave in v with

$$c_1 := m_f + \frac{\underline{\sigma}_B}{\bar{\sigma}_B} \frac{\underline{\sigma}_A \gamma_1}{(2\eta\alpha - \alpha^2 m_f) \bar{\sigma}_A + \rho \underline{\sigma}_B},$$

$$c_2 := \rho \frac{\underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f)}{\underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f) + \bar{\sigma}_B \rho}.$$

Proof. Since $L(x, y)$ is twice differentiable, we are allowed to derive the partial derivatives of $\tilde{L}(u, v)$ as

$$\frac{\partial}{\partial u} \tilde{L}(u, v) = \frac{\partial}{\partial u} f(u - \alpha A^T B^{-1} \tilde{p}) + \eta \tilde{p}^T B^{-T} A,$$

$$\frac{\partial}{\partial v} \tilde{L}(u, v) = \rho (\tilde{p} - v).$$

Using Danskin's Theorem, we can further compute the second-order partial derivatives. Define

$$H := \nabla^2 f(u - \alpha A^T B^{-1} \tilde{p}),$$

which is a symmetric matrix. Then we obtain

$$\begin{aligned} \frac{\partial^2}{\partial u^2} \tilde{L}(u, v) &= H + (\eta I_n - \alpha H) A^T B^{-1} \mathbf{J}_p^u, \\ \frac{\partial^2}{\partial v^2} \tilde{L}(u, v) &= \rho (\mathbf{J}_p^v - I_m), \end{aligned}$$

where $\mathbf{J}_p^u$ and $\mathbf{J}_p^v$ are the partial derivatives of $\tilde{p}(u, v)$ with respect to u and v , respectively. By taking partial derivatives of both sides of the optimality condition (8), we have

$$\begin{aligned} \mathbf{J}_p^u &= \mathbf{J}_c^{-1} B^T A (\eta I_n - \alpha H), \\ \mathbf{J}_p^v &= \rho \mathbf{J}_c^{-1}, \end{aligned}$$

with $\mathbf{J}_c := B^{-T} A (2\eta\alpha I_n - \alpha^2 H) A^T B^{-1} + \rho I_m$.

The explicit second-order partial derivatives from above allow us to characterize the strong convexity-strong concavity of $\tilde{L}(u, v)$:

$$\begin{aligned} \frac{\partial^2}{\partial u^2} \tilde{L}(u, v) &= H + (\eta I_n - \alpha H) A^T B^{-1} \mathbf{J}_c^{-1} B^T A (\eta I_n - \alpha H) \\ &\succeq m_f I + \gamma_1 A^T B^{-1} \mathbf{J}_c^{-1} B^T A \\ &\succeq \left(m_f + \frac{\underline{\sigma}_B}{\bar{\sigma}_B} \frac{\underline{\sigma}_A \gamma_1}{(2\eta\alpha - \alpha^2 m_f) \bar{\sigma}_A + \rho \underline{\sigma}_B} \right) I_n \succ 0, \end{aligned}$$

$$\begin{aligned} \frac{\partial^2}{\partial v^2} \tilde{L}(u, v) &= \rho (\mathbf{J}_p^v - I_m) \\ &= \rho (\rho \mathbf{J}_c^{-1} - I_m) \\ &\preceq -\rho \frac{\underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f)}{\underline{\sigma}_A (2\eta\alpha - \alpha^2 l_f) + \bar{\sigma}_B \rho} I_m \prec 0. \end{aligned}$$

■

Compared with the strong convexity-strong concavity of $\bar{L}(u, p)$ identified in Lemma 4, Lemma 6 reveals an important role of proximal regularization here: it increases the constant of strong convexity while decreasing the constant of strong concavity.

Similarly, we can establish the immediate connection between the saddle points of $\bar{L}(u, p)$ and $\tilde{L}(u, v)$ as follows.

Corollary 7. (u_*, v_*) is the unique saddle point of $\tilde{L}(u, p)$ if and only if (u_*, p_*) is the unique saddle point of $\bar{L}(u, p)$ with $p_* = v_*$.

Note that given Lemmas 4 and 6 as well as Corollaries 5 and 7, Theorem 1 becomes straightforward. We summarize the brief proof in Appendix A.

C. Convergence Analysis

Now it becomes clear that the algorithm (7) is essentially the saddle flow dynamics on $\tilde{L}(u, v)$. The exponential convergence of the algorithm established in Theorem 2 could be built upon our recent result [24], which we briefly present here for completeness.

Lemma 8. Given a twice differentiable function $S(x, y)$ with m -strong convexity in x and q -strong concavity in y , the saddle flow dynamics on $S(x, y)$:

$$\dot{z} = F(z) = \begin{bmatrix} -\nabla_x S(x, y) \\ \nabla_y S(x, y) \end{bmatrix}, \quad (16)$$

with $z := (x, y)$, are globally exponentially stable. More precisely,

$$\|z(t) - z_\star\| \leq \|z(0) - z_\star\| e^{-rt}$$

holds with the lower-bound rate

$$r := \min\{m, q\}.$$

The proof of Lemma 8 is available in [24].

In virtue of Lemma 8, the strong convexity-strong concavity of $\bar{L}(u, v)$ in Lemma 6 immediately implies Theorem 2.

Rate Optimization: According to the parameter conditions in Corollary 3, there exist two constants $0 < \zeta < \beta < 1$ subject to

$$\bar{\sigma}_B = \beta \underline{\sigma}_A \frac{(2\eta\alpha - \alpha^2 l_f)}{m_f}, \quad \underline{\sigma}_B = \zeta \underline{\sigma}_A \frac{(2\eta\alpha - \alpha^2 l_f)}{m_f}.$$

Then the strong convexity and strong concavity constants of $\bar{L}(u, v)$ can be rewritten as

$$c_1 = \left(1 + \frac{\zeta}{\beta} \frac{\gamma_1}{(2\eta\alpha - \alpha^2 m_f) \kappa_A m_f + (2\eta\alpha - \alpha^2 l_f) \rho \zeta}\right) m_f, \quad (17)$$

$$c_2 = \frac{\rho m_f}{\rho \beta + m_f} = \left(1 + \frac{(1 - \beta)\rho - m_f}{\rho \beta + m_f}\right) m_f, \quad (18)$$

where $\kappa_A := \bar{\sigma}_A / \underline{\sigma}_A$ is the condition number of AA^T . Since c_1 decreases while c_2 increases, both monotonically in ρ , and there is bound to be one crossing point (given the respective ranges), we can obtain the fastest convergence rate in terms of ρ by setting $c_1 = c_2$. It will be achieved at $\rho_\star$ that satisfies

$$(\beta - 1)\beta\zeta L \rho_\star^2 + m_f(\zeta\gamma_1 + M\beta m_f \kappa_A) + \rho_\star \beta (\zeta\gamma_1 + m_f (L\zeta + M(\beta - 1)\kappa_A)) = 0,$$

with $L := 2\eta\alpha - \alpha^2 l_f$ and $M := 2\eta\alpha - \alpha^2 m_f$. Further optimization over the other parameters is intricate and will be left for future studies.

IV. EXTENSION TO NONCONVEX-NONCONCAVE MINIMAX PROBLEM

In this section, we study a broader class of bilinearly coupled minimax problems:

$$\min_{x \in \mathbb{R}^n} \max_{y \in \mathbb{R}^m} L(x, y) := f(x) + y^T A x - g(y) \quad (19)$$

where $A \in \mathbb{R}^{m \times n}$ is the coupling matrix while $f(x) : \mathbb{R}^n \mapsto \mathbb{R}$ and $g(y) : \mathbb{R}^m \mapsto \mathbb{R}$ are assumed to be twice differentiable, but not necessarily convex. We assume a stationary point exists.

Searching for saddle/stationary points of nonconvex-nonconcave functions is a popular goal in the machine learning community where such objective functions are common [29], [30]. While in general locating or even approximating stationary points is hard, certain special structures

of objective functions, e.g., bilinear couplings, make this problem still tractable [28], [31], [32].

In this section, we extend the results of this paper to solve the problem (19), in the case where $L(x, y)$ is weakly convex-weakly concave and smooth.

Assumption 3. $f(x) \in \mathcal{C}^2$ is m_f -weakly convex and l_f -smooth, i.e., $l_f I \succeq \nabla^2 f(x) \succeq -m_f I$ with $l_f > 0$ and $m_f > 0$. Similarly, $g(y) \in \mathcal{C}^2$ is m_g -weakly convex and l_g -smooth, i.e., $l_g I \succeq \nabla^2 g(y) \succeq -m_g I$ with $l_g > 0$ and $m_g > 0$.

Remark 3. Note that, the objective function $L(x, y)$ is convex-concave when $m_f = m_g = 0$ holds and strongly convex-strongly concave when $m_f < 0$ and $m_g < 0$ hold. This study specifically addresses the scenarios in which $m_f > 0$ and $m_g > 0$ hold, thereby quantifying the degree of nonconvexity-nonconcavity of $L(x, y)$.

In what follows we use a pipeline similar to that in Sections II and III, highlighting mostly the main changes due to Assumption 3. As before, applying the change of variables (4), we obtain the corresponding objective function $\bar{L}(u, p)$ in the (u, p) space:

$$\bar{L}(u, p) = f(u - \alpha A^T B^{-1} p) + p^T B^{-T} A u - \alpha \|A^T B^{-1} p\|^2 - g(B^{-1} p). \quad (20)$$

Even though in general $\bar{L}(u, p)$ is not necessarily concave in p , one can compute the second-order partial derivatives (similar to the procedure to prove Lemma 4) and show that $\bar{L}(u, p)$ is $\frac{\underline{\sigma}_A(2\alpha - \alpha^2 l_f) - m_g}{\bar{\sigma}_B}$ -strongly concave in p whenever $\underline{\sigma}_A(2\alpha - \alpha^2 l_f) - m_g > 0$ holds. This guarantees the existence and uniqueness of the following Moreau envelop:

$$\begin{aligned} \tilde{L}(u, v) &= \max_{p \in \mathbb{R}^m} \left\{ \bar{L}(u, p) - \frac{\rho}{2} \|p - v\|^2 \right\} \\ &= f(u - \alpha A^T B^{-1} \tilde{p}^\star) + (\tilde{p}^\star)^T B^{-T} A u \\ &\quad - \alpha \|A^T B^{-1} \tilde{p}^\star\|^2 - g(B^{-1} \tilde{p}^\star) - \frac{\rho}{2} \|\tilde{p}^\star - v\|^2, \end{aligned} \quad (21)$$

where $\tilde{p}^\star$, a shorthand for $\tilde{p}^\star(u, v)$, is the unique maximizer in (21).

We will show shortly that under mild conditions, $\tilde{L}(u, v)$ in (21) is strongly convex-strongly concave, and thus the corresponding saddle flow dynamics are guaranteed to convergence to a unique stationary/saddle point $(u_\star, v_\star)$. The following theorem relates such a point with a stationary point of $L(x, y)$.

Theorem 9. If $(u_\star, v_\star)$ is a stationary point of $\tilde{L}(u, v)$, then $(u_\star, p_\star = v_\star)$ is a stationary point of $\bar{L}(u, p)$. Moreover, $(x_\star, y_\star)$, as described by (9), is a stationary point of $L(x, y)$.

The proof is provided in Appendix B.

Theorem 9 guarantees that if the saddle flow dynamics on $\tilde{L}(u, v)$, i.e.,

$$\dot{u} = -\nabla f(u - \alpha A^T B^{-1} \tilde{p}^\star) - A^T B^{-1} \tilde{p}^\star, \quad (22a)$$

$$\dot{v} = \rho(\tilde{p}^\star - v), \quad (22b)$$

converge, then the corresponding $(x_\star, y_\star)$ will be a stationary point of $L(x, y)$.




Fig. 2: Results for a strongly convex-linear minimax problem.

We finalize this section by providing a convergence guarantee for (22).

Theorem 10. *Let Assumptions 2 and 3 hold. Given any positive constants α , ρ , $\bar{\sigma}_B$ and $\underline{\sigma}_B$ that satisfy*

- 1) $\underline{\sigma}_A(2\alpha - \alpha^2 l_f) - m_g > 0$ and $\alpha l_f \neq 1$,
- 2) $\kappa_A := \frac{\bar{\sigma}_A}{\underline{\sigma}_A} < \frac{\gamma_2}{2\alpha + \alpha^2 m_f}$,
- 3) $\kappa_B := \frac{\bar{\sigma}_B}{\underline{\sigma}_B} < \frac{\underline{\sigma}_A \gamma_2}{l_g} - \frac{\bar{\sigma}_A(2\alpha + \alpha^2 m_f)}{l_g}$,
- 4) $0 < \rho < \frac{\underline{\sigma}_A \gamma_2}{\bar{\sigma}_B m_f} - \frac{\bar{\sigma}_A(2\alpha + \alpha^2 m_f)}{\bar{\sigma}_B} - \frac{l_g}{\underline{\sigma}_B}$,

where γ_2 is a constant with

$$0 < \gamma_2 \leq \min\{(1 + \alpha m_f)^2, (1 - \alpha l_f)^2\},$$

then the dynamic law (22) exponentially converges to a unique equilibrium point (u_*, v_*) . More precisely, given $w := (u, v)$,

$$\|w(t) - w_*\| \leq \|w(0) - w_*\| e^{-ct}$$

holds with the lower-bound rate

$$c := \min \left\{ -m_f + \frac{\underline{\sigma}_B}{\bar{\sigma}_B} \frac{\underline{\sigma}_A \gamma_2}{\bar{\sigma}_A(2\alpha + \alpha^2 m_f) + \rho \underline{\sigma}_B + l_g}, \right. \\ \left. \rho \frac{\underline{\sigma}_A(2\alpha - \alpha^2 l_f) - m_g}{\underline{\sigma}_A(2\alpha - \alpha^2 l_f) + \rho \bar{\sigma}_B - m_g} \right\} > 0.$$

The proof is similar to that of Theorem 2.

Remark 4. *Since the original objective function $L(x, y)$ is nonconvex-nonconcave, there might be multiple stationary points. However, $\tilde{L}(u, v)$ becomes strongly convex-strongly concave and has a unique saddle point. As a result, our algorithm only obtains one particular stationary point of $L(x, y)$.*

V. SIMULATION RESULTS

In this section, we numerically validate our algorithm (7) to solve the strongly convex-linear minimax problem (1) and compare it with some existing algorithms in the literature. Precisely, we set $n = 5$ and $m = 4$, and adopt a quadratic objective function $f(x) = \frac{1}{2} x^T Q x$. Here $Q \in \mathbb{R}^{5 \times 5}$ is

given by $Q = 5I_5 + Q_0^T Q_0$, and $Q_0 \in \mathbb{R}^{5 \times 5}$ is generated from a Gaussian random matrix. Similarly, $A \in \mathbb{R}^{4 \times 5}$ and $b \in \mathbb{R}^4$ are also randomly generated. In this particular test run, we have the strong convexity constant $m_f = 0.5$ and the smoothness constant $l_f = 0.693$. The rest of the parameters are set according to (10) and the initial point is chosen arbitrarily.

The convergence results are shown in Fig. 2. It can be observed that the trajectories of the saddle flow dynamics converge rapidly to an equilibrium point in Fig. 2(a). The distance to the equilibrium point exponentially decays, as shown in Fig. 2(b). Our lower-bound estimate is close to the actual rate and almost tight. Moreover, our method is superior to many existing algorithms, including the proximal method [20] and the primal-dual gradient dynamics (PDGD) in [23], the convergence rates of which are limited by the strong convexity constant m_f .

VI. CONCLUSION

This paper studies bilinearly coupled minimax problems and proposes an accelerated algorithm to locate saddle points for strongly convex-linear objective functions. The algorithm implements saddle flow dynamics on a transformed objective function, obtained from two sequential steps: change-of-variable conditioning and proximal regularization. The first step exploits the bilinear couplings to enhance the concavity while maintaining the strong convexity for the objective function. The second step then shifts some of the strong concavity to enhance the strong convexity. We propose the conditions under which our algorithm can achieve a convergence rate beyond the stereotype limit (the strong convexity constant). The same design strategy can also be applied to construct a stationary point searching algorithm for a general class of weakly convex-weakly concave functions.

REFERENCES

- [1] K. J. Arrow, L. Hurwicz, and H. Uzawa, *Studies in Linear and Non-Linear Programming*. Stanford University Press, 1982.

- [2] D. Ding, B. Hu, N. K. Dhirga, and M. R. Jovanović, “An exponentially convergent primal-dual algorithm for nonsmooth composite minimization,” in *IEEE Conference on Decision and Control (CDC)*, 2018, pp. 4927–4932.
- [3] L. Baudin and R. Laraki, “Fictitious play and best-response dynamics in identical interest and zero-sum stochastic games,” in *International Conference on Machine Learning (ICML)*, vol. 162, 2022, pp. 1664–1690.
- [4] B. Ghahesifard and J. Cortés, “Distributed convergence to Nash equilibria in two-network zero-sum games,” *Automatica*, vol. 49, no. 6, pp. 1683–1692, 2013.
- [5] C. Zhao, U. Topcu, N. Li, and S. Low, “Design and stability of load-side primary frequency control in power systems,” *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1177–1189, 2014.
- [6] X. Zhang, N. Li, and A. Papachristodoulou, “Achieving real-time economic dispatch in power networks via a saddle point design approach,” in *IEEE Power & Energy Society General Meeting*, 2015, pp. 1–5.
- [7] F. Paganini and E. Mallada, “A unified approach to congestion control and node-based multipath routing,” *IEEE/ACM Transactions on Networking*, vol. 17, no. 5, pp. 1413–1426, 2009.
- [8] T. Holding and I. Lestas, “Stability and instability in primal-dual algorithms for multi-path routing,” in *IEEE Conference on Decision and Control (CDC)*, 2015, pp. 1651–1656.
- [9] D. Goldszajn and F. Paganini, “Proximal regularization for the saddle point gradient dynamics,” *IEEE Transactions on Automatic Control*, vol. 66, no. 9, pp. 4385–4392, 2021.
- [10] A. Cherukuri, E. Mallada, and J. Cortés, “Asymptotic convergence of constrained primal-dual dynamics,” *Systems & Control Letters*, vol. 87, pp. 10–15, 2016.
- [11] A. Cherukuri, B. Ghahesifard, and J. Cortés, “Saddle-point dynamics: Conditions for asymptotic stability of saddle points,” *SIAM Journal on Control and Optimization*, vol. 55, no. 1, pp. 486–511, 2017.
- [12] A. Cherukuri, E. Mallada, S. Low, and J. Cortés, “The role of convexity in saddle-point dynamics: Lyapunov function and robustness,” *IEEE Transactions on Automatic Control*, vol. 63, no. 8, pp. 2449–2464, 2017.
- [13] P. You and E. Mallada, “Saddle flow dynamics: Observable certificates and separable regularization,” in *IEEE American Control Conference (ACC)*, 2021, pp. 4817–4823.
- [14] X. Chen and N. Li, “Exponential stability of primal-dual gradient dynamics with non-strong convexity,” in *IEEE American Control Conference (ACC)*, 2020, pp. 1612–1618.
- [15] S. Liang, L. Y. Wang, and G. Yin, “Exponential convergence of distributed primal-dual convex optimization algorithm without strong convexity,” *Automatica*, vol. 105, pp. 298–306, 2019.
- [16] P. A. Bansode, V. Chinde, S. R. Wagh, R. Pasumathy, and N. M. Singh, “On the exponential stability of projected primal-dual dynamics on a Riemannian manifold,” *arXiv:1905.04521*, 2019.
- [17] S. K. Niederländer, F. Allgower, and J. Cortés, “Exponentially fast distributed coordination for nonsmooth convex optimization,” in *IEEE Conference on Decision and Control (CDC)*, 2016, pp. 1036–1041.
- [18] Z. Wang, W. Wei, C. Zhao, Z. Ma, Z. Zheng, Y. Zhang, and F. Liu, “Exponential stability of partial primal-dual gradient dynamics with nonsmooth objective functions,” *Automatica*, vol. 129, p. 109585, 2021.
- [19] J. Cortés and S. K. Niederländer, “Distributed coordination for nonsmooth convex optimization via saddle-point dynamics,” *Journal of Nonlinear Science*, vol. 29, no. 4, pp. 1247–1272, 2019.
- [20] N. K. Dhirga, S. Z. Khong, and M. R. Jovanović, “The proximal augmented lagrangian method for nonsmooth composite optimization,” *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2861–2868, 2019.
- [21] D. Ding and M. R. Jovanović, “Global exponential stability of primal-dual gradient flow dynamics based on the proximal augmented lagrangian: A Lyapunov-based approach,” in *IEEE Conference on Decision and Control (CDC)*, 2020, pp. 4836–4841.
- [22] Y. Tang, G. Qu, and N. Li, “Semi-global exponential stability of augmented primal-dual gradient dynamics for constrained convex optimization,” *Systems & Control Letters*, vol. 144, p. 104754, 2020.
- [23] G. Qu and N. Li, “On the exponential stability of primal-dual gradient dynamics,” *IEEE Control Systems Letters*, vol. 3, no. 1, pp. 43–48, 2019.
- [24] P. You, Y. Liu, and E. Mallada, “A unified analysis of saddle flow dynamics: Stability and algorithm design,” *arXiv:2409.05290*, 2024.
- [25] G. Wachsmuth, “On LICQ and the uniqueness of Lagrange multipliers,” *Operations Research Letters*, vol. 41, no. 1, pp. 78–80, 2013.
- [26] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [27] H. D. Nguyen, T. L. Vu, K. Turitsyn, and J.-J. Slotine, “Contraction and robustness of continuous time primal-dual dynamics,” *IEEE Control Systems Letters*, vol. 2, no. 4, pp. 755–760, 2018.
- [28] B. Grimmer, H. Lu, P. Worah, and V. Mirrokni, “The landscape of the proximal point method for nonconvex–nonconcave minimax optimization,” *Mathematical Programming*, vol. 201, no. 1–2, pp. 373–407, 2023.
- [29] F. Farnia and A. Ozdaglar, “Do GANs always have Nash equilibria?” in *International Conference on Machine Learning (ICML)*, vol. 119, 2020, pp. 3029–3039.
- [30] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 27, 2014.
- [31] J. Diakonikolas, C. Daskalakis, and M. I. Jordan, “Efficient methods for structured nonconvex-nonconcave min-max optimization,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021, pp. 2746–2754.
- [32] J. Yang, N. Kiyavash, and N. He, “Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020, pp. 1153–1165.

APPENDIX

A. Proof of Theorem 1

In Section III, we already have that $\tilde{L}(u, v)$ is strongly convex-strongly concave. As a result, the equilibrium point of the dynamic law (7), i.e., the saddle point of $\tilde{L}(u, v)$, is unique and denoted as (u_*, v_*) . According to Corollary 7, (u_*, v_*) is related to (u_*, p_*) with $v_* = p_*$. Meanwhile, Corollary 5 implies that given (u_*, p_*) , we can obtain the unique saddle point (x_*, y_*) of $L(x, y)$ by (11). Putting everything together proves Theorem 1.

B. Proof of Theorem 9

We need to show

$$\begin{aligned} \nabla_u \tilde{L}(u_*, v_*) = 0 & \Rightarrow \nabla_x L(x_*, y_*) = 0 \\ \nabla_v \tilde{L}(u_*, v_*) = 0 & \Rightarrow \nabla_y L(x_*, y_*) = 0 \end{aligned} \quad (23)$$

with the stationary point (x_*, y_*) satisfying (9). Note that any stationary point (u_*, v_*) of $\tilde{L}(u, v)$ satisfies

$$\begin{aligned} \nabla_u \tilde{L}(u_*, v_*) &= \nabla f(u_* - \alpha A^T B^{-1} \tilde{p}^*(u_*, v_*)) \\ &\quad + A^T B^{-1} \tilde{p}^*(u_*, v_*) = 0, \end{aligned} \quad (24a)$$

$$\nabla_v \tilde{L}(u_*, v_*) = \rho(\tilde{p}^* - v_*) = 0, \quad (24b)$$

where $\tilde{p}^*$ is the unique solution to

$$\begin{aligned} -\alpha B^{-T} A \nabla f(u_* - \alpha A^T B^{-1} \tilde{p}^*) - B^{-T} \nabla g(B^{-1} \tilde{p}^*) \\ - 2\alpha B^{-T} A A^T B^{-1} \tilde{p}^* + B^{-T} A u_* - \rho(\tilde{p}^* - v_*) = 0. \end{aligned} \quad (25)$$

Combining (9), (24) and (25), we can readily show

$$\begin{aligned} \nabla_x L(x_*, y_*) &= \nabla f(x_*) + A^T y_* = 0, \\ \nabla_y L(x_*, y_*) &= A x_* - b = 0. \end{aligned}$$

This completes the proof.