

MIT Open Access Articles

On the geometry and refined rate of primal–dual hybrid gradient for linear programming

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Lu, H., Yang, J. On the geometry and refined rate of primal–dual hybrid gradient for linear programming. Math. Program. (2024).

Published Version: 10.1007/s10107-024-02109-9

Publisher: Springer Science and Business Media LLC

Permanent Link: <https://hdl.handle.net/1721.1/155731>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: <https://creativecommons.org/licenses/by/4.0/>





On the geometry and refined rate of primal–dual hybrid gradient for linear programming

Haihao Lu¹ · Jinwen Yang²

Received: 7 July 2023 / Accepted: 10 June 2024
© The Author(s) 2024

Abstract

We study the convergence behaviors of primal–dual hybrid gradient (PDHG) for solving linear programming (LP). PDHG is the base algorithm of a new general-purpose first-order method LP solver, PDLP, which aims to scale up LP by taking advantage of modern computing architectures. Despite its numerical success, the theoretical understanding of PDHG for LP is still very limited; the previous complexity result relies on the global Hoffman constant of the KKT system, which is known to be very loose and uninformative. In this work, we aim to develop a fundamental understanding of the convergence behaviors of PDHG for LP and to develop a refined complexity rate that does not rely on the global Hoffman constant. We show that there are two major stages of PDHG for LP: in Stage I, PDHG identifies active variables and the length of the first stage is driven by a certain quantity which measures how close the non-degeneracy part of the LP instance is to degeneracy; in Stage II, PDHG effectively solves a homogeneous linear inequality system, and the complexity of the second stage is driven by a well-behaved local sharpness constant of the system. This finding is closely related to the concept of partial smoothness in non-smooth optimization, and it is the first complexity result of finite time identification without the non-degeneracy assumption. An interesting implication of our results is that degeneracy itself does not slow down the convergence of PDHG for LP, but near-degeneracy does.

Mathematics Subject Classification 90C05 (linear programming) · 90C47 (minimax problems)

✉ Haihao Lu
haihao@mit.edu

Jinwen Yang
jinweny@uchicago.edu

¹ Sloan School of Management, MIT, Massachusetts, USA

² Department of Statistics, The University of Chicago, Chicago, USA

1 Introduction

Linear programming (LP) is one of the most fundamental and important classes of optimization problems in operation research and computer science with a vast range of applications, such as network flow, revenue management, transportation, scheduling, packing and covering, and many others [3, 9, 10, 14, 17, 29, 47, 53]. Since the 1940s, LP has been extensively studied in both academia and industry. The state-of-the-art methods to solve LP problems, simplex methods [18] and interior-point methods (IPMs) [34], are quite reliable to provide solutions with high accuracy and they serve as the base algorithms for nowadays commercial LP solvers. The success of both methods depends heavily on the efficient factorization methods to solve linear systems arising in the updates, which makes both algorithms highly challenging to further scale up. There are two fundamental reasons for this: (1) the storage of the factorization is quite memory-demanding and it usually requires significantly more memory than storing the original LP instance; (2) it is highly challenging to take advantage of modern computing architectures such as distributed computing and high-performance computing on GPUs when solving linear systems because factorization is sequential in nature.

Recent applications surge the interest in developing new algorithms for LPs with scale far beyond the capability of simplex methods and IPMs. Matrix-free, i.e., no need to solve linear system, is a central feature of promising candidate algorithms, which guarantees low per-iteration computational cost and being friendly to distributed computation. In this sense, first-order methods (FOMs) have become an attractive solution. The update of FOMs requires only the gradient information and is known for its capability of parallelization, thanks to the recent development in deep learning. In the context of LP, the computational bottleneck of FOMs is merely matrix-vector multiplication which is fairly cheap and exhibits easily distributed implementation, in contrast to solving linear systems for simplex methods and IPMs that is costly and highly nontrivial to parallelize.

One notable instance that exemplifies the effectiveness of such methodology is the recent development of an open-source LP solver PDLP¹ [4]. A distributed version of PDLP has been used to solve real-world LP instances with as many as 92 billion non-zeros in the constraint matrix, which is a hundred to a thousand times larger than the scale state-of-the-art commercial LP solvers can solve [54]. The base algorithm of PDLP is primal–dual hybrid gradient (PDHG) [12], a form of operator splitting method with alternating updates between the primal and dual variables. The implementation of PDLP also involves a few other enhancements/heuristics on top of PDHG, such as restart, preconditioning, adaptive step-size, etc, to further improve the numerical performance [4]. A significant difference between PDLP and other commercial LP solvers is that PDLP is totally matrix-free, namely, there is no necessity for PDLP to solve any linear system. The main computational bottleneck is matrix-vector multiplication which is much cheaper for solving large instances and more tailored for large-scale application in the distributed setting.

¹ The solver is open-sourced at [Google OR-Tools](#).

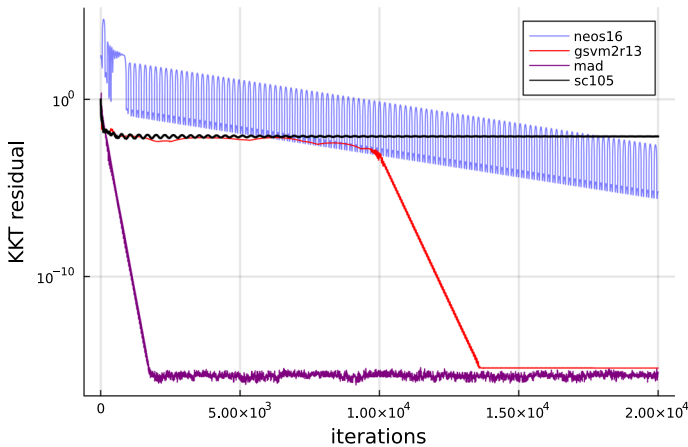


Fig. 1 PDHG on four LP relaxation of instances from MIPLIB

Despite the numerical success, there is a huge gap between the theoretical understanding and the practice of PDHG for LP. Recent works [6, 50] present the complexity theory of (restarted) PDHG for LP, which shows that the iterates of PDHG linearly converge to an optimal solution, but the linear convergence rate depends on the global Hoffman constant of the KKT system, which is known to be exponentially loose. On the other hand, it is evident that the numerical performance of the algorithm does not depend on the overly-conservative Hoffman constant; instead, the algorithm often exhibits a two-stage behavior: an initial sublinear convergence followed by a linear convergence. For instances, Fig. 1 shows the representative behaviors of PDHG on four instances from MIPLIB 2017. The convergence patterns of the algorithm differ dramatically across these LP instances. The instance `mad` converges linearly to optimality within only few thousands of steps. Instances `neos16` and `gsvm2r13` eventually reach the linear convergence stage but beforehand there exists a relatively flat slow stage. The instance `sc105` does not exhibit linear rate within twenty thousand iterations. Furthermore, the eventual linear rates are not equivalent: `gsvm2r13` exhibits much faster linear rate than `neos16`. Similar observation holds for the “warm-up” sublinear stage: `neos16` has much shorter sublinear period than `gsvm2r13`. The global linear convergence with the conservative Hoffman constant [6, 50] is clearly not enough to interpret the diverse empirical behaviors. The goal of this paper is to bridge such gap by answering the following two questions:

- How to understand the two-stage behaviors of PDHG for LP, and what are the geometric quantities that drive the length of the first stage and the convergence rate of the second stage?
- Can we obtain complexity theory of PDHG for LP without using the overly-conservative global Hoffman constant?

More formally, we consider standard-form LP:

$$\begin{aligned} \min_{x \geq 0} \quad & c^T x \\ \text{s.t.} \quad & Ax = b, \end{aligned} \tag{1}$$

and its primal–dual form:

$$\min_{x \geq 0} \max_y \quad c^T x - y^T Ax + b^T y. \tag{2}$$

Algorithm 1 Primal Dual Hybrid Gradient (PDHG) for (2)

Input: Initial point $z^0 = (x^0, y^0)$, step-size $s > 0$.

- 1: **for** $k = 0, 1, \dots$ **do**
 - 2: $x^{k+1} = \text{proj}_{\mathbb{R}_+^n} (x^k - s(c - A^T y^k))$
 - 3: $y^{k+1} = y^k - s(A(2x^{k+1} - x^k) - b)$
 - 4: **end for**
-

Algorithm 1 presents the primal–dual hybrid gradient method (PDHG, a.k.a Chambolle and Pock algorithm [12]) for LP (2). Without loss of generality, we assume the primal and the dual step-sizes are the same throughout the paper. This can be achieved by rescaling the problem instance [6]. The previous works [13, 26, 50] show the following complexity of PDHG for LP to achieve ϵ -accuracy solution:

$$\mathcal{O} \left(\frac{\|A\|_2^2}{1/H^2} \log \left(\frac{R}{\epsilon} \right) \right), \tag{3}$$

where R is essentially an upper bound on the norm of the iterates and H is the Hoffman constant of the KKT system of LP:

$$Ax = b, \quad x \geq 0, \quad A^T y \leq c, \quad \frac{1}{R}(c^T x - b^T y) \leq 0.$$

The major issue of the above complexity is the reliance on the global Hoffman constant H , which is notoriously conservative. Consider a general linear inequality system $Fx \leq g$. A commonly-used characterization to its Hoffman constant [60] is

$$H = \max_{\substack{J \in [m] \\ A_J \text{ full row rank}}} \frac{1}{\min_{\substack{v \in \mathbb{R}^{|J|} \\ \|v\|_2=1}} \|A_J^T v\|_2}, \tag{4}$$

The inner part of the Hoffman constant is an extension of the minimal non-zero singular value of the submatrix F_J , which is expected to appear in the linear convergence rate of first-order methods [55]. However, the outer maximization optimizes over exponentially many subsets of linear constraints, which makes the Hoffman constant overly conservative and cannot characterize the behaviors of the algorithms. Intuitively this is because when calculating the global Hoffman constant, one needs to look at the local geometry on every boundary set (i.e., extreme points, edges, faces, etc) of the feasible region, and consider the worst-case situation. See Appendix B for an example where the Hoffman constant can be arbitrarily loose.

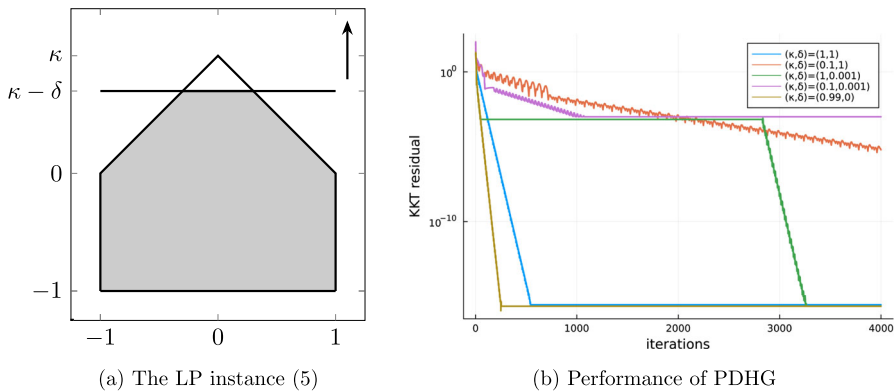


Fig. 2 Plots to illustrate the geometry of the LP instance (5), and the numerical behaviors of PDHG for solving the LP instance with different parameter κ, δ

In this paper, we show that the performance of PDHG for LP does not rely on the overly-conservative global Hoffman constant; instead, the algorithm exhibits a two-stage behavior:

- In the first stage (see Section 3 for more details), the algorithm aims to identify the non-degenerate variables. Notice that the algorithm may not converge to an optimal solution that satisfies strict complementary slackness, which we call degeneracy in this paper, and this definition is consistent with the partial smoothness literature. We further call the primal variables that satisfy strict complementary slackness the non-degenerate variables. This stage finishes within a finite number of iterations, and the convergence rate in this stage is sublinear. The driving force of the first stage is how the non-degenerate part of the LP is close to degeneracy.
- In the second stage (see Section 4 for more details), the algorithm effectively solves a homogeneous linear inequality system. The algorithm converges linearly to an optimal solution, and the driving force of the linear convergence rate is a local sharpness constant for the homogeneous linear inequality system. This local sharpness constant is much better than the global Hoffman constant, and it is a generalization of the minimal non-zero singular value of a certain matrix. Intuitively, this happens due to the structure of the homogeneous linear inequality system so that one just needs to focus on the local geometry around the origin (See [61] for a detailed characterization), avoiding going through the exponentially-many boundary set as in the calculation of global Hoffman constant.

To gain more intuitions of the two-stage convergence behavior, we consider a simple yet representative class of two-dimension house-shaped dual LPs parameterized by $0 \leq \delta \leq \kappa < 1$ (see Fig. 2a):

$$\begin{aligned}
 & \max y_2 \\
 & \text{s.t. } y_1 \geq -1, y_1 \leq 1 \\
 & \quad y_2 \geq -1, y_2 \leq \kappa - \delta \\
 & \quad y_1 + \frac{1}{\kappa}y_2 \leq 1, -y_1 + \frac{1}{\kappa}y_2 \leq 1.
 \end{aligned} \tag{5}$$

Parameter δ essentially measures how close the dual LP is to degeneracy. When $\delta = 0$ (i.e., the top three constraints in Fig. 2a intersect), the problem is degenerate; when δ is close to 0, the problem is near-degenerate; when δ is reasonably large, the problem is far away from degeneracy. Parameter κ roughly measures the condition number of the constraint matrix, i.e., the ratio between the minimal and the maximal singular values of the matrix. Figure 2b presents the numerical performance of PDHG on this two-dimension instance with different choices of δ and κ . Indeed, with a proper choice of the parameters δ and κ , the convergence behavior of Fig. 2b mimics well the real MIPLIB instances shown in Fig. 1. Furthermore, one can see that the smaller the value of κ is, the slower the eventual linear convergence is; the smaller the value of δ is, the longer the initial slow convergence stage. Yet, if $\delta = 0$, i.e., the problem is degenerate, the problem does not have the slow first stage. The above numerical observations are not limited to this two-dimension example, and they are formalized for general LP in Sects. 3 and 4.

Indeed, this two-stage behavior has been documented in the literature of interior-point methods and non-smooth optimization. For example, [28, 71, 72] discuss the two-stage behaviors of interior-point method (IPM) for LP: IPM has linear convergence in the first stage and super-linear convergence in the second stage. The phase transition happens when certain active basis are identified, and the length of the first phase depends on a certain metric that measures how close the analytical center of the optimal solution set (i.e., the converging optimal solution) is to degeneracy. In the context of non-smooth optimization, there are also fruitful results studying identification of first-order methods under partial smoothness [38, 43, 44, 70]. However, almost all of these works require the non-degeneracy condition, which, in the context of LP, refers to the algorithm converges to an optimal solution that satisfies strict complementary slackness. Unfortunately, this condition barely holds for real-world LP instances with PDHG as well as many other problems, and thus it is also called “irrepresentable condition” in the literature [24, 25]. Note that the identification results of IPMs also have an implicit reliance on the non-degeneracy condition, since the converging optimal solution of IPMs (i.e., analytical center of optimal solution face) always satisfies strict complementary slackness [28, 71, 72]. In contrast to these literatures, our work presents the first complexity result on finite time identification without the non-degeneracy assumption, in the context of LP. Indeed, our results suggest that degeneracy itself does not slow down the convergence for first-order methods, but near-degeneracy does.

1.1 Contributions

The goal of the paper is to provide a theoretical foundation for PDLP, a new first-order method LP solver based on PDHG. The contributions of the paper can be summarized as follows:

- Motivated by the empirical behaviors of PDHG for LP, we propose a two-stage characterization of its convergence behaviors:

- In the first stage, PDHG attempts to identify the active variables and the near-to-degeneracy parameter drives the identification complexity.
 - In the second stage, PDHG attempts to solve a homogeneous linear inequality system, and a local sharpness constant controls the linear convergence rate.
- We provide the first linear convergence complexity of PDHG for LP without the dependency on the global Hoffman constant.
 - Of independent interest, we present the first complexity bound on finite-time identification in the partial-smoothness context without the “irrepresentative” non-degeneracy condition.

1.2 Related Literature

FOM solvers for LP. Recently, first-order methods become increasingly appealing for solving large LP due to their low per-iteration cost [5–7, 41, 46, 69, 73].

- PDLP [4] utilizes primal–dual hybrid gradient method (PDHG) as its base algorithm and introduces practical algorithmic enhancements, such as presolving, preconditioning, adaptive restart, adaptive choice of step size, and primal weight, on top of PDHG. Right now, it has three implementations: a prototype implemented in Julia ([FirstOrderLp.jl](#)) for research purposes, a production-level C++ implementation that is included in Google [OR-Tools](#), and an internal distributed version at Google. The internal distributed version of PDLP has been used to solve real-world problems with as many as 92 billion non-zeros [54], which is one of the largest LP instances that are solved by a general-purpose LP solver.
- ABIP [21, 46] is an ADMM-based interior-point method. The core algorithm of ABIP is still a homogeneous self-dual embedded interior-point method. Instead of approximately minimizing the log-barrier penalty function with a Newton step, ABIP utilizes multiple steps of alternating direction method of multipliers (ADMM). The $\mathcal{O}\left(\frac{1}{\epsilon} \log\left(\frac{1}{\epsilon}\right)\right)$ sublinear complexity of ABIP was presented in [46]. Recently, [21] includes new enhancements, i.e., preconditioning, restart, hybrid parameter tuning, on top of ABIP (the enhanced version is called ABIP+). ABIP+ is numerically comparable to the Julia implementation of PDLP. ABIP+ now also supports a more general conic setting when the proximal problem associated with the log-barrier in ABIP can be efficiently computed.
- ECLIPSE [7] is a distributed LP solver designed specifically for addressing large-scale LPs encountered in web applications. These LPs have a certain decomposition structure, and the effective constraints are usually much less than the number of variables. ECLIPSE looks at a certain dual formulation of the problem, then utilizes accelerated gradient descent to solve the smoothed dual problem with Nesterov’s smoothing. This approach is shown to have $\mathcal{O}\left(\frac{1}{\epsilon}\right)$ complexity, and it is used to solve web applications with 10^{12} decision variables [7] and real-world web applications at LinkedIn platform [1, 64].
- SCS [57, 58] tackles the homogeneous self-dual embedding of general conic programming using ADMM. As a special case of conic programming, SCS can also be used to solve LP. Each iteration of SCS involves projecting onto the cone and

solving a system of linear equations with similar forms so that it only needs to store one factorization in memory. Furthermore, SCS supports solving linear equations with an iterative method, which only uses matrix-vector multiplications.

Complexity theory of FOMs for LP. One major drawback of applying first-order methods to solving LP is their slow tail convergence. Due to the lack of strong convexity, one can only obtain sublinear convergence rate when directly applying the classic results of FOMs on LP [8, 55]. This suggests that FOMs cannot produce the high-accuracy solutions that are typically expected from LP solvers within a reasonable time limit. Luckily, it turns out that LP satisfies additional growth condition structures that can improve the convergence of FOMs from sublinear to linear. For example, [22] showed that a variant of ADMM can have linear convergence for LP. It is shown in [43–45] that under non-degeneracy condition, many primal–dual algorithms have eventual linear convergence when solving LP, even though it may take a long time before reaching the linear local regime. Recently, [6] introduced a sharpness condition that is satisfied by LP, and proposed a restarted scheme for a variety of primal–dual methods to solve LP. It turns out that the restarted variants of many FOMs, such as PDHG, extra-gradient method (EGM) and ADMM, can achieve global linear convergence, and such linear convergence is the optimal rate, namely, there is a worst-case LP instance such that no FOMs can achieve better than such linear rate. Later on, [50] shows that PDHG (without restart) also has linear convergence for LP, and the linear convergence rate is slower than that in [6]. However, the obtained linear convergence rates in [6, 50] depend on the global Hoffman constant of the KKT system of LP, which is known to be very loose and cannot characterize the behaviors of the algorithm. In this paper, we aim to develop a tighter complexity theory of PDHG for LP without using global Hoffman constant.

Primal–dual hybrid gradient method (PDHG). PDHG was initially proposed for applications in image processing and computer vision [12, 15, 23, 32, 74]. The first convergence guarantee of PDHG was obtained in [12], which shows the $\mathcal{O}(1/k)$ sublinear ergodic convergence rate of PDHG for convex-concave primal–dual problems. Later on, simplified analyses for the $\mathcal{O}(1/k)$ sublinear ergodic rate of PDHG were presented in [13, 51]. More recently, it is shown that the last iterates of PDHG exhibit linear convergence under a mild regularity condition that is satisfied by many applications, including LP [26, 50]. Moreover, many variants of PDHG have been proposed, including adaptive version [27, 52, 62, 68] and stochastic version [2, 11]. It is also shown that PDHG is equivalent to Douglas-Rachford Splitting up to a linear transformation [48, 59].

Identification and partial smoothness. Active-set identification is a classic topic in constrained optimization and in non-smooth optimization, with both theoretical and computational importance. Partial smoothness, a notion meaning smoothness along some directions and sharpness in other directions, plays a crucial role in the analysis of identification. Under the partial smoothness assumption, a refined analysis is proposed for analyzing the identification property of certain algorithms. More specifically, manifold identification of dual averaging is shown in [35]. It is shown in [42, 43] that the finite time identification of forward-backward-type methods. In [44, 45], it is shown the active set identification and local convergence rate of many primal–dual

algorithms including PDHG and ADMM. In [63], it shows that stochastic methods also share a similar identification property. Similar results are derived for Newton’s method in [39]. The eventual sharp linear rate of Douglas–Rachford Splitting on basis pursuit is derived under non-degeneracy condition in [20]. In [19], the non-smooth dynamic of sub-gradient methods near active manifolds is studied.

However, all the above works assume the non-degeneracy condition of the problem, namely, $0 \in \text{ri}(\mathcal{F}(x^*))$, where $\text{ri}(\cdot)$ represents the relative interior of a set and $\mathcal{F}(x^*)$ denotes the sub-differential set at the converging optimal solution x^* . The only exceptions that do not require the non-degeneracy condition are [24, 25], where they showed that the active set of the iterates can potentially be larger than the active set of optimal solutions because of degeneracy. Unfortunately, it is unclear how to check the non-degeneracy condition in general, thus it is also called “irrepresentable condition” [24, 25]. In the case of LP, this condition means that the converging optimal solution satisfies strict complimentary slackness, which unfortunately is almost never the case when using FOMs to solve real-world LP instances. To the best of our knowledge, this work is the first complexity result of finite-time identification without the non-degeneracy condition.

On a related note, the identification behaviors of interior-point methods (IPMs) for LP have been observed [28, 71, 72], where it was shown that IPM has super-linear convergence after identification. Indeed, these results implicitly utilize the fact that IPM always converges to the analytical center of the optimal face, which satisfies strict complementary slackness [28, 71, 72].

1.3 Organization

We discuss the sharpness of homogeneous linear inequality system and collect necessary convergence results of PDHG in Section 2. The analysis of identification stage (Stage I) and local convergence (Stage II) are presented respectively in Section 3 and 4. Section 5 illustrates and verifies the theoretical results through numerical experiments. We conclude the paper and propose several future research directions in Section 6.

1.4 Notations

For matrix $A \in \mathbb{R}^{m \times n}$, denote A_j the j -th column of matrix A and $A_J = (A_j)_{j \in J} \in \mathbb{R}^{m \times |J|}$, $J \subset [n]$. Denote $\|A\|_2$ the operator norm of a matrix A . Denote $\|x\|_2$ the Euclidean norm for a vector x and $\langle x, y \rangle$ for its associated inner product. For a positive definite matrix $M \succ 0$, let $\langle x, y \rangle_M = \langle x, My \rangle$ and $\|x\|_M = \sqrt{x^T M x}$ the norm induced by the inner product. Let $\text{dist}_M(z, \mathcal{Z})$ represent the distance between point z and set \mathcal{Z} under the norm $\|\cdot\|_M$, that is, $\text{dist}_M(z, \mathcal{Z}) = \min_{u \in \mathcal{Z}} \|u - z\|_M$. In particular, $\text{dist}(z, \mathcal{Z}) = \min_{u \in \mathcal{Z}} \|u - z\|_2$. Denote $P_{\mathcal{Z}}(z) := \arg \min_{u \in \mathcal{Z}} \{\|u - z\|_2\}$ the projection onto set \mathcal{Z} under l_2 norm. Denote the optimal solution set to (2) as \mathcal{Z}^* . $f(x) = \mathcal{O}(g(x))$ means that for sufficiently large x , there exists constant C such that $f(x) \leq Cg(x)$. Let $\iota_{\mathcal{C}}(\cdot)$ be the indicator function of set \mathcal{C} . Let $z = (x, y)$, and denote $\mathcal{F}(z) = \mathcal{F}(x, y) = \left(\begin{array}{c} c - A^T y + \partial \iota_{\mathbb{R}_+^n}(x) \\ b - Ax \end{array} \right)$ the sub-differential of (2). Let

$P_S = \begin{pmatrix} \frac{1}{s}I & -A^T \\ -A & \frac{1}{s}I \end{pmatrix}$ and $\mathbf{1}_n$ represents all 1's vector of length n . For set A and B , let $A \setminus B := \{x \mid x \in A, x \notin B\}$. Let $[n] := \{1, 2, \dots, n\}$.

2 Preliminaries

We here present preliminary results on sharpness of primal–dual problems, convergence properties of PDHG and sharpness of homogeneous linear inequality system that we will use later.

2.1 Sharpness of primal–dual problems

Consider a generic convex-concave primal–dual problem

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} L(x, y), \tag{6}$$

where $L(x, y)$ is a convex function in x and a concave function in y . We use $z = (x, y)$ to represent the primal and the dual solution together.

Normalized duality gap was introduced in [6] as a progress metric for primal–dual problems:

Definition 1 For a primal–dual problem (6) and a solution $z = (x, y)$, the normalized duality gap with radius r is defined as

$$\rho_r(z) = \max_{\hat{z} \in W_r(z)} \frac{L(x, \hat{y}) - L(\hat{x}, y)}{r}, \tag{7}$$

where $W_r(z) = \{\hat{z} \in \mathcal{Z} \mid \|z - \hat{z}\|_2 \leq r\}$ is a ball centered at z with radius r intersected with $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$.

The normalized duality gap $\rho_r(z)$ is a valid progress metric for LP, since it provides an upper bound on the KKT residual:

Proposition 1 ([6, Lemma 4]) *It holds for any $z = (x, y)$ with $x \geq 0$ and $\|z\|_2 \leq R$ that for any $r \in (0, R]$,*

$$\rho_r(z) \geq \left\| \begin{pmatrix} Ax - b \\ [A^T y - c]^+ \\ \frac{1}{R}[c^T x - b^T y]^+ \end{pmatrix} \right\|_2.$$

The sharpness of a primal–dual problem is defined in [6]:

Definition 2 ([6, Definition 1]) We say a primal–dual problem is α -sharp on the set \mathcal{S} if $\rho_r(z)$ is α -sharp on \mathcal{S} for all r , i.e., it holds for all $z \in \mathcal{S}$ that

$$\alpha \text{dist}(z, \mathcal{Z}^*) \leq \rho_r(z). \tag{8}$$

The following proposition shows that the primal–dual form of LP is sharp.

Proposition 2 ([6, Lemma 5]) *LP in the primal–dual form (2) is a sharp problem on set $\{z \mid \|z\|_2 \leq R\}$, i.e., the normalized duality gap $\rho_r(z)$ defined in (7) satisfies*

$$\frac{1}{H} \text{dist}(z, \mathcal{Z}^*) \leq \rho_r(z),$$

for any $z \in \{z \mid \|z\|_2 \leq R\}$ and $r \in (0, R]$, where H is the Hoffman constant of the KKT system of LP:

$$Ax = b, A^T y \leq c, x \geq 0, \frac{1}{R}(c^T x - b^T y) \leq 0.$$

2.2 Convergence of PDHG iterates

In this subsection, we present some basic convergence properties of PDHG on LP in literature. Many of these convergence results hold for generic primal–dual problems, and we here state them in the context of LP for convenience.

The next lemma presents the non-expansiveness and convergence of PDHG on LP:

Lemma 1 ([13]) *Consider the iterations $\{z^k\}_{k=0}^\infty$ of PDHG on LP (Algorithm 1). Let \mathcal{Z}^* be the optimal solution set to (2). Then it holds for any k that*

- (a) (non-expansiveness) $\|z^{k+1} - z^*\|_{P_s} \leq \|z^k - z^*\|_{P_s}$ for any $z^* \in \mathcal{Z}^*$,
- (b) (convergence) there exists $z^* \in \mathcal{Z}^*$ such that $z^* = \lim_{k \rightarrow \infty} z^k$, and
- (c) (property of the optimal solution set) $\|z^k - \tilde{z}^*\|_{P_s} \geq \|z^* - \tilde{z}^*\|_{P_s}$ for any $\tilde{z}^* \in \mathcal{Z}^*$, where z^* is the converging optimal solution defined in (b).

The next theorem presents the sublinear and linear convergence rate of PDHG on LP. These results can be obtained utilizing a similar proof technique as in [50]. We present the formal proof in Appendix A for completeness.

Theorem 1 *Consider the iterations $\{z^k\}_{k=0}^\infty$ of PDHG on LP (Algorithm 1). Suppose the step-size $s \leq \frac{1}{2\|A\|_2}$. Then it holds for any $k \geq 0$ that*

- (a) (sublinear rate)

$$\|z^{k+1} - z^k\|_{P_s} \leq \text{dist}_{P_s^{-1}}(0, \mathcal{F}(z^k)) \leq \frac{\text{dist}_{P_s}(z^0, \mathcal{Z}^*)}{\sqrt{k}}. \tag{9}$$

- (b) (linear rate)

$$\|z^k - z^*\|_{P_s} \leq 4 \exp\left(-\frac{k}{2\lceil 4e/(s\alpha)^2 \rceil}\right) \text{dist}_{P_s}(z^0, \mathcal{Z}^*), \tag{10}$$

where α is the sharpness constant of LP along iterates, i.e., $\alpha \text{dist}(z^k, \mathcal{Z}^*) \leq \rho_r(z^k)$ for any $k \geq 0$.

We comment that we assume the step-size $s \leq \frac{1}{2\|A\|_2}$ in Theorem 1 and throughout the paper later on. While many of our results hold for $s < \frac{1}{\|A\|_2}$, assuming $s \leq \frac{1}{2\|A\|_2}$ makes it easy to convert between P_s norm and ℓ_2 norm, as stated in Lemma 2:

Lemma 2 Suppose $s \leq \frac{1}{2\|A\|_2}$. It holds for any $z \in \mathbb{R}^{m+n}$ that:

$$\sqrt{\frac{1}{2s}}\|z\|_2 \leq \|z\|_{P_s} \leq \sqrt{\frac{2}{s}}\|z\|_2 .$$

2.3 Sharpness of linear inequality system

The existing linear convergence results of PDHG for LP [6, 50] heavily depend on the Hoffman constant of the KKT system. However, it is well-known that the global Hoffman constant is generally very loose and cannot be easily characterized. Recently, [61] presents a simple and tight bound on the Hoffman constant (i.e., inverse of sharpness constant) for homogeneous linear inequality system. We here present these characterizations and discuss their fundamental differences.

First, recall the definition of the sharpness constant of a linear inequality system:

Definition 3 Consider a linear inequality system $Fx = g, \tilde{F}x \leq \tilde{g}$ and its solution set $\mathcal{X}^* = \{x \in \mathbb{R}^n : Fx = g, \tilde{F}x \leq \tilde{g}\}$. We call $\alpha > 0$ a sharpness constant to the system if it holds for any $x \in \mathbb{R}^n$ that

$$\alpha \text{dist}(x, \mathcal{X}^*) \leq \left\| \begin{pmatrix} Fx - g \\ (\tilde{F}x - \tilde{g})^+ \end{pmatrix} \right\|_2 .$$

The results of Hoffman [33, 60] show the existence of a sharpness constant α for any linear inequality system. The value of sharpness constant α is the inverse of the Hoffman constant to the linear inequality system.

Indeed, it is highly difficult to provide simple characterization to sharpness constant α . One characterization to α for system $Fx = g, \tilde{F}x \leq \tilde{g}$ is described in [60]:

$$\alpha = \min_{J \in \mathcal{S}(F, \tilde{F})} \min_{\substack{v \in \mathbb{R}_+^J, z \in F(\mathbb{R}^n) \\ \|(v, z)\|_2=1, \tilde{F}_J^T v + Fz - u = 0}} \|u\|_2 , \tag{11}$$

where $\mathcal{S}(F, \tilde{F}) = \{J \subset [m] : \{(\tilde{F}x + s, Fx) : s \in \mathbb{R}^m, s_J \geq 0, x \in \mathbb{R}^n\}$ is a linear subspace}.

Informally speaking, the inner minimization in (11) computes an extension of minimal positive singular value for a certain matrix, which is specified by an ‘‘active set’’ J of constraints. The outer minimization takes the minimum of these extended minimal positive singular values over all possible ‘‘active sets’’ (intuitively, it goes through every extreme point, edge, face, etc., of \mathcal{X}^*). While the inner minimization is expected when characterizing the behaviors of an algorithm, similar to the strong convexity in a minimization problem, there are usually exponentially many ‘‘active

sets” for a polytope \mathcal{X}^* , thus α defined in (11) is known to be a loose bound and it is generally NP-hard to compute its exact value or even just a reasonable bound.

On the other hand, it turns out a simple characterization exists for the sharpness of a homogeneous linear inequality system as shown recently in [61]. It does not require going through the exponential many “active sets” as in (11) and dramatically simplifies the characterization of the sharpness constant.

Proposition 3 (Proposition 1-3, Theorem 1 in [61]) *Consider a linear inequality system $Fx = 0, \tilde{F}x \leq 0$ where $F \in \mathbb{R}^{m \times n}$ and $\tilde{F} \in \mathbb{R}^{\tilde{m} \times n}$, and its solution set $\mathcal{X}^* = \{x \in \mathbb{R}^n : Fx = 0, \tilde{F}x \leq 0\}$. Denote α the sharpness constant to the system,*

i.e., it holds for any $x \in \mathbb{R}^n$ that $\alpha \text{dist}(x, \mathcal{X}^) \leq \left\| \begin{pmatrix} Fx - 0 \\ (\tilde{F}x - 0)^+ \end{pmatrix} \right\|_2$. Then:*

(a) *There exists a unique partition $P \cup Q = \{1, \dots, \tilde{m}\}$ such that for $\tilde{F} = \begin{pmatrix} \tilde{F}_P \\ \tilde{F}_Q \end{pmatrix}$,*

$$Fv = 0, \tilde{F}_P v = 0, \tilde{F}_Q v < 0 \text{ for some } v \in \mathbb{R}^n,$$

and

$$F^T w_1 + \tilde{F}_P^T w_2 = 0 \text{ for some } w_1, w_2 > 0.$$

(b) *Moreover, define*

$$L := \{v : Fv = 0, \tilde{F}_P v = 0\}, K := \{v : \tilde{F}_Q v \leq 0\}.$$

It holds that

$$\alpha(L, K) \cdot \min\{\alpha_0(L), \alpha_0(K)\} \leq \alpha \leq \min\{\alpha_0(L), \alpha_0(K)\}, \tag{12}$$

where

$$\alpha_0(K) \geq \min_{w \geq 0, \|w\|_2=1} \|\tilde{F}_Q^T w\|_2, \alpha_0(L) \geq \min_{\|v\|_2=1, Fv-w_1 \leq 0, \tilde{F}_P v-w_2 \leq 0} \|(w_1, w_2)\|_2, \tag{13}$$

and

$$\alpha(L, K) = \inf_{u \notin L \cap K} \frac{\max\{\text{dist}(u, L), \text{dist}(u, K)\}}{\text{dist}(u, L \cap K)} > 0. \tag{14}$$

Remark 1 Similar to the interpretation of the inner minimum in (11), $\alpha_0(K)$ and $\alpha_0(L)$ can be interpreted as an extension of the minimal non-zero singular value of the matrices. They are indeed highly related to and can be lower-bounded by the distance-to-infeasibility defined in Renegar’s condition number [65]. $\alpha(L, K)$ essentially measures the angle between K and L and it is strictly positive by the construction of L and K [61].

Compared to the characterization of sharpness constant of a non-homogeneous linear system (11), which requires computing the minimum over potentially exponentially many values, the bound in (12) for homogeneous linear system essentially just looks at the corresponding values of two sub-systems and combines them using the angle measurement $\alpha(L, K)$. As a result, the sharpness constant of a homogeneous system is much simpler and tighter than the non-homogeneous counterpart, and can be computed in polynomial time [61].

3 Stage I: finite time identification

In this section, we characterize the initial slow convergence stage and the phase transition between the two stages as described in Section 1. It turns out that PDHG can identify the “active basis” of the converging optimal solution in the first stage. We show that this stage always finishes (i.e., the phase transition happens) within a finite number of iterations, and the length of the first phase is driven by a quantity that characterizes how close the non-degeneracy part of the optimal solution is to degeneracy.

Such phase transition behavior is closely related to the concept of partial smoothness that was proposed in previous literature to understand non-smooth optimization and primal–dual algorithms [16, 30, 31, 36, 37, 40, 45, 56, 70]. In contrast to the previous literature, we do not assume the non-degeneracy condition of the problem. In the context of LP, the non-degeneracy condition in partial smoothness refers to that the iterates converge to an optimal solution which satisfies strict complimentary slackness. Such condition is also called “irrepresentable condition” in the optimization literature [24, 25], because it is highly rare for first-order methods to converge to a strictly complementary optimal solution in practice.

Indeed, the major difficulty of our analysis is to handle the degeneracy of the problem. To the best of our knowledge, this work is the first complexity result on finite-time identification without the non-degeneracy assumption.

To describe the identification property of the algorithm, we first introduce some new notations.

Definition 4 For a linear programming (2) and an optimal primal–dual solution $z^* = (x^*, y^*)$, we define the index partition (N, B_1, B_2) of the primal variable coordinates as

$$\begin{aligned} N &= \{i \in [n] : c_i - A_i^T y^* > 0\} \\ B_1 &= \{i \in [n] : c_i - A_i^T y^* = 0, x_i^* > 0\} \\ B_2 &= \{i \in [n] : c_i - A_i^T y^* = 0, x_i^* = 0\} . \end{aligned}$$

Since (x^*, y^*) is an optimal solution pair to the LP, it follows from complementary slackness that $x_N^* = 0$. However, (x^*, y^*) may not satisfy strict complementary slackness, i.e., B_2 can be a non-empty set. The lack of strict complementary slackness is called degeneracy herein.

In analogy to the notions in simplex method, we call the elements in N the non-basic variables, the elements in B_1 the non-degenerate basic variables, and the elements in

B_2 the degenerate basic variables. Furthermore, we call the elements in $N \cup B_1$ the non-degenerate variables and the elements in B_2 the degenerate variables. We also denote $B = B_1 \cup B_2$ as the set of basic variables.

Next, we introduce the non-degeneracy metric δ , which essentially measures how close the non-degenerate part of the optimal solution is to degeneracy:

Definition 5 For a linear programming (2) and an optimal primal–dual solution pair $z^* = (x^*, y^*)$, we define the non-degeneracy metric δ as:

$$\delta := \min \left\{ \min_{i \in N} \frac{c_i - A_i^T y^*}{\|A\|_2}, \min_{i \in B_1} x_i^* \right\}.$$

The non-degeneracy metric δ looks at the non-degenerate variables (i.e. $N \cup B_1$) and measures how close they are to degeneracy. For non-degenerate non-basic variables in N , it computes the minimal scaled reduced cost $\min_{i \in N} \frac{c_i - A_i^T y^*}{\|A\|_2}$; for non-degenerate basic variables in B_1 , it computes the minimal primal variable slackness $\min_{i \in B_1} x_i^*$; and then δ takes the minimum of the two. For every optimal solution z^* to a linear programming, the non-degeneracy metric δ is always strictly positive. The smaller the δ is, the closer the non-degenerate part of the solution is to degeneracy. In the rest of this section, we show that the value of δ drives the identification of the algorithm.

Next, we introduce the sharpness to a homogeneous conic system α_{L_1} that arises in the complexity of identification. We can view it as a local sharpness condition of a homogeneous linear inequality system around the converging optimal solution. Following the results stated in Section 2.3, α_L is an extension of the minimal non-zero singular value of a certain matrix and does not depend on the overly-conservative global Hoffman constant.

More formally, we consider the following homogeneous linear inequality system

$$A_B^T v \leq 0, Au = 0, u_{N \cup B_2} \geq 0, \frac{1}{R}(c^T u - b^T v) \leq 0, \tag{15}$$

and denote $\mathcal{K} := \{(u, v) \mid A_B^T v \leq 0, Au = 0, u_{N \cup B_2} \geq 0, \frac{1}{R}(c^T u - b^T v) \leq 0\}$ the feasible set of (15). Denote α_{L_1} the sharpness constant to (15), i.e., for any $(u, v) \in \mathbb{R}^{m+n}$,

$$\alpha_{L_1} \text{dist}((u, v), \mathcal{K}) \leq \left\| \begin{pmatrix} [A_B^T v]^+ \\ Au \\ [-u_{N \cup B_2}]^+ \\ \frac{1}{R}[c^T u - b^T v]^+ \end{pmatrix} \right\|_2. \tag{16}$$

Theorem 2 presents the main result for this section. It essentially states that after a certain number of iterations, PDHG iterates are reasonably close to the converging optimal solution and the algorithm can identify the elements in set N and B_1 . Furthermore, we present the complexity theory for identification in terms of the non-degenerate metric δ and the local sharpness constant α_{L_1} . Notice that δ does not rely

on the degenerate coordinates, i.e., those in B_2 , which shows that degeneracy does not slow down the identification, in contrast to previous literature on partial smoothness. We also comment that the set B_2 may not be identifiable by PDHG iterates due to the degeneracy of the problem.

Theorem 2 (Identification Complexity) *Consider PDHG for solving (2) with step-size $s \leq \frac{1}{2\|A\|_2}$. Let $\{z^k = (x^k, y^k)\}_{k=0}^\infty$ be the iterates of the algorithm and let z^* be the converging optimal solution, i.e., $z^k \rightarrow z^*$. Denote (N, B_1, B_2) the partition of the primal variable indices using Definition 4. Denote α_{L_1} the sharpness constant to the homogeneous linear inequality system (15). Denote $R = 2(\|z^0 - z^*\|_2 + \|z^*\|_2) + 1$. Then, it holds for any*

$$k \geq K := \max \left\{ 4, \frac{1}{s^2 \alpha_{L_1}^2} \right\} \frac{256R^2}{\delta^2} + \frac{2}{s\|A\|_2}, \tag{17}$$

that

$$\|z^k - z^*\|_2 \leq \frac{\delta}{2}, \text{ and } x_N^k = 0, x_{B_1}^k > 0, c_N - A_N^T y^k > 0.$$

Remark 2 Theorem 2 shows that after $K = \mathcal{O}\left(\max\left\{1, \frac{1}{s^2 \alpha_{L_1}^2}\right\} \frac{R^2}{\delta^2} + \frac{1}{s\|A\|_2}\right)$ iterations, the solution z^k is $\delta/2$ -close to the convergent optimal solution z^* . Furthermore, the set N and B_1 can be identified after K iterations.

The rest of this section presents the proof of Theorem 2. First, we consider the linear inequality system

$$A_B^T y \leq c_B, Ax = b, x_{N \cup B_2} \geq 0, \frac{1}{R}(c^T x - b^T y) \leq 0, \tag{18}$$

and denote $\mathcal{Z}_L^* := \{(x, y) \mid A_B^T y \leq c_B, Ax = b, x_{N \cup B_2} \geq 0, \frac{1}{R}(c^T x - b^T y) \leq 0\}$ the feasible set to (18).

The next lemma builds up the connections between \mathcal{Z}_L^* and \mathcal{K} . More specifically, it states that (a) \mathcal{Z}_L^* is a shift of the cone \mathcal{K} , (b) the sharpness constant of the two linear inequality systems (15) and (18) are the same, (c) one can lower bound $\text{dist}(0, \mathcal{F}(z))$ using $\text{dist}(z, \mathcal{Z}_L^*)$ and the sharpness constant of the systems.

Lemma 3 (a). *It holds that*

$$\mathcal{Z}_L^* = z^* + \mathcal{K}.$$

(b). α_{L_1} is the sharpness constant to system (18), that is, for any $z = (x, y) \in \mathbb{R}^{m+n}$,

$$\alpha_{L_1} \text{dist}(z, \mathcal{Z}_L^*) \leq \left\| \begin{pmatrix} [A_B^T y - c_B]^+ \\ Ax - b \\ [-x_{N \cup B_2}]^+ \\ \frac{1}{R}[c^T x - b^T y]^+ \end{pmatrix} \right\|_2.$$

(c). For any $z \in B_R(0) \cap \{x : x \geq 0\}$ we have

$$\alpha_{L_1} \text{dist}(z, \mathcal{Z}_L^*) \leq \text{dist}(0, \mathcal{F}(z)) .$$

Proof (a). For any $z \in z^* + \mathcal{K}$, there exists $w = (u, v) \in \mathcal{K}$ such that $z = z^* + w$. By the definition of set N, B, B_1, B_2 , we have

$$\begin{aligned} A_B^T(y^* + v) - c_B &= A_B^T v \leq 0 \\ A(x^* + u) &= b + Au = b \\ (x^* + u)_{N \cup B_2} &= u_{N \cup B_2} \geq 0 \\ \frac{1}{R} \left(c^T(x^* + u) - b^T(y^* + v) \right) &= \frac{1}{R}(c^T u - b^T v) \leq 0 , \end{aligned}$$

which implies

$$z^* + \mathcal{K} \subset \mathcal{Z}_L^* . \tag{19}$$

On the other hand, for any $z \in \mathcal{Z}_L^*$, we have $z = z^* + (z - z^*)$. Notice that

$$\begin{aligned} A_B^T(y - y^*) &= A_B^T y - c_B \leq 0 \\ A(x - x^*) &= Ax - b = 0 \\ (x - x^*)_{N \cup B_2} &= x_{N \cup B_2} \geq 0 \\ \frac{1}{R} \left(c^T(x - x^*) - b^T(y - y^*) \right) &= \frac{1}{R}(c^T x - b^T y) \leq 0 , \end{aligned}$$

thus $z - z^* \in \mathcal{K}$, which implies

$$\mathcal{Z}_L^* \subset z^* + \mathcal{K} . \tag{20}$$

Combining (19) and (20), we arrive at

$$\mathcal{Z}_L^* = z^* + \mathcal{K} .$$

(b). By definition of α_{L_1} , we have

$$\begin{aligned} \alpha_{L_1} \text{dist}(z, \mathcal{Z}_L^*) &= \alpha_{L_1} \text{dist}(z - z^*, \mathcal{Z}_L^* - z^*) \\ &= \alpha_{L_1} \text{dist}(z - z^*, \mathcal{K}) \\ &\leq \left\| [A_B^T(y - y^*)]^+, A(x - x^*), [-(x - x^*)_{N \cup B_2}]^+ , \right. \\ &\quad \left. \frac{1}{R} [c^T(x - x^*) - b^T(y - y^*)]^+ \right\|_2 \\ &= \left\| ([A_B^T y - c_B]^+, Ax - b, [-x_{N \cup B_2}]^+, \frac{1}{R} [c^T x - b^T y]^+) \right\|_2 . \end{aligned}$$

(c). From the definition of sharpness and Lemma 3, it holds for any $z \in B_R(0)$ with $x \geq 0$ that

$$\begin{aligned} \alpha_{L_1} \text{dist}(z, \mathcal{Z}_L^*) &\leq \left\| \left([A_B^T y - c_B]^+, Ax - b, \frac{1}{R} [c^T x - b^T y]^+ \right) \right\|_2 \\ &\leq \left\| \left([A^T y - c]^+, Ax - b, \frac{1}{R} [c^T x - b^T y]^+ \right) \right\|_2 \\ &\leq \rho_r(z) \leq \text{dist}(0, \mathcal{F}(z)) , \end{aligned}$$

where the third and fourth inequalities follow from Proposition 1 and [49, Proposition 1] respectively. □

The next lemma shows when z is close enough to z^* , the non-basic dual slacks in set N and the non-degenerate basic variables in set B_1 are bounded away from zero.

Lemma 4 For any z such that $\|z - z^*\|_2 \leq \frac{\delta}{2}$, it holds for any $i \in N, j \in B_1$ that

$$c_i - A_i^T y \geq \frac{\delta}{2} \|A\|_2, \quad x_j \geq \frac{\delta}{2} .$$

Proof It follows from $\|z - z^*\|_2 \leq \frac{\delta}{2}$ that

$$\|x - x^*\|_2 \leq \frac{\delta}{2}, \quad \|y - y^*\|_2 \leq \frac{\delta}{2} .$$

As a result, it holds for any $j \in B_1$ and $i \in N$ that

$$\begin{aligned} |x_j - x_j^*| &\leq \|x - x^*\|_2 \leq \frac{\delta}{2} , \\ |c_i - A_i^T y - c_i - A_i^T y^*| &\leq \|(c - A^T y) - (c - A^T y^*)\|_2 \\ &\leq \|A\|_2 \|y - y^*\|_2 \leq \|A\|_2 \frac{\delta}{2} , \end{aligned}$$

and thus by definition of δ , we have

$$\begin{aligned} x_j &\geq -\frac{\delta}{2} + x_j^* \geq \frac{\delta}{2} , \\ c_i - A_i^T y &\geq -\|A\|_2 \frac{\delta}{2} + (c_i - A_i^T y^*) \geq \|A\|_2 \left(-\frac{\delta}{2} + \delta \right) = \|A\|_2 \frac{\delta}{2} . \end{aligned}$$

□

The next lemma connects \mathcal{Z}_L^* and the optimal solution set \mathcal{Z}^* .

Lemma 5 $\mathcal{Z}_L^* \cap B_\delta(z^*) \subset \mathcal{Z}^*$.

Proof Following the same proof of Lemma 4, we know that for any $z \in B_\delta(z^*)$,

$$c_N - A_N^T y \geq 0, \quad x_{B_1} \geq 0 . \tag{21}$$

Recall the definition of \mathcal{Z}_L^* as the solution set to the system

$$A_B^T y - c_B \leq 0, Ax = b, x_{N \cup B_2} \geq 0, \frac{1}{R}(c^T x - b^T y) \leq 0. \tag{22}$$

Note that combining (21) and (22) gives exactly the KKT system of original LP (2) and thus

$$\mathcal{Z}_L^* \cap B_\delta(z^*) \subset \mathcal{Z}^* .$$

□

The following lemma ensures that after a certain number of iterations, the iterates $\{z^k\}$ stay in the ball centered at z^* with radius $\frac{\delta}{2}$.

Lemma 6 *Under the same condition of Theorem 2, for any iteration k such that*

$$k \geq K_1 := \max \left\{ 4, \frac{1}{s^2 \alpha_{L_1}^2} \right\} \frac{256 \text{dist}^2(z^0, \mathcal{Z}^*)}{\delta^2} ,$$

we have

$$\|z^k - z^*\|_2 \leq \frac{\delta}{2} .$$

Proof We first show that for any $k \geq K_1$, it holds that $\|z^{k+1} - z^k\|_2 \leq \frac{\delta}{16}$ and $\text{dist}(z^k, \mathcal{Z}_L^*) \leq \frac{\delta}{8}$, then we show $\|z^k - z^*\|_2 \leq \frac{\delta}{2}$. The proof heavily relies on the non-expansive nature of PDHG iterates. Notice that

$$\|z^{k+1} - z^k\|_2 \leq \sqrt{2s} \|z^{k+1} - z^k\|_{P_s} \leq \frac{\sqrt{2s} \text{dist}_{P_s}(z^0, \mathcal{Z}^*)}{\sqrt{k}} \leq \frac{2 \text{dist}(z^0, \mathcal{Z}^*)}{\sqrt{k}} ,$$

where the first and third inequalities use Lemma 2 while the second one follows from part (a) in Theorem 1. Furthermore, it holds that

$$\begin{aligned} \text{dist}(z^k, \mathcal{Z}_L^*) &\leq \frac{1}{\alpha_{L_1}} \text{dist}(0, \mathcal{F}(z^k)) \leq \frac{1}{\alpha_{L_1}} \sqrt{\frac{2}{s}} \text{dist}_{P_s^{-1}}(0, \mathcal{F}(z^k)) \leq \sqrt{\frac{2}{s}} \frac{\text{dist}_{P_s}(z^0, \mathcal{Z}^*)}{\alpha_{L_1} \sqrt{k}} \\ &\leq \frac{2 \text{dist}(z^0, \mathcal{Z}^*)}{s \alpha_{L_1} \sqrt{k}} , \end{aligned}$$

where the first inequality follows from part (c) in Lemma 3, the second and fourth inequalities utilize Lemma 2, and the third inequality is from part (b) of Theorem 1. Therefore, we have $\|z^{k+1} - z^k\|_2 \leq \frac{\delta}{16}$ and $\text{dist}(z^k, \mathcal{Z}_L^*) \leq \frac{\delta}{8}$ whenever iteration number satisfies $k \geq K_1$.

In the rest of the proof, we consider two cases, separating by whether $\|P_{\mathcal{Z}_L^*}(z^k) - z^*\|_2 \leq \frac{\delta}{2}$ or not. We will show that the first case implies $\|z^k - z^*\|_2 \leq \frac{\delta}{2}$, and it is impossible for the second case to happen, which then finishes the proof of the lemma.

Case I: $\|P_{\mathcal{Z}_L^*}(z^k) - z^*\|_2 \leq \frac{\delta}{2}$. In this case, we have $P_{\mathcal{Z}_L^*}(z^k) \in \mathcal{Z}_L^* \cap B_{\delta/2}(z^*) \subset \mathcal{Z}^*$ from Lemma 5 and therefore it follows from part (c) in Lemma 1 that

$$\|z^k - P_{\mathcal{Z}_L^*}(z^k)\|_{P_s} \geq \|z^* - P_{\mathcal{Z}_L^*}(z^k)\|_{P_s} .$$

Thus, it holds that

$$\begin{aligned} \|z^k - z^*\|_{P_s} &\leq \|z^k - P_{\mathcal{Z}_L^*}(z^k)\|_{P_s} + \|P_{\mathcal{Z}_L^*}(z^k) - z^*\|_{P_s} \leq 2\|z^k - P_{\mathcal{Z}_L^*}(z^k)\|_{P_s} \\ &\leq \frac{2\sqrt{2}}{\sqrt{s}} \|z^k - P_{\mathcal{Z}_L^*}(z^k)\|_2 . \end{aligned}$$

whereby

$$\|z^k - z^*\|_2 \leq \sqrt{2s} \|z^k - z^*\|_{P_s} \leq 4\|z^k - P_{\mathcal{Z}_L^*}(z^k)\|_2 = 4\text{dist}(z^k, \mathcal{Z}_L^*) \leq \frac{\delta}{2} .$$

Case II: $\|P_{\mathcal{Z}_L^*}(z^k) - z^*\|_2 > \frac{\delta}{2}$. We aim to show this case is impossible under conditions $\|z^{k+1} - z^k\|_2 \leq \frac{\delta}{16}$ and $\text{dist}(z^k, \mathcal{Z}_L^*) \leq \frac{\delta}{8}$.

First, since the iterates converge to an optimal solution (i.e. $z^k \rightarrow z^*$), we know that $\|z^k - z^*\|_2 \geq \|P_{\mathcal{Z}_L^*}(z^k) - z^*\|_2 > \frac{\delta}{2}$ and $\|z^k - z^*\|_2$ decreases to zero. Then there exists the smallest $N \geq k$ such that

$$\|z^N - z^*\|_2 > \frac{\delta}{2}, \quad \|z^{N+1} - z^*\|_2 \leq \frac{\delta}{2} ,$$

and hence

$$\|P_{\mathcal{Z}_L^*}(z^N) - z^*\|_2 > \frac{\delta}{2}, \quad \|P_{\mathcal{Z}_L^*}(z^{N+1}) - z^*\|_2 \leq \frac{\delta}{2} , \tag{23}$$

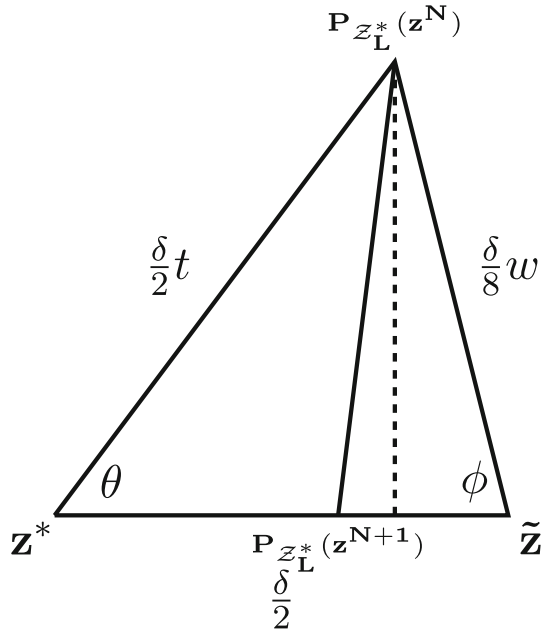
where the first one is due to the conclusion in Case I and the second one is the consequence of non-expansiveness of projection. It then follows from Lemma 5 that $P_{\mathcal{Z}_L^*}(z^{N+1}) \in \mathcal{Z}_L^* \cap B_{\delta/2}(z^*) \subset \mathcal{Z}^*$.

If $P_{\mathcal{Z}_L^*}(z^{N+1}) = z^*$, then from (23) we have

$$\|P_{\mathcal{Z}_L^*}(z^N) - P_{\mathcal{Z}_L^*}(z^{N+1})\|_2 = \|P_{\mathcal{Z}_L^*}(z^N) - z^*\|_2 > \frac{\delta}{2} \tag{24}$$

Otherwise, denote $\tilde{z} = z^* + \frac{\delta}{2\|P_{\mathcal{Z}_L^*}(z^{N+1}) - z^*\|_2} (P_{\mathcal{Z}_L^*}(z^{N+1}) - z^*) \in \mathcal{Z}_L^*$ and then $\|\tilde{z} - z^*\|_2 = \frac{\delta}{2}$, which implies $\tilde{z} \in \mathcal{Z}_L^* \cap B_{\delta/2}(z^*) \subset \mathcal{Z}^*$. By triangle inequality, we

Fig. 3 Illustration of the proof of Case II in Lemma 6



have

$$\|P_{Z_L^*}(z^N) - \tilde{z}\|_2 \geq \|z^N - \tilde{z}\|_2 - \|P_{Z_L^*}(z^N) - z^N\|_2 \geq \frac{\delta}{4} - \text{dist}(z^N, Z_L^*) \geq \frac{\delta}{8}, \tag{25}$$

where the second inequality follows from

$$\|z^N - \tilde{z}\|_2 \geq \sqrt{\frac{s}{2}} \|z^N - \tilde{z}\|_{P_s} \geq \sqrt{\frac{s}{2}} \|z^* - \tilde{z}\|_{P_s} \geq \sqrt{\frac{s}{2}} \sqrt{\frac{1}{2s}} \|z^* - \tilde{z}\|_2 = \frac{1}{2} \times \frac{\delta}{2} = \frac{\delta}{4},$$

and $\|P_{Z_L^*}(z^N) - z^N\|_2 = \text{dist}(z^N, Z_L^*) \leq \frac{\delta}{8}$ as $N \geq k \geq K_1$.

Consider the two-dimensional plane that z^* , \tilde{z} and $P_{Z_L^*}(z^N)$ lay (see Fig. 3). We know by construction of \tilde{z} that $P_{Z_L^*}(z^{N+1})$ lies on the segment between z^* and \tilde{z} . Let $\|P_{Z_L^*}(z^N) - z^*\|_2 = t \frac{\delta}{2}$ and $\|P_{Z_L^*}(z^N) - \tilde{z}\|_2 = w \frac{\delta}{8}$, then we have $t \geq 1$ and $w \geq 1$ by (23) and (25). Define θ as the angle between vector $P_{Z_L^*}(z^N) - z^*$ and $\tilde{z} - z^*$, and ϕ as the angle between vector $P_{Z_L^*}(z^N) - \tilde{z}$ and $z^* - \tilde{z}$. There are three cases:

(a) If $\cos(\theta) \leq 0$, then we have $\|P_{Z_L^*}(z^N) - P_{Z_L^*}(z^{N+1})\|_2 \geq \|P_{Z_L^*}(z^N) - z^*\|_2 \geq \frac{\delta}{2} t \geq \frac{\delta}{2}$.

(b) If $\cos(\phi) \leq 0$, then we have $\|P_{Z_L^*}(z^N) - P_{Z_L^*}(z^{N+1})\|_2 \geq \|P_{Z_L^*}(z^N) - \tilde{z}\|_2 \geq \frac{\delta}{8} w \geq \frac{\delta}{8}$.

(c) It $\cos(\theta) > 0$ and $\cos(\phi) > 0$, then it follows from law of cosine that

$$\frac{1}{16}w^2 = 1 + t^2 - 2t\cos(\theta) \leq 1 + t^2 ,$$

thus

$$t^2 \geq \frac{1}{16}w^2 - 1 . \tag{26}$$

Again by law of cosine and note that $t \geq 1$ and (26), we have

$$1 + \frac{1}{16}w^2 - \frac{1}{2}w\cos(\phi) = t^2 \geq \max \left\{ 1, \frac{1}{16}w^2 - 1 \right\} ,$$

which implies

$$\cos(\phi) \leq \min \left\{ \frac{w}{8}, \frac{4}{w} \right\} \leq \frac{\sqrt{2}}{2} ,$$

whereby

$$\sin(\phi) \geq \frac{\sqrt{2}}{2} .$$

By two-dimensional geometry, we have

$$\|P_{\mathcal{Z}_L^*}(z^N) - P_{\mathcal{Z}_L^*}(z^{N+1})\|_2 \geq \frac{\delta}{8}w\sin(\phi) \geq \frac{\sqrt{2}}{16}\delta .$$

Combining (a)-(c) and (24), we have

$$\|z^N - z^{N+1}\|_2 \geq \|P_{\mathcal{Z}_L^*}(z^N) - P_{\mathcal{Z}_L^*}(z^{N+1})\|_2 > \min \left\{ \frac{\delta}{2}, \frac{\delta}{8}, \frac{\sqrt{2}\delta}{16} \right\} > \frac{\delta}{16} ,$$

However, since $N \geq k \geq K_1$, we know $\|z^{N+1} - z^N\|_2 \leq \frac{\delta}{16}$. This leads to the contradiction.

To sum up, for any iteration k , when $\|z^{k+1} - z^k\|_2 \leq \frac{\delta}{16}$ and $\text{dist}(z^k, \mathcal{Z}_L^*) \leq \frac{\delta}{8}$, it must have $\|P_{\mathcal{Z}_L^*}(z^k) - z^*\|_2 \leq \frac{\delta}{2}$ (i.e., Case II cannot happen) and thus $\|z^k - z^*\|_2 \leq \frac{\delta}{2}$ following the argument in Case I, which completes the proof of Lemma 6. \square

By combining Lemmas 4 and 6, we are ready to prove Theorem 2.

Proof of Theorem 2 Consider $k \geq K \geq K_1$. It follows from Lemma 6 that $\|z^k - z^*\|_2 \leq \frac{\delta}{2}$. Thus, we have from Lemma 4 that $x_{B_1}^k > 0, c_N - A_N^T y^k > 0$. We just need to show

that $x_N^k = 0$. Notice it holds from Lemma 4 that for any $i \in N$ and $k \geq K_1$

$$c_i - A_i^T y^k \geq \|A\|_2 \frac{\delta}{2}.$$

Consider the update formula of PDHG

$$x_N^{k+1} = \text{proj}_{\mathbb{R}_+^{|N|}}(x_N^k - s(c_N - A_N^T y^k)).$$

Thus, it follows from Lemma 2 that

$$\max_{i \in N} x_i^{K_1} \leq \|z_N^{K_1}\|_2 \leq \sqrt{2s} \|z_N^{K_1}\|_{P_s} = \sqrt{2s} \|z_N^{K_1} - z_N^*\|_{P_s} \leq \sqrt{2s} \|z^{K_1} - z^*\|_{P_s} \leq \delta.$$

Then for any $q > \frac{2\delta}{s\|A\|_2} = \frac{2}{s\|A\|_2}$ and $i \in N$, we have $\sum_{k=0}^{q-1} s(c - A^T y^{K_1+k})_i \geq \frac{q\delta s\|A\|_2}{2}$, thus it holds that

$$x_N^{K_1+q} = 0,$$

that is, after at most $K := K_1 + \frac{2}{s\|A\|_2}$ steps, the coordinates in set N are fixed to zero. This finishes the proof of Theorem 2. □

4 Stage II: local convergence

In the previous section, we show that after a certain number of iterations, the iterates of PDHG stay close to an optimal solution, and the iterates identify the non-degenerate coordinates set N and B_1 . In this section, we characterize the local behaviors of PDHG for LP afterward. More specifically, we show that PDHG iterates converge to an optimal solution linearly after identification, and the local linear convergence rate is characterized by a local sharpness constant α_{L_2} for a homogeneous linear inequality system. The local convergence can be much faster than the global linear convergence rate obtained in [6], which is characterized by the Hoffman constant of the KKT system of the LP. This provides a theoretical explanation of the eventual fast rate of PDHG for LP.

Throughout this section, we consider the iterates z^k after identification, i.e., for $k \geq K$. In this stage, it holds that $\|z^k - z^*\|_2 \leq \frac{\delta}{2}$ and $x_N^k = 0$, $x_{B_1}^k > 0$, $c_N - A_N^T y^k > 0$ from Theorem 2. We call this stage the local setting. To develop the eventual faster linear convergence rate in the local setting, we first reformulate the KKT system of LP to remove the non-basic variables. Then we show the faster local linear convergence rate using the sharpness of a homogeneous linear inequality system.

Recall that the KKT system for original linear programming (2) is

$$\begin{cases} Ax = b, x \geq 0 \\ A^T y \leq c \\ c^T x - b^T y \leq 0 \end{cases}. \tag{27}$$

Every solution that satisfies (27) is an optimal primal–dual solution pair to the original LP (2). In the local setting, note that $x_N^k = 0$, $c_N - A_N^T y^k > 0$, (27) can be reduced to the following system by removing non-basic coordinates N :

$$\begin{cases} A_B x_B = b, x_B \geq 0 \\ A_B^T y \leq c_B \\ c_B^T x_B - b^T y \leq 0 \end{cases} . \tag{28}$$

In particular, for any local iterate z^k for $k \geq K$, if (28) is satisfied, then it is an optimal solution to (27).

Denote $\widetilde{\mathcal{X}}_B^* \times \widetilde{\mathcal{Y}}^*$ the solution set to (28). Recall that z^* is the optimal solution that the algorithm converges to. By the definition of set B and N , we have $z^* = (x^*, y^*) = (x_N^*, x_B^*, y^*)$, where $x_N^* = 0$. Since z^* satisfy (27), it must hold that $(x_B^*, y^*) \in \widetilde{\mathcal{X}}_B^* \times \widetilde{\mathcal{Y}}^*$. Moreover, from the definition of set N and B_1, B_2 , it holds that

$$\begin{cases} A_B x_B^* = b, x_{B_1}^* > 0, x_{B_2}^* = 0 \\ A_B^T y^* = c_B \\ c_B^T x_B^* - b^T y^* = 0 \end{cases} . \tag{29}$$

Next, we discuss the characterization of local sharpness properties for (28) and its critical role in the local convergence rate. The idea is to connect (28) in the local setting with a homogeneous conic system. The results explain why the local rate is much faster than the global rate that is dependent on the global Hoffman constant.

In this section, we aim to derive the local convergence complexity of PDHG on LP. To develop the eventual faster linear convergence rate, we first analyze the sharpness of the reduced primal–dual system of LP and build the connection with a homogeneous linear inequality system. Based on the understanding of local sharpness, a refined local convergence rate is then derived.

We consider the following homogeneous linear inequality system

$$A_B u_B = 0, u_{B_2} \geq 0, A_B^T v \leq 0, \frac{1}{R}(c_B^T u_B - b^T v) \leq 0. \tag{30}$$

and denote $\mathcal{W}^* := \{(u_B, v) \mid A_B u_B = 0, A_B^T v \leq 0, \frac{1}{R}(c_B^T u_B - b^T v) \leq 0, u_{B_2} \geq 0\}$ the feasible set of (30). Denote $R_2 = \delta + \|z^*\|$. Denote α_{L_2} the sharpness constant to (30), i.e., for any $(u_B, v) \in \mathbb{R}^{|B|+m}$,

$$\alpha_{L_2} \text{dist}((u_B, v), \mathcal{W}^*) \leq \left\| \begin{pmatrix} A_B u_B \\ [A_B^T v]^+ \\ [-u_{B_2}]^+ \\ \frac{1}{R_2} [c_B^T u_B - b^T v]^+ \end{pmatrix} \right\|_2.$$

There is a slight difference of the sharpness constants α_{L_1} and α_{L_2} , because they correspond to slightly different linear inequality systems, (15) and (30), respectively. The difference is that in the Stage II (i.e., α_{L_2}), the non-basic set N is identified, i.e., $x_N^k = 0$. Thus in Stage II, the iterates of PDHG are driven by sharpness α_{L_2} of system (30) which can ignore the constraints in non-basic set N , while sharpness α_{L_1} of (15) does depend on the whole constraint matrix. Nevertheless, α_{L_1} and α_{L_2} are sharpness of homogeneous linear inequality systems (15) and (30) respectively and thus exhibit simpler and tighter characterizations, as stated in Section 2.3. We present a simple example in Appendix B that demonstrates α_{L_2} can be arbitrarily larger than the global sharpness constant α .

Theorem 3 presents the local linear convergence of PDHG after identification.

Theorem 3 (Local Linear Convergence) *Consider PDHG for solving (2) with step-size $s \leq \frac{1}{2\|A\|_2}$. Let $\{z^k = (x^k, y^k)\}_{k=0}^\infty$ be the iterates of the algorithm and let z^* be the converging optimal solution, i.e., $z^k \rightarrow z^*$. Then it holds for any $k > K$ that*

$$\|z^k - z^*\|_2 \leq 4\delta \exp\left(-\frac{k - K}{2\lceil 4e/(s^2\alpha_{L_2}^2) \rceil}\right),$$

where K is the maximal length of the first stage defined in (17).

Theorem 3 shows that after identification, PDHG can obtain an ϵ -close optimal solution within

$$\mathcal{O}\left(\frac{1}{s^2\alpha_{L_2}^2} \log\left(\frac{\delta}{\epsilon}\right)\right)$$

iterations. Combining Theorems 2 and 3, the complexity of PDHG finding an ϵ -close optimal solution is

$$\mathcal{O}\left(\max\left\{1, \frac{1}{s^2\alpha_{L_1}^2}\right\} \frac{R^2}{\delta^2} + \frac{1}{s\|A\|_2} + \frac{1}{s^2\alpha_{L_2}^2} \log\left(\frac{\delta}{\epsilon}\right)\right).$$

In particular, when setting step-size $s = \frac{C}{\|A\|_2}$ with $C \leq \frac{1}{2}$, the total complexity becomes

$$\mathcal{O}\left(\max\left\{1, \frac{\|A\|_2^2}{\alpha_{L_1}^2}\right\} \frac{R^2}{\delta^2} + \frac{\|A\|_2^2}{\alpha_{L_2}^2} \log\left(\frac{\delta}{\epsilon}\right)\right).$$

The rest of this section presents a proof of Theorem 3. First, we show that α_{L_2} is a local sharpness constant of the non-homogeneous linear inequality system (29) for (x_B, y) that is close enough to (x_B^*, y^*) .

Lemma 7 *It holds for any (x_B, y) such that $\|(x_B, y) - (x_B^*, y^*)\|_2 \leq \frac{\delta}{2}$ and $x_B \geq 0$ that*

$$\alpha_{L_2} \text{dist}((x_B, y), \widetilde{\mathcal{X}}_B^* \times \widetilde{\mathcal{Y}}^*) \leq \left\| \begin{pmatrix} A_B x_B - b \\ [A_B^T y - c_B]^+ \\ [\frac{1}{R_2} (c_B^T x_B - b^T y)]^+ \end{pmatrix} \right\|_2.$$

Proof Denote $\widetilde{\mathcal{Z}} = \{(x_B, y) : \|(x_B, y) - (x_B^*, y^*)\|_2 \leq \frac{\delta}{2}, x_B \geq 0\}$. Notice by the definition of δ , we have $x_i^* \geq \delta$ for any $i \in B_1$. Thus, it holds for any $(x, y) \in \widetilde{\mathcal{Z}}$ and $i \in B_1$ that $x_i \geq x_i^* - \|x_B - x_B^*\|_2 \geq \frac{\delta}{2} > 0$. In other words, the constraint $\|(x_B, y) - (x_B^*, y^*)\|_2 \leq \frac{\delta}{2}$ implies $x_{B_1} \geq 0$. Therefore, we have

$$\widetilde{\mathcal{Z}} = \{(x_B, y) : \|(x_B, y) - (x_B^*, y^*)\|_2 \leq \frac{\delta}{2}, x_{B_2} \geq 0\}. \tag{31}$$

Now denote $u_B = (u_{B_1}, u_{B_2}) := x_B - x_B^*$ as the shift of x_B from x_B^* , namely, $u_{B_1} = x_{B_1} - x_{B_1}^*$ and $u_{B_2} = x_{B_2} - x_{B_2}^* = x_{B_2}$ (by definition of B_2 we have $x_{B_2}^* = 0$). Denote $v := y - y^*$ as the shift of y from y^* . It then holds that $(u_B, v) \in \widetilde{\mathcal{W}}^* := \{(u_B, v) \mid \|(u_B, v)\|_2 \leq \frac{\delta}{2}, u_{B_2} \geq 0\}$.

Next, notice $0 \in \mathcal{W}^*$, thus $P_{\mathcal{W}^*}(0) = 0$. It then follows from the nonexpansiveness of the projection operator that

$$\begin{aligned} \|P_{\mathcal{W}^*}(u_B, v)\|_2 &= \|P_{\mathcal{W}^*}(u_B, v) - P_{\mathcal{W}^*}(0)\|_2 \leq \|(u_B, v) - (0, 0)\|_2 \\ &= \|(u_B, v)\|_2 \leq \delta/2. \end{aligned}$$

Thus, we have $P_{\mathcal{W}^*}(u_B, v) \in B_{\delta/2}(0)$ for any $(u_B, v) \in \mathcal{W}$. As a result, it holds for all $(u_B, v) \in \mathcal{W}$ that

$$\text{dist}((u_B, v), \mathcal{W}^*) = \text{dist}((u_B, v), \mathcal{W}^* \cap B_{\delta/2}(0)). \tag{32}$$

Since (x_B^*, y^*) satisfies (29), we have

$$\begin{aligned} &\|A_B x_B - b\|_2^2 + \|[A_B^T y - c_B]^+\|_2^2 + \left\| \left[\frac{1}{R_2} (c_B^T x_B - b^T y) \right]^+ \right\|_2^2 \\ &= \|A_B(x_B - x_B^*)\|_2^2 + \|[A_B^T(y - y^*)]^+\|_2^2 + \left\| \frac{1}{R_2} [c_B^T(x_B - x_B^*) - b^T(y - y^*)]^+ \right\|_2^2 \\ &= \|A_B u_B\|_2^2 + \|[A_B^T v]^+\|_2^2 + \left\| \frac{1}{R_2} [c_B^T u_B - b^T v]^+ \right\|_2^2. \end{aligned} \tag{33}$$

Furthermore, for any $(u_B, v) \in \mathcal{W}^* \cap B_{\frac{\delta}{2}}(0)$, we have $(x_B^* + u_B, y^* + v) \in \widetilde{\mathcal{X}}_B^* \times \widetilde{\mathcal{Y}}^*$, because

$$\begin{aligned} A_B(x_B^* + u_B) - b &= A_Bx_B^* - b + A_Bu_B = 0 \\ A_B^T(y^* + v) - c_B &= A_B^Ty^* - c_B + A_B^Tv \leq 0 \\ c_B^T(x_B^* + u_B) - b^T(y^* + v) &= c_B^Tx_B^* - b^Ty^* + c_B^Tu_B - b^Tv \leq 0 \\ x_{B_1}^* + u_{B_1} &\geq \frac{\delta}{2} \mathbf{1}_{|B_1|} > 0, \quad x_{B_2}^* + u_{B_2} = u_{B_2} \geq 0, \end{aligned}$$

thus it holds that

$$(x_B^*, y^*) + \mathcal{W}^* \cap B_{\frac{\delta}{2}}(0) \subset \widetilde{\mathcal{X}}_B^* \times \widetilde{\mathcal{Y}}^*. \tag{34}$$

Finally, by the definition of α_{L_1} , we have

$$\begin{aligned} &\|A_Bu_B\|_2^2 + \|[A_B^Tv]^+\|_2^2 + \|\left[\frac{1}{R_2}(c_B^Tu_B - b^Tv)\right]^+\|_2^2 \\ &\geq \alpha_{L_2}^2 \text{dist}^2((u_B, v), \mathcal{W}^*) \\ &= \alpha_{L_2}^2 \text{dist}^2((u_B, v), \mathcal{W}^* \cap B_{\delta/2}(0)) \\ &= \alpha_{L_2}^2 \text{dist}^2((x_B - x_B^*, y - y^*), (x_B^*, y^*) + \mathcal{W}^* \cap B_{\delta/2}(0) - (x_B^*, y^*)) \\ &= \alpha_{L_2}^2 \text{dist}^2((x_B, y), (x_B^*, y^*) + \mathcal{W}^* \cap B_{\delta/2}(0)) \\ &\geq \alpha_{L_2}^2 \text{dist}^2((x_B, y), \widetilde{\mathcal{X}}_B^* \times \widetilde{\mathcal{Y}}^*), \end{aligned} \tag{35}$$

where the first equality uses (32) and the last inequality utilizes (34). We finish the proof by combining (33) and (35). \square

Now we are ready to prove Theorem 3.

Proof of Theorem 3 Notice that $k > K$. It follows from Theorem 2 that $\|z^k - z^*\|_2 \leq \frac{\delta}{2}$ and $x_N^k = 0$. Denote the local normalized duality gap as

$$\rho_r^B(x_B, y) = \max_{(\hat{x}_B, \hat{y}) \in W_r(x_B, y)} \frac{L^B(x_B, \hat{y}) - L^B(\hat{x}_B, y)}{r},$$

where $L^B(x_B, y) = c_B^Tx_B - y^T A_Bx_B + b^Ty$. Then we have for any $r \in (0, \tilde{R})$ that

$$\begin{aligned} \alpha_{L_2} \text{dist}((x^k, y^k), \mathcal{X}^* \times \mathcal{Y}^*) &\leq \alpha_{L_2} \text{dist}((0_N, x_B^k, y^k), \{0_N\} \times \widetilde{\mathcal{X}}_B^* \times \widetilde{\mathcal{Y}}^*) \\ &= \alpha_{L_2} \text{dist}((x_B^k, y^k), \widetilde{\mathcal{X}}_B^* \times \widetilde{\mathcal{Y}}^*) \\ &\leq \left\| \begin{pmatrix} A_Bx_B^k - b \\ [A_B^Ty^k - c_B]^+ \\ \left[\frac{1}{R_2}(c_B^Tx_B^k - b^Ty^k)\right]^+ \end{pmatrix} \right\|_2 \\ &\leq \rho_r^B(x_B^k, y^k) = \rho_r(x^k, y^k), \end{aligned}$$

where the second inequality is due to definition of α_{L_2} and the last inequality utilizes Proposition 1 and $\|(x_B^k, y^k)\|_2 \leq \|(x^k, y^k)\|_2 \leq \|(x^k, y^k) - (x^*, y^*)\|_2 + \|(x^*, y^*)\|_2 \leq \delta + \|z^*\|_2 = R_2$. Applying part (b) of Theorem 1, it holds for any $k > K$ that

$$\begin{aligned} \|z^k - z^*\|_2 &\leq \sqrt{2s}\|z^k - z^*\|_{P_s} \\ &\leq 4\sqrt{2s} \exp\left(-\frac{k - K}{2\lceil 4e/(s^2\alpha_{L_2}^2)\rceil}\right) \text{dist}_{P_s}((x^K, y^K), \mathcal{X}^* \times \mathcal{Y}^*) \\ &\leq 8 \exp\left(-\frac{k - K}{2\lceil 4e/(s^2\alpha_{L_2}^2)\rceil}\right) \frac{\delta}{2}, \end{aligned}$$

which finishes the proof. □

5 Numerical experiments

In this section, we present numerical experiments that verify our theoretical results in the previous sections.

Dataset. In the experiments, we utilize the root-node LP relaxation of instances from the [MIPLIB 2017](#). We convert the instances to the following form:

$$\begin{aligned} \min_x \quad & c^T x \\ \text{s.t.} \quad & A_E x = b_E \\ & A_I x \leq b_I \\ & x \geq 0, \end{aligned} \tag{36}$$

and consider its primal–dual formulation

$$\min_{x \geq 0} \max_{y \leq 0} c^T x - y^T A x + b^T y, \tag{37}$$

where $A_E \in \mathbb{R}^{m_E \times n}$, $A_I \in \mathbb{R}^{m_I \times n}$, $A = \begin{pmatrix} A_E \\ A_I \end{pmatrix} \in \mathbb{R}^{(m_E+m_I) \times n}$ and $b = \begin{pmatrix} b_E \\ b_I \end{pmatrix} \in \mathbb{R}^{(m_E+m_I)}$. We choose the instances that can be solved by PDHG (after preconditioning, see below for more details) to 10^{-10} accuracy of KKT residual (see Progress metric for formal definition) within 3×10^5 iterations, and there are 50 such instances. Among all 50 instances, the number of variables ranges from 86 to 3648, the number of constraints ranges from 29 to 6972, and the number of nonzeros in the constraint matrices ranges from 376 to 22584. Notice that there is a difference between (36) and the standard form (1) we present in the paper due to the existence of the inequality constraints. Indeed, all of our theory can be generalized to this form without loss of generality. We choose to present the results for the standard form LP (1) because that leads to theoretical results in relatively clean forms.

Preprocessing. As a first-order method, PDHG suffers from slow convergence on many instances from MIPLIB due to the ill-conditioning nature of the instances. To overcome this issue, we leverage the diagonal preconditioning heuristic developed in PDLP [4]. More specifically, we rescale the constraint matrix A to $\tilde{A} = D_1 A D_2$ with positive definite diagonal matrices D_1 and D_2 . Both vectors b and c are correspondingly rescaled as $\tilde{b} = D_1 b$ and $\tilde{c} = D_2 c$. The matrices D_1 and D_2 are obtained by running 10 steps of Ruiz scaling [66] followed by an l_2 Pock-Chambolle scaling [62] on the constraint matrix A . More details about this preconditioning scheme can be found in [4].

Progress metric. We use KKT residual of (37), i.e., a combination of primal infeasibility, dual infeasibility and primal–dual gap, to measure the performance of current iterates. More formally, the KKT residual of (36) is given by

$$\text{KKT}(x, y) = \left\| \begin{pmatrix} A_E x - b_E \\ [A_I x - b_I]^+ \\ [-x]^+ \\ [A^T y - c]^+ \\ [y_I]^+ \\ [c^T x - b^T y]^+ \end{pmatrix} \right\|_2.$$

Computing identification time and δ . In the experiments, we run PDHG to 10^{-8} accuracy to obtain the converging optimal solution $z^* = (x^*, y^*)$. With z^* , we can then identify the non-basic solution set N , non-degenerate basic solution set B_1 , and the degenerate basic solution set B_2 with respect to z^* for the primal and dual variables, respectively, by

$$N^P = \{i \in [n] : c_i - A_i^T y^* > 0\}, B_1^P = \{i \in [n] \setminus N^P : x_i^* > 0\}, B_2^P = [n] \setminus N^P \setminus B_1^P,$$

and

$$N^D = \{j \in I : (b - Ax^*)_j > 0\}, B_1^D = \{j \in I \setminus N^D : y_j^* < 0\}, B_2^D = I \setminus N^D \setminus B_1^D.$$

Then we go backward from the last iteration to find the first moment when either N or B_1 changes. This iteration is referred to as the empirical identification moment. The non-degeneracy metric δ for (37) can be computed via

$$\delta = \min \left\{ \min_{i \in N^P} \frac{c_i - A_i^T y^*}{\|A\|_2}, \min_{i \in B_1^P} x_i^*, \min_{j \in N^D} \frac{(b - Ax^*)_j}{\|A\|_2}, \min_{j \in B_1^D} -y_j^* \right\}. \quad (38)$$

Results. Figure 4 presents KKT residual versus PDHG iterations on six representative instances. The orange vertical line in Fig. 4 represents the empirical identification moment, i.e., subsequently the active set will be fixed and identical to the limit optimal point. As illustrated in Fig. 4, there is a phase transition happening around the identification moment (orange line): PDHG identifies active variables in Stage I with a sublinear rate; PDHG has a fast linear convergence in Stage II after identification.

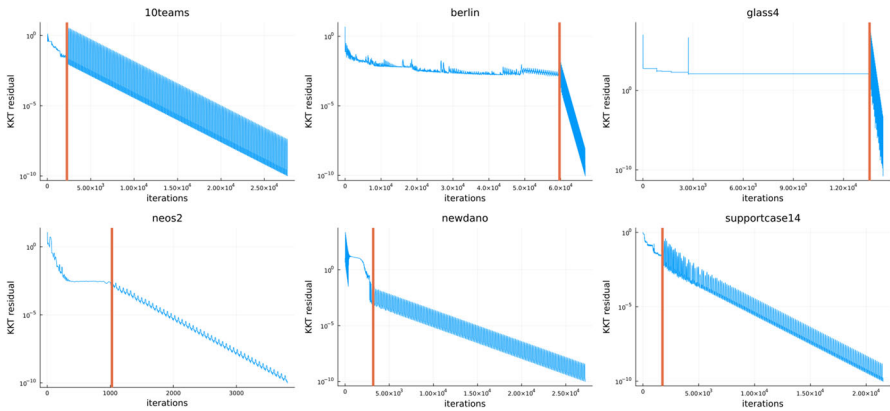


Fig. 4 Plots showing the KKT residual in log scale versus number of iterations of PDHG on six representative LP instances from MIPLIB

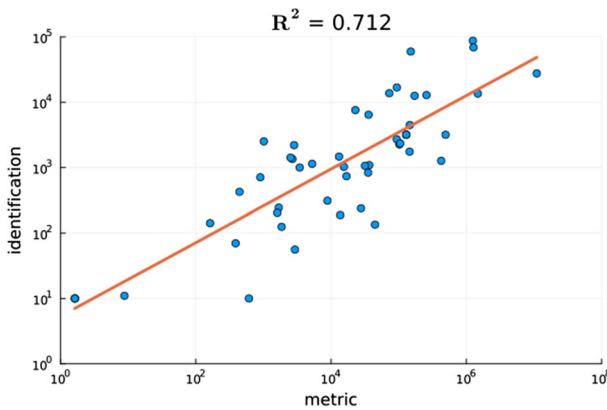


Fig. 5 Plot showing the iteration number at identification in log scale versus non-degeneracy metric R/δ in log scale on the 50 MIPLIB instances

Theorem 2 shows a bound on the identification time with the non-degeneracy parameter δ . Figure 5 presents the scatter plot of the metric R/δ versus the empirical iteration number for identification, where $R = 2\|z^0 - z^*\|_2 + 2\|z^*\|_2 + 1$ and δ is defined as (38). Each point in Fig. 5 represents an LP instance from our dataset. We observe from Fig. 5 that there is a strong correlation between the empirical identification time and the metric R/δ . Indeed, when running a linear regression model, the R^2 is 0.712, which means that 71.2% of the variability in the observed identification complexity is captured by this single quantity.

An interesting implication of our result is that degeneracy itself does not slow down the convergence, but the near-degeneracy of the non-degenerate part does. Figure 6 demonstrates this theoretical implication by preliminary numerical experiments. We choose three representative LP instances from the dataset. Then, we slightly perturb the instances by adding a small Gaussian noise in A , b , and c , then run PDHG on both the original and the perturbed instances. We use Gurobi to verify the feasibility of the

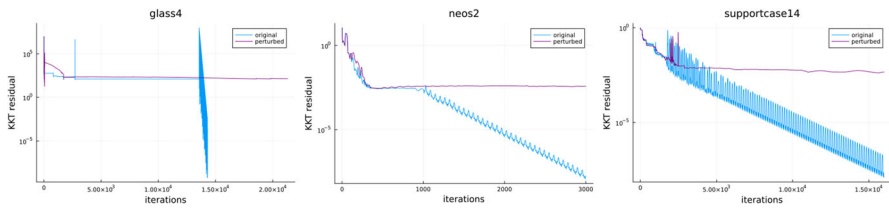


Fig. 6 Comparison of original and perturbed LP instances

three perturbed problems. Figure 6 plots the KKT residual and the number of PDHG iterations for the original LP and the perturbed LP. For each instance in Fig. 6, the blue line represents the performance of PDHG on the original LP, and the purple line is on the perturbed LP. While the original LP and the perturbed LP are almost identical (because the perturbation is tiny), PDHG converges much faster on the original LP than the perturbed LP. Indeed, almost all practical LP, including the three instances in Fig. 6, are degenerate. Adding the tiny perturbation makes the problem non-degenerate. This showcases that degeneracy itself does not slow down the convergence of PDHG, but near-degeneracy can make the instance much harder to solve, which verifies our theory.

In summary, the numerical experiment verifies our geometric understandings of PDHG on LPs: (1) the convergence behavior of PDHG for LP have two distinct stages, that is, Stage I for active variable identification and Stage II for local linear convergence; (2) the non-degeneracy metric δ plays a crucial role in the length of the first stage; (3) degeneracy itself does not deteriorate the convergence of first-order methods, but near-to-degeneracy does.

6 Conclusion and future directions

In conclusion, this paper aims to bridge the gap between the practical success and loose complexity bound of PDHG for LP, and to identify geometric quantities of LP that drive the performance of the algorithm. To achieve this, we show that the behaviors of PDHG for LP have two stages: in the first sublinear stage, the algorithm identifies the non-basic and non-degenerate basic set (i.e., the set N and B_1) and we present the complexity result of identification in terms of the near-degeneracy metric δ ; in the second linear stage, the algorithm effectively solves a homogeneous linear inequality system, and the linear convergence rate is driven by the local sharpness parameter of the system. Compared to existing complexity results of PDHG on LP, our results do not depend on the global Hoffman’s constant of the KKT system, which is known to be too loose in the literature. Such two-stage behavior is also tied with the concept of partial smoothness in the non-smooth optimization literature. We introduce a new framework of complexity analysis without assuming the “irrepresentable” non-degeneracy condition.

We end the paper by presenting a few interesting open questions for further investigation:

- Extensions to non-smooth optimization.** The partial smoothness literature in non-smooth optimization always assumes the non-degeneracy condition. Unfortunately, the iterates of first-order methods usually converge to a degenerate solution, thus violating this condition. A typical argument for avoiding this issue in theory is by adding a small perturbation to the original problem. However, as shown in our numerical experiments (Fig. 6), adding the small perturbation can significantly slow down the convergence of the algorithm. Instead, we show in this paper that degeneracy itself does not slow down the convergence in the context of LP. How to extend this result to general non-smooth optimization can be an interesting future direction.
- Extensions to other primal–dual algorithms.** We believe the two-stage behaviors are not limited to PDHG. Indeed, most of the theoretical results can be extended to ADMM. This should provide new theoretical understandings for ADMM-based solvers, such as OSQP [67] and SCS [58].
- Extensions to the restarted algorithm.** PDLP utilizes a new technique, restarting, to obtain the optimal linear convergence rate with global sharpness condition. We believe our theoretical results can be extended to the restarted algorithms, and we leave it as future work.
- Extensions to infeasible problems.** In the paper, we consider the case when the LP instance is feasible and bounded. Such two-stage behaviors are also numerically observed in infeasible problems. How to theoretically analyze it without the non-degeneracy condition can be another interesting question.

Acknowledgements This work was supported by the ONR Program.

Funding Open Access funding provided by the MIT Libraries.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Proof of Theorem 1

In this section, we provide a proof to Theorem 1. We begin with a lemma that bridges the distance to converging point with the distance to optimality.

Lemma 8 *Consider the iterations $\{z^k\}_{k=0}^{\infty}$ of PDHG. Suppose the step-size $s < \frac{1}{\|A\|}$. Then it holds for any $k \geq 0$ that*

$$\|z^k - z^*\|_{P_s} \leq 2 \text{dist}_{P_s}(z^k, \mathcal{Z}^*),$$

where $z^* = \lim_{k \rightarrow \infty} z^k$.

Proof By non-expansiveness of the iterates, we have

$$\|z^n - \text{proj}_{\mathcal{Z}^*}(z^k)\|_{P_s} \leq \|z^k - \text{proj}_{\mathcal{Z}^*}(z^k)\|_{P_s} = \text{dist}_{P_s}(z^k, \mathcal{Z}^*).$$

Let $n \rightarrow \infty$. It holds that

$$\|z^* - \text{proj}_{\mathcal{Z}^*}(z^k)\|_{P_s} \leq \text{dist}_{P_s}(z^k, \mathcal{Z}^*).$$

Thus we finish the proof by triangle inequality

$$\|z^k - z^*\|_{P_s} \leq \|z^k - \text{proj}_{\mathcal{Z}^*}(z^k)\|_{P_s} + \|\text{proj}_{\mathcal{Z}^*}(z^k) - z^*\|_{P_s} \leq 2\text{dist}_{P_s}(z^k, \mathcal{Z}^*).$$

□

It turns out that distance to optimality has linear convergence under sharpness (8).

Lemma 9 Consider the iterations $\{z_k\}_{k=0}^\infty$ of PDHG (Algorithm 1) to solve a convex-concave primal–dual problem. Suppose the step-size $s \leq \frac{1}{2\|A\|}$, and the primal–dual problem satisfies sharpness condition (8) on a set \mathcal{S} that contains $\{z_k\}_{k=0}^\infty$. Then, it holds for any iteration $k \geq 0$ that

$$\text{dist}_{P_s}^2(z^k, \mathcal{Z}^*) \leq \exp\left(1 - \frac{k}{\lceil 4e/(s\alpha)^2 \rceil}\right) \text{dist}_{P_s}^2(z^0, \mathcal{Z}^*)$$

Proof By [49, Proposition 1] and Lemma 2, we know the α -sharpness condition implies that the following equation holds for any $k \geq 0$

$$\text{dist}_{P_s}(z^k, \mathcal{Z}^*) \leq \frac{2}{s\alpha} \text{dist}_{P_{s-1}}(0, \mathcal{F}(z^k)) \tag{39}$$

Then the proof is similar to [50, Theorem 2]. For any iteration $k \geq 1$, suppose $c \lceil \frac{4e}{s^2\alpha^2} \rceil \leq k < (c + 1) \lceil \frac{4e}{s^2\alpha^2} \rceil$ for a non-negative integer c .

$$\begin{aligned} \text{dist}_{P_s}^2(z^k, \mathcal{Z}^*) &\leq \frac{4}{s^2\alpha^2} \text{dist}_{P_{s-1}}^2(0, \mathcal{F}(z^k)) \\ &\leq \frac{4}{s^2\alpha^2} \frac{1}{\lceil 4e/(s^2\alpha^2) \rceil} \text{dist}_{P_s}^2(z^{k-\lceil 4e/(s\alpha)^2 \rceil}, \mathcal{Z}^*) \\ &\leq \frac{1}{e} \text{dist}_{P_s}^2(z^{k-\lceil 4e/(s\alpha)^2 \rceil}, \mathcal{Z}^*) \\ &\leq \dots \\ &\leq \left(\frac{1}{e}\right)^c \text{dist}_{P_s}^2(z^{k-c\lceil 4e/(s\alpha)^2 \rceil}, \mathcal{Z}^*) \\ &\leq \left(\frac{1}{e}\right)^c \text{dist}_{P_s}^2(z^0, \mathcal{Z}^*) \\ &\leq \exp\left(1 - \frac{k}{\lceil 4e/(s\alpha)^2 \rceil}\right) \text{dist}_{P_s}^2(z^0, \mathcal{Z}^*) \end{aligned}$$

where the second inequality follows from [50, Theorem 1]. □

Proof of Theorem 1 (1) The result follows from [50, Theorem 1].

(2) From Lemma 9, we have

$$\text{dist}_{P_s}(z^k, \mathcal{Z}^*) \leq \exp\left(\frac{1}{2} - \frac{k}{2\lceil 4e/(s\alpha)^2 \rceil}\right) \text{dist}_{P_s}(z^0, \mathcal{Z}^*). \tag{40}$$

Therefore we reach

$$\begin{aligned} \|z^k - z^*\|_{P_s} &\leq 2\text{dist}_{P_s}(z^k, \mathcal{Z}^*) \\ &\leq 2 \exp\left(\frac{1}{2} - \frac{k}{2\lceil 4e/(s\alpha)^2 \rceil}\right) \text{dist}_{P_s}(z^0, \mathcal{Z}^*) \\ &\leq 4 \exp\left(-\frac{k}{2\lceil 4e/(s\alpha)^2 \rceil}\right) \text{dist}_{P_s}(z^0, \mathcal{Z}^*) \end{aligned}$$

where the first inequality is exactly Lemma 8 and the second one utilizes (40). □

B Examples to illustrate the looseness of global sharpness constant

In this section, we present a series of simple LP instances which exhibits very bad global sharpness behaviors (so as the Hoffman constant), whereas the three quantities we use in the analysis, δ , α_{L_1} and α_{L_2} , are of reasonable sizes. This indicates that the previous complexity result based on Hoffman constant argument [6] can be arbitrarily loose, and may not be able to explain the behaviors of the algorithm.

Consider a two dimension dual LP parameterized by $\kappa > 0$:

$$\begin{aligned} \max \quad &y_2 \\ \text{s.t.} \quad &y_1 \geq 0, \quad y_2 \leq 1 \\ &\kappa y_1 - y_2 \leq \kappa, \end{aligned} \tag{41}$$

where the primal form (1) is given by

$$A = \begin{pmatrix} 0 & -1 & \kappa \\ 1 & 0 & -1 \end{pmatrix}, \quad b = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad c = \begin{pmatrix} 1 \\ 0 \\ \kappa \end{pmatrix}. \tag{42}$$

The feasible region of (41) is plotted in Fig. 7. It is easy to check that there are multiple dual optimal solutions given by $y^* = (y_1, 1)$ with $0 \leq y_1 \leq 1 + \frac{1}{\kappa}$, and the optimal primal solution is unique given by $x^* = (1, 0, 0)$. Denote α the sharpness constant (i.e., the reciprocal of the Hoffman constant) of the KKT system

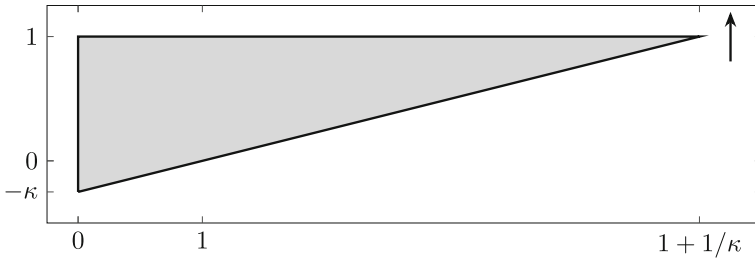


Fig. 7 Illustration of the LP instance (41)

Table 1 Comparison of global sharpness and metrics proposed for small κ

κ	α	α_{L_1}	α_{L_2}	δ
$= 10^{-10}$	$\leq 10^{-10}$	≥ 0.00821	≥ 0.0276	≥ 0.363

$$Ax = b, A^T y \leq c, x \geq 0, \frac{1}{R}(c^T x - b^T y) \leq 0,$$

where (A, b, c) is given in (42) and R is essentially the distance between initial point to origin. This is the term that appears in the complexity of previous results [6]. Table 1 presents bounds on $\alpha, \alpha_{L_1}, \alpha_{L_2}$ and δ for the LP when κ is set to 10^{-10} . The value of α can be obtained by noticing that $\alpha \leq \kappa$ (see Lemma 10 for a proof). The bound of α_{L_1} and α_{L_2} are obtained from the algorithm presented in [61]. The bound on δ can be computed via Lemma 11. As we can see, for this specific example, the value of α is tiny, while the value of $\alpha_{L_1}, \alpha_{L_2}$ and δ are of reasonable sizes.

Lemma 10 $\alpha \leq \kappa$.

Proof By definition of sharpness constant, it holds for any z that

$$\alpha \text{dist}(z, \mathcal{Z}^*) \leq \left\| \begin{pmatrix} Ax - b \\ [A^T y - c]^+ \\ \frac{1}{R}[c^T x - b^T y]^+ \end{pmatrix} \right\|_2,$$

where $\mathcal{Z}^* = \{(1, 0, 0, y_1, 1) \mid 0 \leq y_1 \leq 1 + \frac{1}{\kappa}\}$. Then for any $\zeta > 0$, consider $z = (1, 0, 0, \frac{1+\kappa}{\kappa} + \zeta, 1)$ and we have

$$\alpha^2 \zeta^2 \leq \left(\kappa \left(\frac{1+\kappa}{\kappa} + \zeta \right) - 1 - \kappa \right)^2,$$

which implies

$$\alpha^2 \leq \kappa^2.$$

□

Lemma 11 Suppose $\kappa \leq 0.1$ and the initial point z^0 satisfies $\|z^0 - \tilde{z}\|_2 \leq 1.5$, where $\tilde{z} = (1, 0, 0, 4, 1) \in \mathcal{Z}^*$. Then it holds that

$$\delta \geq \frac{4}{11}.$$

Proof From the non-expansiveness of PDHG iterates and Lemma 2, we have

$$\|z^k - \tilde{z}\|_2 \leq \sqrt{2s} \|z^k - \tilde{z}\|_{P_s} \leq \sqrt{2s} \|z^0 - \tilde{z}\|_{P_s} \leq 2 \|z^0 - \tilde{z}\|_2 \leq 3.$$

Denote $z^* = (x^*, y^*)$ the converging optimal point of $\{z^k\}_{k=0}^\infty$. Then it holds that

$$0 < 1 \leq y_1^* \leq 7 < 1 + \frac{1}{\kappa}.$$

Thus

$$N = \{2, 3\}, \quad B_1 = \{1\}, \quad B_2 = \emptyset,$$

and

$$c_2 - A_2^T y^* \geq 1, \quad c_3 - A_3^T y^* \geq \frac{1 + \frac{1}{\kappa} - 7}{1 + \frac{1}{\kappa}} \geq \frac{4}{11}, \quad x_1^* = 1,$$

where $\frac{1 + \frac{1}{\kappa} - 7}{1 + \frac{1}{\kappa}} \geq \frac{4}{11}$ by $\kappa \leq 0.1$. Note that $\|A\|_2 \leq \|A\|_F \leq 2$ and we reach at

$$\delta = \min \left\{ \min_{i \in N} \frac{c_i - A_i^T y^*}{\|A\|_2}, \min_{i \in B_1} x_i^* \right\} \geq \frac{4}{22} = \frac{2}{11}.$$

□

References

1. Acharya, A., Gao, S., Ocejjo, B., Basu, K., Saha, A., Selvaraj, K., Mazumdar, R., Agrawal, P., Gupta, A.: Promoting inactive members in edge-building marketplace. In: Companion Proceedings of the ACM Web Conference, vol. 2023, pp. 945–949 (2023)
2. Alacaoglu, A., Fercoq, O., Cevher, V.: On the convergence of stochastic primal-dual hybrid gradient. *SIAM J. Optim.* **32**(2), 1288–1318 (2022)
3. Anderson, R.I., Fok, R., Scott, J.: Hotel industry efficiency: an advanced linear programming examination. *Am. Bus. Rev.* **18**(1), 40 (2000)
4. Applegate, D., Díaz, M., Hinder, O., Lu, H., Lubin, M., O’Donoghue, B., Schudy, W.: Practical large-scale linear programming using primal-dual hybrid gradient. In: *Advances in Neural Information Processing Systems*, vol. 34 (2021)
5. Applegate, D., Díaz, M., Lu, H., Lubin, M.: Infeasibility detection with primal-dual hybrid gradient for large-scale linear programming, arXiv preprint [arXiv:2102.04592](https://arxiv.org/abs/2102.04592) (2021)
6. Applegate, D., Hinder, O., Haihao, L., Lubin, M.: Faster first-order primal-dual methods for linear programming using restarts and sharpness. *Math. Program.* **201**(1–2), 133–184 (2023)
7. Basu, K., Ghoting, A., Mazumdar, R., Pan, Y.: ECLIPSE: an extreme-scale linear program solver for web-applications. In: *International Conference on Machine Learning*, pp. 704–714. PMLR (2020)

8. Beck, A.: First-Order Methods in Optimization. SIAM, Philadelphia (2017)
9. Bowman, E.H.: Production scheduling by the transportation method of linear programming. *Oper. Res.* **4**(1), 100–103 (1956)
10. Boyd, S.P., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, Cambridge (2004)
11. Chambolle, A., Ehrhardt, M.J., Richtárik, P., Schonlieb, C.-B.: Stochastic primal-dual hybrid gradient algorithm with arbitrary sampling and imaging applications. *SIAM J. Optim.* **28**(4), 2783–2808 (2018)
12. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vis.* **40**(1), 120–145 (2011)
13. Chambolle, A., Pock, T.: On the ergodic convergence rates of a first-order primal-dual algorithm. *Math. Program.* **159**(1), 253–287 (2016)
14. Charnes, A., Cooper, W.W.: The stepping stone method of explaining linear programming calculations in transportation problems. *Manag. Sci.* **1**(1), 49–69 (1954)
15. Condat, L.: A primal-dual splitting method for convex optimization involving Lipschitzian, proximable and linear composite terms. *J. Optim. Theory Appl.* **158**(2), 460–479 (2013)
16. Daniilidis, A., Drusvyatskiy, D., Lewis, A.S.: Orthogonal invariance and identifiability. *SIAM J. Matrix Anal. Appl.* **35**(2), 580–598 (2014)
17. Dantzig, G.B.: Linear programming. *Oper. Res.* **50**(1), 42–47 (2002)
18. Dantzig, G.B.: *Linear Programming and Extensions*, vol. 48. Princeton University Press, Princeton (1998)
19. Davis, D., Drusvyatskiy, D., Jiang, L.: Subgradient methods near active manifolds: saddle point avoidance, local convergence, and asymptotic normality, arXiv preprint [arXiv:2108.11832](https://arxiv.org/abs/2108.11832) (2021)
20. Demanet, L., Zhang, X.: Eventual linear convergence of the Douglas–Rachford iteration for basis pursuit. *Math. Comput.* **85**(297), 209–238 (2016)
21. Deng, Q., Feng, Q., Gao, W., Ge, D., Jiang, B., Jiang, Y., Liu, J., Liu, T., Xue, C., Ye, Y. et al.: New developments of ADMM-based interior point methods for linear programming and conic programming, arXiv preprint [arXiv:2209.01793](https://arxiv.org/abs/2209.01793) (2022)
22. Eckstein, J, Bertsekas, D.P. et al.: An alternating direction method for linear programming (1990)
23. Esser, E., Zhang, X., Chan, T.F.: A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM J. Imaging Sci.* **3**(4), 1015–1046 (2010)
24. Fadili, J, Garrigos, G, Malick, J., Peyré, G.: Model consistency for learning with mirror-stratifiable regularizers. In: *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1236–1244. PMLR (2019)
25. Fadili, J., Malick, J., Peyré, G.: Sensitivity analysis for mirror-stratifiable convex functions. *SIAM J. Optim.* **28**(4), 2975–3000 (2018)
26. Fercoq, O.: Quadratic error bound of the smoothed gap and the restarted averaged primal-dual hybrid gradient, arXiv preprint [arXiv:2206.03041](https://arxiv.org/abs/2206.03041) (2022)
27. Goldstein, T., Li, M., Yuan, X.: Adaptive primal-dual splitting methods for statistical learning and image processing. In: *Advances in neural information processing systems*, pp. 2089–2097 (2015)
28. Güler, O., Ye, Y.: Convergence behavior of interior-point algorithms. *Math. Program.* **60**(1–3), 215–228 (1993)
29. Hanssmann, F., Hess, S.W.: A linear programming approach to production and employment scheduling. *Manag. Sci.* **1**, 46–51 (1960)
30. Hare, W.L., Lewis, A.S.: Identifying active constraints via partial smoothness and prox-regularity. *J. Convex Anal.* **11**(2), 251–266 (2004)
31. Hare, W.L., Lewis, A.S.: Identifying active manifolds. *Algorithmic Oper. Res.* **2**(2), 75–82 (2007)
32. He, B., Yuan, X.: Convergence analysis of primal-dual algorithms for a saddle-point problem: from contraction perspective. *SIAM J. Imaging Sci.* **5**(1), 119–149 (2012)
33. Hoffman, A.J.: On approximate solutions of systems of linear inequalities. *J. Res. Natl. Bureau Stand.* **49**, 263–265 (1952)
34. Karmarkar, N.: A new polynomial-time algorithm for linear programming. In: *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, pp. 302–311 (1984)
35. Lee, S., Wright, S.J., Bottou, L.: Manifold identification in dual averaging for regularized stochastic online learning. *J. Mach. Learn. Res.* **13**(6), 1705–1744 (2012)
36. Lewis, A.S., Liang, J., Tian, T.: Partial smoothness and constant rank. *SIAM J. Optim.* **32**(1), 276–291 (2022)
37. Lewis, A.S., Wright, S.J.: Identifying activity. *SIAM J. Optim.* **21**(2), 597–614 (2011)

38. Lewis, A.S., Wright, S.J.: A proximal method for composite minimization. *Math. Program.* **158**(1), 501–546 (2016)
39. Lewis, A.S., Wylie, C.: Active-set newton methods and partial smoothness. *Math. Oper. Res.* **46**(2), 712–725 (2021)
40. Lewis, A.S., Zhang, S.: Partial smoothness, tilt stability, and generalized Hessians. *SIAM J. Optim.* **23**(1), 74–94 (2013)
41. Li, X., Sun, D., Toh, K.-C.: An asymptotically superlinearly convergent semismooth Newton augmented Lagrangian method for linear programming. *SIAM J. Optim.* **30**(3), 2410–2440 (2020)
42. Liang, J., Fadili, J., Peyré, G.: Local linear convergence of forward–backward under partial smoothness. In: *Advances in neural Information Processing Systems*, vol. 27 (2014)
43. Liang, J., Fadili, J., Peyré, G.: Activity identification and local linear convergence of forward–backward-type methods. *SIAM J. Optim.* **27**(1), 408–437 (2017)
44. Liang, J., Fadili, J., Peyré, G.: Local convergence properties of Douglas–Rachford and alternating direction method of multipliers. *J. Optim. Theory Appl.* **172**(3), 874–913 (2017)
45. Liang, J., Fadili, J., Peyré, G.: Local linear convergence analysis of primal-dual splitting methods. *Optimization* **67**(6), 821–853 (2018)
46. Lin, T., Ma, S., Ye, Y., Zhang, S.: An ADMM-based interior-point method for large-scale linear programming. *Optim. Methods Softw.* **36**(2–3), 389–424 (2021)
47. Liu, Q., Van Ryzin, G.: On the choice-based linear programming model for network revenue management. *Manuf. Service Oper. Manag.* **10**(2), 288–310 (2008)
48. Liu, Y., Xu, Y., Yin, W.: Acceleration of primal-dual methods by preconditioning and simple subproblem procedures. *J. Sci. Comput.* **86**(2), 1–34 (2021)
49. Lu, H., Yang, J.: Nearly optimal linear convergence of stochastic primal-dual methods for linear programming, arXiv preprint [arXiv:2111.05530](https://arxiv.org/abs/2111.05530) (2021)
50. Lu, H., Yang, J.: On the infimal sub-differential size of primal-dual hybrid gradient method, arXiv preprint [arXiv:2206.12061](https://arxiv.org/abs/2206.12061) (2022)
51. Lu, H., Yang, J.: On a unified and simplified proof for the ergodic convergence rates of ppm, pdhg and admm, arXiv preprint [arXiv:2305.02165](https://arxiv.org/abs/2305.02165) (2023)
52. Malitsky, Y., Pock, T.: A first-order primal-dual algorithm with linesearch. *SIAM J. Optim.* **28**(1), 411–432 (2018)
53. Manne, A.S.: Linear programming and sequential decisions. *Manag. Sci.* **6**(3), 259–267 (1960)
54. Mirrokni, V.: Google research, 2022 & beyond: Algorithmic advances, <https://ai.googleblog.com/2023/02/google-research-2022-beyond-algorithmic.html>, (2023-02-10)
55. Nesterov, Y.: *Introductory Lectures on Convex Optimization: A Basic Course*, vol. 87. Springer Science & Business Media, New York (2013)
56. Oberlin, C., Wright, S.J.: Active set identification in nonlinear programming. *SIAM J. Optim.* **17**(2), 577–605 (2006)
57. O’Donoghue, B.: Operator splitting for a homogeneous embedding of the linear complementarity problem. *SIAM J. Optim.* **31**(3), 1999–2023 (2021)
58. O’Donoghue, B., Chu, E., Parikh, N., Boyd, S.: Conic optimization via operator splitting and homogeneous self-dual embedding. *J. Optim. Theory Appl.* **169**(3), 1042–1068 (2016)
59. O’Connor, D., Vandenberghe, L.: On the equivalence of the primal-dual hybrid gradient method and Douglas–Rachford splitting. *Math. Program.* **179**(1), 85–108 (2020)
60. Pena, J., Vera, J.C., Zuluaga, L.F.: New characterizations of Hoffman constants for systems of linear constraints. *Math. Program.* **187**(1), 79–109 (2021)
61. Peña, J.F.: An easily computable upper bound on the Hoffman constant for homogeneous inequality systems. *Comput. Optim. Appl.* 1–13 (2023)
62. Pock, T., Chambolle, A.: Diagonal preconditioning for first order primal-dual algorithms in convex optimization. In: 2011 International Conference on Computer Vision, pp. 1762–1769. IEEE (2011)
63. Poon, C., Liang, J., Schoenlieb, C.: Local convergence properties of SAGA/Prox-SVRG and acceleration. In: *International Conference on Machine Learning*, pp. 4124–4132. PMLR (2018)
64. Ramanath, R., Keerthi, S.S., Pan, Y., Salamat, K., Basu, K.: Efficient vertex-oriented polytopic projection for web-scale applications. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, pp. 3821–3829 (2022)
65. Renegar, J.: Incorporating condition measures into the complexity theory of linear programming. *SIAM J. Optim.* **5**(3), 506–524 (1995)

66. Ruiz, D.: A scaling algorithm to equilibrate both rows and columns norms in matrices, Tech. report, CM-P00040415, (2001)
67. Stellato, B., Banjac, G., Goulart, P., Bemporad, A., Boyd, S.: OSQP: An operator splitting solver for quadratic programs. *Math. Program. Comput.* **12**(4), 637–672 (2020)
68. Vladarean, M.-L., Malitsky, Y., Cevher, V.: A first-order primal-dual method with adaptivity to local smoothness. In: *Advances in neural information processing systems*, vol. 34, pp. 6171–6182 (2021)
69. Wang, S., Shroff, N.: A new alternating direction method for linear programming. In: *Advances in neural information processing systems*, vol. 30 (2017)
70. Wright, S.J.: Identifiable surfaces in constrained optimization. *SIAM J. Control. Optim.* **31**(4), 1063–1079 (1993)
71. Wright, S.J.: *Primal-Dual Interior-point Methods*. SIAM, Philadelphia (1997)
72. Ye, Y.: On the finite convergence of interior-point algorithms for linear programming. *Math. Program.* **57**(1–3), 325–335 (1992)
73. Yen, I.E.-H., Zhong, K., Hsieh, C.-J., Ravikumar, P. K., Dhillon, I. S.: Sparse linear programming via primal and dual augmented coordinate descent. In: *Advances in neural information processing systems*, vol. 28 (2015)
74. Zhu, M., Chan, T.: An efficient primal-dual hybrid gradient algorithm for total variation image restoration. *UCLA Cam Report* **34**, 8–34 (2008)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.