Accessible Complexity Bounds for Restarted PDHG on Linear Programs with a Unique Optimizer

Zikai Xiong
MIT Operations Research Center, 77 Massachusetts Avenue, Cambridge, MA 02139, USA, zikai@mit.edu.

Abstract. The restarted primal-dual hybrid gradient method (rPDHG) has recently emerged as an important tool for solving large-scale linear programs (LPs). For LPs with unique optima, we present an iteration bound of $O\left(\kappa\Phi \cdot \ln\left(\frac{\kappa\Phi\|w^\star\|}{\varepsilon}\right)\right)$, where ε is the target tolerance, κ is the standard matrix condition number, $\|w^\star\|$ is the norm of the optimal solution, and Φ is a geometric condition number of the LP sublevel sets. This iteration bound is "accessible" in the sense that computing it is typically no more difficult than computing the optimal solution itself. Indeed, we present a closed-form and tractably computable expression for Φ. This enables an analysis of the "two-stage performance" of rPDHG: we show that the first stage identifies the optimal basis in $O\left(\kappa\Phi \cdot \ln(\kappa\Phi)\right)$ iterations, and the second stage computes an ε-optimal solution in $O\left(\|B^{-1}\|\|A\| \cdot \ln\left(\frac{\xi}{\varepsilon}\right)\right)$ additional iterations, where A is the constraint matrix, B is the optimal basis and ξ is the smallest nonzero in the optimal solution. Furthermore, computational tests mostly confirm the tightness of our iteration bounds. We also show a reciprocal relation between the iteration bound and stability under data perturbation, which is also equivalent to (i) proximity to multiple optima, and (ii) the LP sharpness of the instance. Finally, we analyze an "optimized" primal-dual reweighting which offers some intuition concerning the step-size heuristics used in practice.

Funding: This research was supported by AFOSR Grant No. FA9550-22-1-0356.
Key words: Linear optimization, First-order method, Convergence guarantees, Condition number
History: Updated: December 22, 2025

1. Introduction

Linear programming has been a cornerstone of optimization since the 1950s, with far-reaching applications across diverse fields, including economics (see, e.g., Greene [27]), transportation (see, e.g., Charnes and Cooper [13]), manufacturing (see, e.g., Bowman [10], Hanssmann and Hess [30]), computer science (see, e.g., Cormen et al. [15]), and medicine (see, e.g., Wagner et al. [62]) among many others (see, e.g., Dantzig [16]). Linear programming algorithms have also been extensively researched in the past several decades. Almost all linear programming algorithms to date are based on either simplex methods or interior-point methods (IPMs). These classic methods form the foundation of modern solvers due to their reliability and robustness in providing high-quality solutions. However, both of them require repeatedly solving linear systems at each iteration using matrix factorizations, whose cost usually grows superlinearly in the size of the instance (as measured in the problem dimensions or the number of nonzeros in the data), unless the data sparsity pattern is suitable. Consequently, as problem size increases, these methods often become computationally impractical. Furthermore, matrix factorizations cannot efficiently leverage modern computational architectures, such as parallel computing on graphics processing units (GPUs). For these reasons, first-order methods (FOMs) are emerging as attractive solution algorithms because they are "matrix-free," meaning they require no or perhaps only very few matrix factorizations, while their primary computational cost lies just in computing matrix-vector products when computing gradients and related quantities. As a result, FOMs can often exploit sparse linear algebra and parallel architectures (e.g., GPUs) more readily, and their per-iteration cost typically scales roughly linearly with the number of nonzeros.

The restarted primal-dual hybrid gradient method (rPDHG) has emerged as a particularly successful FOM for solving Linear Programs (LPs). It directly addresses the saddle-point formulation of LP (see Applegate et al. [6]), automatically detects infeasibility (see Applegate et al. [5]), and has natural extensions to conic linear programs in Xiong and Freund [67] and convex quadratic programs in Lu and Yang [48] and Huang et al. [33]. This algorithm has led to various implementations on CPUs (PDLP by Applegate et al. [4]) and GPUs (cuPDLP by Lu and Yang [46] and cuPDLP-C by Lu et al. [50]), equipped with several effective heuristics. Notably, the performance of the GPU implementations has surpassed classic algorithms (simplex methods and IPMs) on a significant number of problem instances as shown by Lu and Yang [46] and Lu et al. [50]. The strong practical performance has also sparked considerable industrial interest from mathematical

optimization software companies. To date, rPDHG has been integrated into the state-of-art commercial solvers COPT (see [11]), Xpress (see [21]) and Gurobi (see [29]) as a new base algorithm for LPs, complementing simplex methods and IPMs. It has also been added to Google OR-Tools (see Applegate et al. [4]), HiGHS (see Ge et al. [24]), and NVIDIA cuOpt (see [53]). With rPDHG, many problems previously considered too large-scale for classic algorithms are now solvable; for instance, it is reported by Mirrokni [52] that a distributed version of rPDHG has been used to solve practical LP instances with more than 9.2×10^{10} nonzero entries in the constraint matrix, a scale far beyond the capabilities of traditional methods. Another example is a large-scale benchmark instance called zib03. This instance, which took 16.5 hours to solve in 2021 as reported by Koch et al. [34], can now be solved in 15 minutes using rPDHG by Lu et al. [50]. Furthermore, it has recently unlocked the potentials of large-scale linear programming in real applications, including making targeted marketing policies by Lu et al. [42], solving large-scale integer programming instances by De Rosa and Khajavirad [17], and optimizing data center network traffic engineering by Lu and Applegate [40], the latter of which has been deployed in Google's production environment.

Despite the strong performance of rPDHG on many LP instances, certain aspects of its practical behavior remain poorly understood. Indeed, rPDHG sometimes performs poorly, even for some very small LP instances. Additionally, minor data perturbation of some easily solvable instances can lead to instances with substantially increased computational cost. Also, it has been observed that rPDHG often exhibits a "two-stage performance" phenomenon in which the second stage exhibits much faster local convergence, but this phenomenon has not been adequately explained or otherwise addressed by suitable theory.

To better understand the underlying behavior of rPDHG, it is important to have theory that is consistent with practical performance. However, many aspects of the existing theory cannot be adequately evaluated for practical relevance due to the difficulty of actually computing the quantities in the theoretical bounds. Applegate et al. [6] establish the linear convergence rate of rPDHG using the global Hoffman constant of the matrix K of the KKT system corresponding to the LP instance. Roughly speaking, the Hoffman constant is equal to the reciprocal of the smallest nonzero singular values of the submatrices of K, of which there are exponentially many (see Pena et al. [56]). While intuition suggests that the Hoffman constant is itself an overly conservative quantity in the computational complexity, we do not know this empirically on nontrivial LP instances because the Hoffman constant is not computable in reasonable time. Xiong and Freund [65] provide a tighter computational guarantee for rPDHG using two natural properties of the LP problem: LP sharpness and the "limiting error ratio." Furthermore, for LPs with extremely poor sharpness and the broader family of conic LPs, Xiong and Freund [67] provide computational guarantees for rPDHG based on three geometric measures of the primal-dual (sub)level set geometry. In addition, Lu and Yang [47] study the vanilla primal-dual hybrid gradient method (PDHG) using a trajectory-based analysis approach, and show the two-stage performance of PDHG based on the Hoffman constant of a smaller linear system. However, despite these studies, none provides an iteration bound that is reasonably easy to compute, and so we cannot ascertain the extent to which any of these iteration bounds align with computational practice. To compute the iteration bounds, all existing works require prohibitively expensive operations, such as directly computing Hoffman constants (e.g., Applegate et al. [6], Lu and Yang [47]), solving multiple additional optimization problems (e.g., Lu and Yang [47], Xiong and Freund [65, 67]), or running a first-order method beforehand to obtain the solution trajectory (e.g., Lu and Yang [47]).

Because existing iteration bounds are hard to evaluate in practice, it remains unclear how well current theory matches computational practice. This also makes it harder to analyze rPDHG and improve its empirical performance. More broadly, without a convenient, practically evaluable iteration bound, it is difficult to understand rPDHG's behavior on specific LP families, to provide theoretical support for effective heuristics, or to develop further practical enhancements.

This paper aims to make progress on the above issues by posing and trying to answer the following questions:

- Can we derive an *accessible* iteration bound for rPDHG? By *accessible*, we mean that computing the quantities in the bound should rely on intrinsic properties of the data of the LP instance and its optimal solutions. The computation is typically not more costly than solving the LP instance itself. (For example, the norm of the optimal solution is accessible, whereas Hoffman constants typically are not.)

- If we have an accessible iteration bound, can we use the bound to provide deeper insights into the practical performance of rPDHG? – particularly regarding the two-stage performance phenomenon, the sensitivity to minor data perturbations, and to possibly improve current heuristic components of practical implementations?

This paper focuses on LP instances with unique optimal solutions (of the primal and dual) and hence unique optimal bases, and proves an accessible iteration bound that is easily computable and indeed has a closed-form expression. We acknowledge this unique optimum assumption is restrictive and is often violated in real-world LP instances due to special problem structures (e.g., network-flow structure can induce degeneracy and multiple optima). Our goal is therefore not to develop a universal accessible iteration bound for all practical LP instances, but rather to provide a fully computable analysis in a simplified setting, and provide new insights into the performance of rPDHG. As the first accessible iteration bound for rPDHG, it may inform future work on accessible bounds beyond the unique-optimum setting.

Although the unique optimum assumption is often violated in real applications, it is a standard simplifying assumption and frequently used to simplify analysis and convey insights. See, e.g., large-scale LP (e.g., Liu et al. [39] and Xiong and Freund [66]) and other optimization problems, such as semidefinite programs (e.g., Alizadeh et al. [1]) and general convex optimization (e.g., Drusvyatskiy and Lewis [19]). The property of unique optima for LP holds almost surely under most models of randomly generated LP instances. The unique optimum assumption is also strictly weaker than the nondegeneracy assumption, because once the primal and dual optimal basic feasible solutions are nondegenerate, then the primal and dual solutions are unique and nondegenerate, which means the LP instance has one unique optimal basis. This nondegeneracy assumption is also used by almost all classic optimization textbooks, such as Bertsimas and Tsitsiklis [8], to simplify analysis and convey insights.

1.1. Outline and contributions In this paper, we consider the following standard form LPs:

$$\min_{x\in\mathbb{R}^n} \quad c^\top x \qquad \text{s.t. } Ax = b\ ,\ x \geq 0\ . \tag{1.1}$$

where $A \in \mathbb{R}^{m\times n}$ is the constraint matrix, $b \in \mathbb{R}^m$ is the right-hand side vector, and $c \in \mathbb{R}^n$ is the objective vector. The corresponding dual problem is:

$$\max_{y\in\mathbb{R}^m, s\in\mathbb{R}^n} \quad b^\top y \qquad \text{s.t. } A^\top y + s = c\ ,\ s \geq 0\ . \tag{1.2}$$

We will assume the optimal basis is unique (formally in Assumption 2.1), denoted by B, and let $x^\star$, $y^\star$ and $s^\star$ denote the optimal solutions. In Section 2, we revisit its saddle-point formulation and its symmetric reformulation on the space of x and s. We also review rPDHG for solving LPs.

Section 3 presents the main result of the paper: an accessible iteration bound of rPDHG that has a closed-form expression. The bound takes the form

$$O\left(\kappa\Phi\cdot\ln\left(\kappa\Phi\frac{\|(x^\star, s^\star)\|_2}{\varepsilon}\right)\right)\ ,$$

where ε is the target tolerance, κ is the standard matrix condition number, and Φ is a geometric condition number of the LP sublevel sets that admits a closed-form expression. This new bound is actually proven equivalent to a bound in Xiong and Freund [67] (under the unique optimum assumption) but has a closed-form expression. Furthermore, Φ has an even simpler upper bound:

$$\Phi \leq \frac{\|x^\star + s^\star\|_1}{\min_{1\leq i\leq n}\left\{x_i^\star + s_i^\star\right\}}\cdot\left\|B^{-1}A\right\|\ .$$

Here and throughout, for matrices $\|\cdot\|$ denotes the spectral norm.

In Section 4, using the established accessible iteration bound, we provide a mathematical analysis of rPDHG's "two-stage performance." Specifically, we show that Stage I achieves finite-time optimal basis identification in

$$O\left(\kappa\Phi\cdot\ln\left(\kappa\Phi\right)\right)$$

iterations, and Stage II exhibits a faster local convergence rate and computes an ε-optimal solution in

$$O\left(\|B^{-1}\|\|A\| \cdot \ln\left(\frac{\min_{1\le i\le n}\{x_i^\star + s_i^\star\}}{\varepsilon}\right)\right)$$

additional iterations. The iteration bound of Stage II is independent of Φ and may thus be significantly lower than that of Stage I. This provides at least a partial explanation for the "two-stage performance" in theory.

In Section 5, using the expression of the new iteration bound, we study the relation between the iteration bound of rPDHG and stability under data perturbations, which is also equivalent to two other types of condition measures: proximity to multiple optima, and the LP sharpness of the instance. Specifically, we show that

$$\Phi = \frac{\|x^\star\|_1 + \|s^\star\|_1}{\min\{\zeta_p, \zeta_d\}},$$

where ζ_p and ζ_d denote the stability measures for the primal and dual problems. This relationship yields a new computational guarantee, and also interprets the impact of tiny data perturbations on the convergence rate of rPDHG.

In Section 6, since the new iteration bounds can now be easily computed, we confirm their tightness via computational tests on LP instances. As predicted by the new iteration bounds, experiments show that tiny perturbations may indeed significantly alter Φ and the overall convergence rates. Additionally, the new iteration bounds are also confirmed matching the practice as our experiments show $\kappa\Phi$ and $\|B^{-1}\|\|A\|$ indeed play important roles in the global linear convergence rates and the two-stage performance. Since our new iteration bounds are proven equivalent to a bound in Xiong and Freund [67], the latter is also confirmed to match the practical behavior by our experiments.

In Section 7, we show a new insight of a practical heuristic of rPDHG can these new bounds enable. We demonstrate that the reweighting that equalizes the ℓ_1-norms of the primal and dual optimal solutions can approximately minimize Φ and the overall iteration bound. This finding provides some intuition concerning the very effective heuristic of balancing the "primal weights" that has been used in many PDHG-based LP solvers.

1.2. Other related works In addition to the previously discussed papers, several other studies have analyzed the performance of PDHG and its variants. Hinder [31] and Lu and Yang [44] present instance-independent worst-case complexity bounds of rPDHG on totally-unimodular LPs and optimal transport problems. Lu and Yang [43] show that the last iterate of the vanilla PDHG without restarts also exhibits a linear convergence rate, dependent on the global Hoffman constant of the KKT system matrix. A recent concurrent work Lu and Yang [45] propose a new restart scheme for PDHG by restarting from the Halpern iterate instead of the average iterate. They prove an accelerated refined complexity bound compared to that of the vanilla PDHG proven in Lu and Yang [47]. This new bound is still based on the Hoffman constant of the reduced KKT system and employs a trajectory-based analysis approach.

Furthermore, after the release of this paper, Xiong [64] shows that rPDHG has a high-probability polynomial iteration bound on LP instances whose data follows certain random distributions, via directly analyzing the accessible iteration bound presented in this paper on the random LP instances.

There has been increasing interest in developing FOMs for LPs. Xiong and Freund [67] propose to use central-path Hessian-based rescaling to accelerate rPDHG, Li et al. [35] design a learning-to-optimize method to emulate PDHG for solving LPs, and Lin et al. [37] propose a PDHG-based solver for addressing general conic LP problems, which include LPs. Beyond PDHG, several other FOMs have been studied recently. Lin et al. [36] and Deng et al. [18] develop ABIP (and ABIP+), an ADMM-based interior-point method that leverages the framework of the homogeneous self-dual interior-point method and employs ADMM to solve the inner log-barrier problems. O'Donoghue et al. [55] and O'Donoghue [54] develop SCS, applying ADMM directly to the homogeneous self-dual formulation for general conic LP problems. Basu et al. [7] utilize accelerated gradient descent to solve a smoothed dual form of LP. Wang et al. [63] use overparametrized neural networks to solve entropically regularized LPs. Very recently, Hough and Vavasis [32] use a Frank-Wolfe method to address the saddle-point problem formulation, and Chen et al. [14] implement a Halpern Peaceman-Rachford method with semi-proximal terms to solve LPs.

1.3. Notation In this paper, we use $[n]$ as shorthand for $\{1, 2, \ldots, n\}$. For a matrix $A \in \mathbb{R}^{m\times n}$, $\mathrm{Null}(A) := \{x \in \mathbb{R}^n : Ax = 0\}$ denotes the null space of A and $\mathrm{Im}(A) := \{Ax : x \in \mathbb{R}^n\}$ denotes the image of A. For any $i \in [m]$ and $j \in [n]$, $A_{\cdot,j}$ and $A_{i,\cdot}$ denote the jth column and ith row of A, respectively. For any subset Θ of $[n]$, A_Θ denotes the submatrix of A formed by the columns indexed by Θ. We use $\|A\|_{\alpha,\beta}$ to denote the induced operator norm, i.e., $\|A\|_{\alpha,\beta} := \sup_{x\neq 0} \frac{\|Ax\|_\beta}{\|x\|_\alpha}$. Specifically, $\|A\|_{2,\infty} = \max_{1\le i\le m} \left\|A_{i,\cdot}\right\|_2$ and $\|A\|_{1,2} = \max_{j\in[n]} \left\|A_{\cdot,j}\right\|_2$. We let $\|\cdot\|_M$ denote the inner product "norm" induced by M, namely, $\|z\|_M := \sqrt{z^\top M z}$. Unless otherwise specified, for a vector v, $\|v\|$ denotes the Euclidean norm, and for a matrix A, $\|A\|$ denotes $\|A\|_{2,2}$, the spectral norm of A. For any set $\mathcal{X} \subset \mathbb{R}^n$, $P_\mathcal{X} : \mathbb{R}^n \to \mathbb{R}^n$ denotes the Euclidean projection onto $\mathcal{X}$, namely, $P_\mathcal{X}(x) := \arg\min_{\hat{x}\in\mathcal{X}} \|x - \hat{x}\|$. For any $x \in \mathbb{R}^n$ and set $\mathcal{X} \subset \mathbb{R}^n$, the Euclidean distance between x and $\mathcal{X}$ is denoted by $\mathrm{Dist}(x, \mathcal{X}) := \min_{\hat{x}\in\mathcal{X}} \|x - \hat{x}\|$ and the M-norm distance between x and $\mathcal{X}$ is denoted by $\mathrm{Dist}_M(x, \mathcal{X}) := \min_{\hat{x}\in\mathcal{X}} \|x - \hat{x}\|_M$. For any set $\mathcal{X} \subseteq \mathbb{R}^n$, $\partial\mathcal{X}$ denotes the boundary of $\mathcal{X}$. For $x \in \mathbb{R}^n$, $x^+ \in \mathbb{R}^n$ denotes the componentwise positive part, i.e., $(x^+)_i = \max\{x_i, 0\}$. For any affine subspace V, we use $\vec{V}$ to denote the associated linear subspace corresponding to V. For any linear subspace $\vec{S}$ in $\mathbb{R}^n$, we use $\vec{S}^\perp$ to denote the corresponding complementary linear subspace of $\vec{S}$. In this paper, we use $O(\cdot)$ to hide factors of only absolute constants.

2. Preliminaries and Background Throughout this paper, we consider the primal problem (1.1) and its dual problem (1.2). For simplicity, we assume the rows of A are linearly independent and the LP is feasible and bounded. We impose the unique-optimum assumption later (Assumption 2.1). A primal-dual solution pair x and (y, s) is optimal if and only if they are feasible and the duality gap is zero, i.e.,

$$\mathrm{Gap}(x, y) := c^\top x - b^\top y = 0 \,. \tag{2.1}$$

Furthermore, (1.1) and (1.2) are equivalent to the following saddle-point problem:

$$\min_{x\in\mathbb{R}^n_+} \max_{y\in\mathbb{R}^m} \quad L(x, y) := c^\top x + b^\top y - (Ax)^\top y \,. \tag{2.2}$$

A pair $(x^\star, y^\star)$ is a saddle point of (2.2) if and only if $x^\star$ is primal optimal and $y^\star$ is dual optimal; in that case the associated dual slack is $s^\star := c - A^\top y^\star$. Conversely, any primal-dual optimal pair $(x^\star, y^\star)$ is a saddle point of (2.2).

2.1. Symmetric formulation of LP Given that (1.2) includes the constraint $A^\top y + s = c$, it follows that for any feasible (y, s), $y = (AA^\top)^{-1}A(c - s)$. Let us define $q := A^\top (AA^\top)^{-1} b$; then the objective function of y is equivalent to an objective function of s, i.e., $b^\top y = q^\top(c - s)$, and (1.2) is thus equivalent to the following (dual) problem on s :

$$\max_{s\in\mathbb{R}^n} \; q^\top(c - s) \quad \text{s.t. } s \in c + \mathrm{Im}(A^\top), s \ge 0. \tag{2.3}$$

We denote the duality gap for the pair (x, s) as $\mathrm{Gap}(x, s) := c^\top x - q^\top(c - s)$, which is equivalent to $\mathrm{Gap}(x, y)$ when $A^\top y + s = c$.

Note that the feasible set of the primal problem (1.1) is the intersection of the affine subspace $V_p := q + \mathrm{Null}(A)$ and the nonnegative orthant. Similarly, the feasible set of (2.3) is the intersection of $V_d := c + \mathrm{Im}\,(A^\top)$ and the nonnegative orthant. We can thus rewrite the primal and dual problems in the following symmetric form:

$$\begin{aligned} &\min_{x\in\mathbb{R}^n} \; c^\top x && \max_{s\in\mathbb{R}^n} \; q^\top(c - s) \\ &\quad \text{s.t. } x \in \mathcal{F}_p := V_p \cap \mathbb{R}^n_+ && \quad \text{s.t. } s \in \mathcal{F}_d := V_d \cap \mathbb{R}^n_+ \end{aligned} \tag{2.4}$$

Let $\vec{V}_p$ and $\vec{V}_d$ denote the linear subspaces associated with the affine subspaces V_p and V_d, respectively. These subspaces are orthogonal complements. We then use $\mathcal{X}^\star$ and $\mathcal{S}^\star$ to denote the optimal solutions for the primal and the dual problem, respectively. Notably, any change in c within the space of $\vec{V}_p$ does not affect $\mathcal{X}^\star$, $\mathcal{S}^\star$, V_p, V_d, $\mathcal{F}_p$, or $\mathcal{F}_d$. Without loss of generality, we may sometimes assume that c is in $\mathrm{Null}(A)$, which can be achieved by replacing c with $P_{\vec{V}_p}(c)$ beforehand. This leads to the following symmetric properties for (1.1) and (2.3).

FACT 2.1. Suppose that $Ac = 0$. Then $\vec{V}_d$ is the orthogonal complement of $\vec{V}_p$, i.e., $\vec{V}_d = \vec{V}_p^{\perp}$. Furthermore, $q^\top c = 0$, and the objective function of (2.3) is equal to $-q^\top s$, and $\text{Gap}(x, s)$ is equal to $c^\top x + q^\top s$. Additionally, $c \in \vec{V}_p$ and $c = \arg\min_{v \in V_d} \|v\|$, and $q \in \vec{V}_d$ and $q = \arg\min_{v \in V_p} \|v\|$.
This re-formulation of the dual was, to the best of our knowledge, first introduced in Todd and Ye [60]. We will use the notation $\mathcal{W}^\star$ to denote the primal-dual pairs, i.e.,

$$\mathcal{W}^\star := \mathcal{X}^\star \times \mathcal{S}^\star = \left\{(x^\star, s^\star) : x^\star \in \mathcal{X}^\star, s^\star \in \mathcal{S}^\star\right\} . \tag{2.5}$$

Our focus is on computing ε-optimal solutions, which are essentially solutions sufficiently close to $\mathcal{W}^\star$, as defined below.

DEFINITION 2.1 (ε-OPTIMAL SOLUTION). A solution w is said to be ε-optimal if the Euclidean distance between w and $\mathcal{W}^\star$ is less than ε, i.e.,

$$\text{Dist}(w, \mathcal{W}^\star) \le \varepsilon .$$

2.2. Restarted Primal-dual hybrid gradient method (rPDHG) The vanilla primal-dual hybrid gradient method (abbreviated as PDHG) was introduced by Esser et al. [20], Pock et al. [57] to solve general convex-concave saddle-point problems, of which (2.2) is a specific subclass. For LP problems, let z denote the primal-dual pair (x, y), and then iteration of PDHG, denoted by $z^{k+1} = (x^{k+1}, y^{k+1}) \leftarrow \text{OnePDHG}(z^k)$, is defined as follows:

$$\begin{cases} x^{k+1} \leftarrow \left(x^k - \tau\left(c - A^\top y^k\right)\right)^+ \\ y^{k+1} \leftarrow y^k + \sigma\left(b - A\left(2x^{k+1} - x^k\right)\right) \end{cases} \tag{2.6}$$

where τ and σ are the primal and dual step-sizes, respectively.

Algorithm 1 presents the general restart scheme for PDHG. We refer to this algorithm as "rPDHG," short for restarted-PDHG.

Algorithm 1: rPDHG: restarted-PDHG

1 **Input:** Initial iterate $z^{0,0} := (x^{0,0}, y^{0,0}) = (0, 0)$, $n \leftarrow 0$, step-size τ, σ, and $\beta \in (0, 1)$;
2 **repeat**
3 **initialize the inner loop:** inner loop counter $k \leftarrow 0$;
4 **repeat**
5 **conduct one step of PDHG:** $z^{n,k+1} \leftarrow \text{OnePDHG}(z^{n,k})$;
6 **compute the average iterate in the inner loop.** $\bar{z}^{n,k+1} \leftarrow \frac{1}{k+1}\sum_{i=1}^{k+1} z^{n,i}$;
7 $k \leftarrow k + 1$;
8 **until** *satisfying the β-restart condition*;
9 **restart the outer loop:** $z^{n+1,0} \leftarrow \bar{z}^{n,k}$, $n \leftarrow n + 1$;
10 **until** *Either $z^{n,0}$ is a saddle point or $z^{n,0}$ satisfies some other convergence condition* ;
11 **Output:** $z^{n,0}$ ($= (x^{n,0}, y^{n,0})$)

Line 5 of Algorithm 1 is an iteration of the vanilla PDHG as described in (2.6). For each iterate $z^{n,k} = (x^{n,k}, y^{n,k})$, we define $s^{n,k} := c - A^\top y^{n,k}$ and $\bar{s}^{n,k} := c - A^\top \bar{y}^{n,k}$. The double superscript indexes the outer iteration counter followed by the inner iteration counter, so that $z^{n,k}$ is the k-th inner iteration of the n-th outer loop. Line 8 of Algorithm 1 specifies an easily verifiable restart condition proposed by Applegate et al. [6] and also used by Xiong and Freund [65, 67] and the practical implementation by Applegate et al. [4]. We will formally specify the β-restart condition and recall its key properties in Section 4.1.

The primary computational effort of Algorithm 1 is the OnePDHG in Line 5, which involves two matrix-vector products. In contrast to traditional methods such as simplex and interior-point methods, rPDHG does not require any matrix factorizations. It is worth noting that the step-sizes τ and σ need to be sufficiently small to ensure convergence. In particular, if $M := \begin{pmatrix} \frac{1}{\tau} I_n & -A^\top \\ -A & \frac{1}{\sigma} I_m \end{pmatrix}$ is positive semi-definite, then Chambolle and

Pock [12] prove rPDHG's iterates will converge to a saddle point of (2.2). The above requirement can be equivalently expressed as:

$$\tau > 0,\ \sigma > 0,\ \text{ and }\ \tau\sigma \le \frac{1}{\|A\|^2}\ . \tag{2.7}$$

Furthermore, the matrix M turns out to be particularly useful in analyzing the convergence of rPDHG through its induced inner product norm defined as $\|z\|_M := \sqrt{z^\top M z}$. This norm will be extensively employed throughout the remainder of this paper.

2.3. LPs with unique optima This paper focuses particularly on LPs with unique optima, the problems satisfying the following assumption:

ASSUMPTION 2.1. *The linear optimization problem* (1.1) *has a unique optimal solution* $x^\star$*, and the dual problem* (1.2) *has a unique optimal solution* $(y^\star, s^\star)$*, i.e.,* $\mathcal{X}^\star = \{x^\star\}$*,* $\mathcal{Y}^\star = \{y^\star\}$ *and* $\mathcal{S}^\star = \{s^\star\}$*.*

When $\mathcal{S}^\star$ is a singleton, actually $\mathcal{Y}^\star$ is a singleton if and only if the rows of the constraint matrix A are linearly independent. This assumption is equivalent to having a unique optimal basis, and is also equivalent to the case that the primal and dual optimal basic feasible solutions are nondegenerate. Randomly generated instances are known to be nondegenerate almost surely (see Borgwardt [9]). The unique optimum assumption is actually weaker than the nondegeneracy assumption, as it requires nondegeneracy only at the optimal solutions. The unique optimum assumption and the stronger nondegeneracy assumption are often used in large-scale linear programming (see, e.g., Liu et al. [39], Xiong and Freund [66]), semidefinite programs (see, e.g., Alizadeh et al. [1]), and general convex optimization (see, e.g., Drusvyatskiy and Lewis [19]). But in practice, due to special structures of real problems, this assumption rarely holds.

Under Assumption 2.1, the primal-dual pair of optimal solutions, $x^\star$ and $(y^\star, s^\star)$, are optimal basic feasible solutions, corresponding to the optimal basis $\Theta := \{i \in [n] : x_i^\star > 0\}$. Let $\bar{\Theta}$ denote the complement of Θ, i.e., $\bar{\Theta} := [n] \setminus \Theta$. Due to strict complementary slackness, $\bar{\Theta} = \{i \in [n] : s_i^\star > 0\}$. As $x^\star$ is an optimal basic feasible solution, there are exactly m components in Θ and $n - m$ components in $\bar{\Theta}$.

Since the algorithm is invariant under permutation of the variables, for simplicity of notation in this paper we assume that the optimal basis is $\{1, 2, \ldots, m\}$ and use B and N to denote the submatrices A_Θ and $A_{\bar{\Theta}}$, respectively. In other words,

$$\Theta = [m] = \{1, 2, \ldots, m\},\ \ \bar{\Theta} = [n] \setminus [m] = \{m+1, m+2, \ldots, n\}\ \text{ and }\ A = \begin{pmatrix} B & N \end{pmatrix}\ . \tag{2.8}$$

With the above Θ and $\bar{\Theta}$, the indices of the nonzero entries of $x^\star$ are exactly $[m]$, and the indices of the nonzero entries of $s^\star$ are exactly $[n] \setminus [m]$.

Later in the paper we will frequently use the following quantities of the matrix A:

$$\lambda_{\max} := \sigma^+_{\max}(A),\ \lambda_{\min} := \sigma^+_{\min}(A),\ \kappa := \frac{\lambda_{\max}}{\lambda_{\min}} \tag{2.9}$$

where $\sigma^+_{\max}(A)$ and $\sigma^+_{\min}(A)$ denote the largest and the smallest nonzero singular values of A, respectively. And κ is often referred to as the matrix condition number of A.

3. Closed-form Complexity Bound of rPDHG This section presents the main result of the paper: an iteration bound of the global linear convergence that has a closed-form expression. First of all, we define the following quantity Φ:

$$\Phi := \left(\|x^\star\|_1 + \|s^\star\|_1\right) \cdot \max\left\{ \max_{1 \le j \le n-m} \frac{\sqrt{\left\|(B^{-1}N)_{\cdot,j}\right\|^2 + 1}}{s^\star_{m+j}},\ \max_{1 \le i \le m} \frac{\sqrt{\left\|(B^{-1}N)_{i,\cdot}\right\|^2 + 1}}{x^\star_i} \right\}\ . \tag{3.1}$$

Notably, it leads to the following iteration bound of rPDHG.

THEOREM 3.1. *Suppose Assumption 2.1 holds and $Ac = 0$. When running Algorithm 1 (rPDHG) with $\tau = \frac{1}{2\kappa}$, $\sigma = \frac{1}{2\lambda_{\max}\lambda_{\min}}$ and $\beta := 1/e$ to solve the LP, the total number of OnePDHG iterations required to compute an ε-optimal solution is at most*

$$O\left(\kappa\Phi \cdot \ln\left(\kappa\Phi \cdot \frac{\|w^\star\|}{\varepsilon}\right)\right) . \tag{3.2}$$

This new computational guarantee for rPDHG is an accessible iteration bound, as it has a closed-form expression that can be easily computed once the optimal solution has been identified. Examining the definition of Φ in (3.1), $B^{-1}A$ is the simplex tableau at the optimal basis B. Overall, Φ is in closed form of the optimal solution/basis, making its computation almost as easy as solving the LP itself. Given the optimal basis, the matrix $B^{-1}N$ can also be easily computed via one matrix factorization followed by one matrix multiplication. Overall, computing Φ is almost as easy as solving the LP itself. In addition, κ can be computed by one singular value decomposition of A. Consequently, the bound (3.2) in Theorem 3.1 is accessible because the main calculation is solving the linear program, computing $B^{-1}N$ and the condition number of A. Except for the very small-scale LP instances, it is usually not way more expensive than solving the LP itself.

Among κ and Φ, κ is a standard definition and easy to compute and analyze. It is solely determined by the matrix A and is independent of the problem's geometry. Conversely, although Φ is defined in terms of the matrix A and the optimal solutions, it is equivalent to an intrinsic measure of the geometry detailed later in Section 3.1. Indeed, Φ is not affected by any parameters of Algorithm 1. In addition, replacing the constraint $Ax = b$ with any preconditioned constraint $DAx = Db$ does not change Φ either.

Furthermore, Φ has the following simplified upper bound:

PROPOSITION 3.1. *The following inequality holds for Φ:*

$$\Phi \le \frac{\|x^\star + s^\star\|_1}{\min_{1\le i\le n}\left\{x_i^\star + s_i^\star\right\}} \cdot \|B^{-1}A\| .$$

Due to strict complementary slackness, all components of $x^\star + s^\star$ are strictly positive, and $\min_{1\le i\le n}\left\{x_i^\star + s_i^\star\right\}$ represents the minimum nonzero entry among $x^\star$ and $s^\star$. This upper bound is the product of two factors: (i) $\frac{\|x^\star+s^\star\|_1}{\min_{1\le i\le n}\left\{x_i^\star+s_i^\star\right\}}$, the ratio between the ℓ_1-norm and the smallest nonzero of the optimal solution, and (ii) $\|B^{-1}A\|$, the spectral norm of $B^{-1}A$. For readers familiar with simplex methods, $B^{-1}A$ is the simplex tableau at the optimal basis B. Its proof directly computes the relaxation of Φ; we defer it to Appendix A.

It should be noted that Φ is also relevant to condition numbers of other methods, beyond its connection to the tableau in simplex methods. Firstly, $\min_{1\le i\le n}\left\{x_i^\star + s_i^\star\right\}$ and the ratio $\frac{\|w^\star\|}{\min_{1\le i\le n}\left\{x_i^\star+s_i^\star\right\}}$ appear in classic complexity analyses of interior-point methods, including the convergence behavior (e.g., Güler and Ye [28]), finite convergence to optimal solutions (e.g., Ye [69]), and identification of the optimal face (e.g, Mehrotra and Ye [51]). Additionally, Lu and Yang [47] demonstrate that PDHG (without restarts) exhibits faster local linear convergence within a neighborhood whose size relates to $\min_{1\le i\le n}\left\{x_i^\star + s_i^\star\right\}$. Furthermore, $\frac{\|w^\star\|}{\min_{1\le i\le n}\left\{x_i^\star+s_i^\star\right\}}$ also appears in finite termination analysis of interior-point methods by Anstreicher et al. [3], Potra [58], in the form that is multiplied by certain norms of $B^{-1}N$. Later in Section 4, we will show that Φ also plays an important role in rPDHG's finite time identification of the optimal basis. Notably, while these condition numbers typically appear inside logarithmic terms in the complexity of interior-point methods, rPDHG's complexity is linear with respect to Φ. This suggests that Φ has more profound implications for the complexity and practical convergence rates of rPDHG compared to interior-point methods. Beyond interior-point methods, the upper bound $\frac{\|x^\star\|_1}{\min_{1\le i\le m} x_i^\star}$ also plays a crucial role in the complexity analysis of simplex and policy-iteration methods for discounted Markov decision problems (see Ye [71]).

The rest of this section presents the proof of Theorem 3.1. Section 3.1 recalls the sublevel set condition numbers defined by Xiong and Freund [67] and their roles in rPDHG. Furthermore, Section 3.1 shows a key lemma of the equivalence relationship between Φ and the sublevel set condition numbers, which helps prove Theorem 3.1. After that, Section 3.2 proves the key lemma of the equivalence relationship.

3.1. Sublevel-set condition numbers and the proof of Theorem 3.1 Recall that $\mathcal{F}_p = V_p \cap \mathbb{R}^n_+$ and $\mathcal{F}_d = V_d \cap \mathbb{R}^n_+$ denote the feasible sets of the primal and dual problems, respectively. Let $\mathcal{F} := \mathcal{F}_p \times \mathcal{F}_d$ represent the primal-dual feasible set of the solution pair (x, s). The optimal solution can then be characterized as

$$\mathcal{W}^\star := \mathcal{F} \cap \{w = (x, s) \in \mathbb{R}^{2n} : \mathrm{Gap}(w) = 0\} ,$$

the set of feasible solutions with zero duality gap. For $w = (x, s)$, we write $\mathrm{Gap}(w) := \mathrm{Gap}(x, s)$. The δ-sublevel set is similarly characterized as the feasible solutions whose duality gap is less than or equal to δ, formally defined as follows:

DEFINITION 3.1 (δ-SUBLEVEL SET $\mathcal{W}_\delta$). For $\delta \geq 0$, the δ-sublevel set $\mathcal{W}_\delta$ is defined as:

$$\mathcal{W}_\delta := \mathcal{F} \cap \{w = (x, s) \in \mathbb{R}^{2n} : \mathrm{Gap}(w) \leq \delta\} . \tag{3.3}$$

Based on $\mathcal{W}_\delta$, Xiong and Freund [67] introduce the following two geometric condition numbers: the diameter D_δ and the conic radius r_δ.

DEFINITION 3.2 (CONDITION NUMBERS OF $\mathcal{W}_\delta$). For $\delta \geq 0$, the diameter of $\mathcal{W}_\delta$ is defined as

$$D_\delta := \max_{u,v \in \mathcal{W}_\delta} \|u - v\| . \tag{3.4}$$

And the conic radius of $\mathcal{W}_\delta$ is the optimal objective value of the optimization problem

$$r_\delta := \left(\max_{w \in \mathcal{W}_\delta, r \geq 0} r \qquad \text{s.t. } \left\{\hat{w} : \|\hat{w} - w\| \leq r\right\} \subseteq \mathbb{R}^{2n}_+ \right) , \tag{3.5}$$

which is also equal to $\left(\max_{w \in \mathcal{W}_\delta} \mathrm{Dist}(w, \partial\mathbb{R}^{2n}_+)\right)$ and $\left(\max_{w \in \mathcal{W}_\delta} \min_{i \in [2n]} w_i\right)$. (Indeed, for any $w \in \mathbb{R}^{2n}_+$, the largest Euclidean ball centered at w contained in $\mathbb{R}^{2n}_+$ has radius $\min_{i \in [2n]} w_i$, which equals $\mathrm{Dist}(w, \partial\mathbb{R}^{2n}_+)$.)

These condition numbers play a crucial role in the iteration bound of rPDHG as follows:

LEMMA 3.1. *Suppose Assumption 2.1 holds and $Ac = 0$. When running Algorithm 1 (rPDHG) with $\tau = \frac{1}{2\kappa}$, $\sigma = \frac{1}{2\lambda_{\max}\lambda_{\min}}$ and $\beta := 1/e$ to solve the LP, the total number of OnePDHG iterations T required to obtain the first outer iteration N such that $w^{N,0} = (x^{N,0}, c - A^\top y^{N,0})$ is ε-optimal is bounded above by*

$$T \leq 380\kappa \left(\liminf_{\delta \searrow 0} \frac{D_\delta}{r_\delta} \right) \left[\ln\left(380\kappa \left(\liminf_{\delta \searrow 0} \frac{D_\delta}{r_\delta} \right) \right) + \ln\left(\frac{\|w^\star\|}{\varepsilon} \right) \right] . \tag{3.6}$$

This lemma is a simple extension of Theorem 3.5 of Xiong and Freund [67], and its complete proof is deferred to Appendix A. For simplicity of notation, the rest of the paper uses $\hat{\Phi}$ to denote $\left(\liminf_{\delta \searrow 0} \frac{D_\delta}{r_\delta}\right)$:

$$\hat{\Phi} := \liminf_{\delta \searrow 0} \frac{D_\delta}{r_\delta} . \tag{3.7}$$

Lemma 3.1 implies that the linear convergence rate is mostly determined by κ, a condition number of the constraint matrix, and $\hat{\Phi}$, a condition number of the sublevel set. When $\mathcal{W}^\star$ is a singleton, the sublevel set $\mathcal{W}_\delta$ is always inside the tangent cone to $\mathcal{F}$ at $w^\star$. Although looking similar, $\hat{\Phi}$ is not equivalent to the "width" of the tangent cone at $w^\star$ (see a formal definition in Freund and Vera [23]). The latter is an inherent property of the tangent cone, but the former is also influenced by the direction of (c, q).

Actually, we have the following critical lemma, which states the approximate equivalence between Φ and $\hat{\Phi}$:

LEMMA 3.2. *Suppose that Assumption 2.1 holds, and let Φ and $\hat{\Phi}$ be defined in* (3.1) *and* (3.7)*, respectively. The geometric condition number $\hat{\Phi}$ and Φ are equivalent up to a constant factor of* 2*, i.e.,*

$$\Phi \leq \hat{\Phi} \leq 2\Phi . \tag{3.8}$$

Because of Lemma 3.2, all previous discussions for Φ also apply to $\hat{\Phi}$. In addition, we can now directly prove Theorem 3.1.

Proof of Theorem 3.1. Directly applying Lemma 3.2 in the iteration bound of Lemma 3.1 yields the desired iteration bound (3.2). □

Furthermore, since Φ is equivalent to $\hat{\Phi}$ up to a constant, the new accessible iteration bound of Theorem 3.1 is also equivalent to the iteration bound of Lemma 3.1 (an iteration bound of Xiong and Freund [67]) up to a constant. Xiong and Freund [67] point out that proper central path based Hessian rescaling can improve $\hat{\Phi}$ to at most $2n$, so Lemma 3.2 indicates that this rescaling can also improve Φ to at most $2n$.

We now prove Lemma 3.2 in Section 3.2.

3.2. Proof of the approximate equivalence between Φ and $\hat{\Phi}$ The sublevel set can be equivalently regarded as the "primal sublevel set" for an artificial LP problem whose variables contain both the primal and dual variables. Section 3.2.1 extends the definitions of the sublevel set to the primal space only and demonstrates how to compute its condition numbers approximately. Subsequently, Section 3.2.2 illustrates how to approximate the sublevel-set condition numbers by treating them as primal sublevel-set condition numbers of an artificial problem, thereby proving Lemma 3.2.

3.2.1. Condition numbers of the primal sublevel set Recall that the primal feasible set is $\mathcal{F}_p := V_p \cap \mathbb{R}^n_+$, the intersection of the nonnegative orthant $\mathbb{R}^n_+$ and the affine subspace of the linear equality constraints. We define the objective error of x as $\mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x) := c^\top x - f^\star$, where $f^\star$ is the optimal objective $c^\top x^\star$ for an optimal $x^\star$ of (1.1). The optimal primal solution is the feasible solution with zero objective error. The primal δ-sublevel set is then defined as:

$$\mathcal{X}_\delta := \mathcal{F}_p \cap \left\{ x \in \mathbb{R}^n : \mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x) := c^\top x - f^\star \le \delta \right\} , \tag{3.9}$$

the sets of feasible primal solutions whose objective error does not exceed δ. Analogous to Definition 3.2, we define the diameter and conic radius of $\mathcal{X}_\delta$ as

$$D^p_\delta := \max_{u,v \in \mathcal{X}_\delta} \|u - v\| \quad \text{and} \quad r^p_\delta := \max_{x \in \mathcal{X}_\delta} \mathrm{Dist}(x, \partial \mathbb{R}^n_+) . \tag{3.10}$$

Now we show the representation of $\mathcal{X}_\delta$ and how to compute D^p_δ and r^p_δ.

Convex hull representation of the primal sublevel set $\mathcal{X}_\delta$. Under Assumption 2.1, the optimal primal and dual solutions are unique and nondegenerate, corresponding to a unique optimal basis. Each edge emanating from $\mathcal{X}^\star = \{x^\star\}$ connects to a basic feasible solution of an adjacent basis. There are exactly $n - m$ entering basic variables. Let the corresponding directions of these edges be given by the vectors $u^1, u^2, \ldots, u^{n-m} \in \mathbb{R}^n$. Since $\Theta = [m]$ and $\bar{\Theta} = [n] \setminus [m]$, these vectors can be computed as follows:

$$u^j_{[m]} := -B^{-1} N_{\cdot,j}, \;\; u^j_{m+j} := 1, \;\; \text{and} \;\; u^j_k := 0 \text{ for all } k \notin [m] \text{ and } k \neq m + j \tag{3.11}$$

for each $j \in [n - m]$. Therefore, the $n - m$ edges are as follows:

$$\mathcal{E}^j := \left\{ x^\star + \theta \cdot u^j : \theta \ge 0, \; x^\star + \theta \cdot u^j \in \mathcal{F}_p \right\} \quad \text{for each } j \in [n - m] . \tag{3.12}$$

If $u^j \ge 0$, then $\mathcal{E}^j$ is an extreme ray. Otherwise, $\mathcal{E}^j$ connects to an adjacent basic feasible solution of $x^\star$. Based on these edges, when δ is sufficiently small so that it is no larger than the extreme points' best nonzero objective error $\bar{\delta}_p$ (which is always strictly positive for LP) defined by

$$\bar{\delta}_p := \begin{cases} \min\{\mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x) : x \in \mathrm{EP}_{\mathcal{F}_p} \setminus \mathcal{X}^\star\} & \text{if } \mathrm{EP}_{\mathcal{F}_p} \setminus \mathcal{X}^\star \neq \emptyset \\ +\infty & \text{if } \mathrm{EP}_{\mathcal{F}_p} \setminus \mathcal{X}^\star = \emptyset \end{cases} \tag{3.13}$$

then $\mathcal{X}_\delta$ can be represented as the convex hull of these edges. Here we use $\mathrm{EP}_{\mathcal{F}_p}$ to denote the set of extreme points of $\mathcal{F}_p$.

LEMMA 3.3. *Suppose that Assumption 2.1 holds and $\delta \in (0,\infty)$ lies in $(0,\bar{\delta}_p]$. Then $\mathcal{X}_\delta$ is represented by the following convex hull formulation:*

$$\mathcal{X}_\delta = \mathrm{Conv}\left(\{x^\star\} \cup \{x^j : j \in [n-m]\}\right) \tag{3.14}$$

where

$$x^j := x^\star + \frac{\delta}{s^\star_{m+j}} \cdot u^j \text{ for each } j \in [n-m] \ . \tag{3.15}$$

Proof. Because $\delta \in (0,\bar{\delta}_p]$, the halfspace $\{x : c^\top x < f^\star + \bar{\delta}\}$ contains no other basic feasible solutions of $\mathcal{F}_p$ and intersects no other edges except the $n-m$ edges emanating from $\mathcal{X}^\star$, namely $\mathcal{E}^1, \mathcal{E}^2, \ldots, \mathcal{E}^{n-m}$. Therefore, in addition to $x^\star$, the other extreme points of $\mathcal{X}_\delta$ are the intersection points of the hyperplane $\{x : c^\top x = f^\star + \delta\}$ and the edges $\mathcal{E}^1, \mathcal{E}^2, \ldots, \mathcal{E}^{n-m}$. These intersection points all exist because $\delta \in (0,\bar{\delta}_p] \cap (0,\infty)$. Moreover, they are precisely the $\{x^j : j \in [n-m]\}$ defined in (3.15), as the objective errors $\mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x^j)$ all equal δ (which is because $\mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x^j) = (x^j)^\top s^\star = x^j_{m+j} s^\star_{m+j} = \frac{\delta}{s^\star_{m+j}} \cdot u^j_{m+j} \cdot s^\star_{m+j} = \delta u^j_{m+j} = \delta$). Therefore, $\mathcal{X}_\delta$ is indeed the convex hull of the $n-m+1$ points in $\{x^\star\}$ and $\{x^j : j \in [n-m]\}$. □

Computing the diameter and conic radius of $\mathcal{X}_\delta$. We now provide an approximation of D^p_δ for sufficiently small δ.

LEMMA 3.4. *Suppose that Assumption 2.1 holds and $\delta \in (0,\bar{\delta}_p]$. Then we have*

$$\delta \cdot \max_{1\le j\le n-m} \frac{\sqrt{\left\|B^{-1}N_{\cdot,j}\right\|^2 + 1}}{s^\star_{m+j}} \le D^p_\delta \le 2\delta \cdot \max_{1\le j\le n-m} \frac{\sqrt{\left\|B^{-1}N_{\cdot,j}\right\|^2 + 1}}{s^\star_{m+j}} \ . \tag{3.16}$$

Proof. The diameter of a polyhedron is the maximum distance between any two extreme points of the polyhedron. Thus, $D^p_\delta \le \max_{i,j\in[n-m]} \|x^i - x^\star\| + \|x^j - x^\star\| \le 2 \cdot \max_{1\le j\le n-m} \|x^j - x^\star\| = 2\delta \cdot \max_{1\le j\le n-m} \frac{\|u^j\|}{s^\star_{m+j}}$, where the last equality uses Lemma 3.3. Conversely, $D^p_\delta \ge \max_{1\le j\le n-m} \|x^j - x^\star\|$, which equals $\delta \cdot \max_{1\le j\le n-m} \frac{\|u^j\|}{s^\star_{m+j}}$ (by Lemma 3.3). Finally, note from (3.11) that $\|u^j\| = \sqrt{\left\|B^{-1}N_{\cdot,j}\right\|^2 + 1}$ so the proof is completed. □

Next, we show how to exactly compute r_δ when δ is sufficiently small. Specifically, we study δ small enough so that for $\{x^j : 1 \le j \le n-m\}$ defined in (3.15):

$$\min_{1\le k\le m} x^j_k \ge x^j_{m+j} \text{ for all } j \in [n-m] \ . \tag{3.17}$$

In other words, δ is small enough such that x^j_{m+j} is one of the smallest nonzeros of x^j for all $j \in [n-m]$. Using (3.15) and (3.11), the inequalities in (3.17) are equivalent to $x^\star_k + \frac{\delta}{s^\star_{m+j}} u^j_k \ge \frac{\delta}{s^\star_{m+j}}$ for all $j \in [n-m]$ and $k \in [m]$. If $u^j_k \ge 1$ this inequality holds for every $\delta \ge 0$, while if $u^j_k < 1$ it is further equivalent to $0 \le \delta \le \frac{x^\star_k s^\star_{m+j}}{1-u^j_k}$. We therefore introduce

$$\hat{\delta}_p := \min\left\{ \frac{x^\star_k s^\star_{m+j}}{1-u^j_k} \ : \ 1 \le j \le n-m,\ 1 \le k \le m, \text{ and } u^j_k < 1 \right\} \tag{3.18}$$

where the minimum is interpreted as $+\infty$ if the set is empty. For every $0 < \delta \le \hat{\delta}_p$, condition (3.17) holds. Assumption 2.1 implies $x^\star_k > 0$ and $s^\star_{m+j} > 0$, and thus $\hat{\delta}_p > 0$.

LEMMA 3.5. *Suppose that Assumption 2.1 holds. For any $\delta > 0$ and $\delta \le \min\{\bar{\delta}_p, \hat{\delta}_p\}$, we have*

$$r^p_\delta = \frac{\delta}{\|s^\star\|_1} \ . \tag{3.19}$$

Proof. By definition, $r_\delta^p = \max_{x\in\mathcal{X}_\delta} \operatorname{Dist}(x, \partial\mathbb{R}^n_+) = \max_{x\in\mathcal{X}_\delta} \min_{1\le l\le n} x_l$. Since $\delta \le \bar{\delta}_p$, using the convex hull formulation of $\mathcal{X}_\delta$ presented in Lemma 3.3, r_δ^p can be equivalently written as:

$$r_\delta^p = \max_{\substack{\lambda\in\mathbb{R}^{n-m+1}_+ \\ \sum_{j=1}^{n-m+1}\lambda_j=1}} \min_{1\le l\le n} x(\lambda)_l \quad \text{in which} \quad x(\lambda) := \lambda_{n-m+1}\cdot x^\star + \sum_{j=1}^{n-m} \lambda_j \cdot x^j \ . \tag{3.20}$$

Due to the above definition of $x(\lambda)$ and the definition of x^j in (3.15), for each $m+j$ in $\bar{\Theta} = [n]\setminus[m]$, the component $x(\lambda)_{m+j}$ is given by $\lambda_j \cdot x^j_{m+j}$.

We now claim that a smallest component of $x(\lambda)$ is of an index in $[n]\setminus[m]$. On the one hand,

$$\min_{1\le i\le m} x(\lambda)_i \overset{(3.20)}{\ge} \lambda_{n-m+1}\cdot \min_{1\le i\le m} x^\star_i + \sum_{j=1}^{n-m} \lambda_j \cdot \min_{1\le i\le m} x^j_i \ge 0 + \sum_{j=1}^{n-m} \lambda_j \cdot x^j_{m+j} \ . \tag{3.21}$$

where the last inequality uses (3.17) because $\delta < \hat{\delta}_p$. On the other hand, because $x(\lambda)_{m+j} = \lambda_j \cdot x^j_{m+j}$,

$$\sum_{j=1}^{n-m} \lambda_j \cdot x^j_{m+j} = \sum_{j=1}^{n-m} x(\lambda)_{m+j} \ge \min_{1\le j\le n-m} x(\lambda)_{m+j} \ . \tag{3.22}$$

Overall, for this small δ, (3.21) and (3.22) ensure that a smallest component of $x(\lambda)$ is of an index in $[n]\setminus[m]$.

Consequently, when computing r_δ^p we only need to consider the components in $[n]\setminus[m]$.

$$r_\delta^p \overset{(3.20)}{=} \max_{\substack{\lambda\in\mathbb{R}^{n-m+1}_+ \\ \sum_{j=1}^{n-m+1}\lambda_j=1}} \min_{1\le l\le n} x(\lambda)_l = \max_{\substack{\lambda\in\mathbb{R}^{n-m+1}_+ \\ \sum_{j=1}^{n-m+1}\lambda_j=1}} \min_{m+1\le l\le n} x(\lambda)_l = \max_{\substack{\lambda\in\mathbb{R}^{n-m+1}_+ \\ \lambda_{n-m+1}=0 \\ \sum_{j=1}^{n-m}\lambda_j=1}} \min_{m+1\le l\le n} x(\lambda)_l \tag{3.23}$$

where the last equality follows from $x^\star_{[n]\setminus[m]} = 0$, which implies that λ_{n-m+1} in an optimal λ for r_δ^p must be 0. The value of r_δ^p in (3.23) is then equal to the optimal objective of

$$\begin{pmatrix} \max\limits_{\lambda\in\mathbb{R}^{n-m}} \min\limits_{1\le j\le n-m} \lambda_j \cdot x^j_{m+j} \\ \text{s.t.} \quad \sum\limits_{1\le j\le n-m} \lambda_j = 1,\ \lambda \ge 0 \end{pmatrix} = \begin{pmatrix} \max\limits_{\lambda\in\mathbb{R}^{n-m}} \min\limits_{1\le j\le n-m} \lambda_j \cdot \dfrac{\delta}{s^\star_{m+j}} \\ \text{s.t.} \quad \sum\limits_{1\le j\le n-m} \lambda_j = 1,\ \lambda \ge 0 \end{pmatrix} \tag{3.24}$$

where the equality uses $x^j_{m+j} = \frac{\delta}{s_{m+j}}$ according to (3.15) and (3.11). Finally, the optimal solution $\lambda^\star$ of (3.24) is given by $\lambda^\star_j = \frac{s^\star_{m+j}}{\sum_{k=1}^{n-m} s^\star_{m+k}} = \frac{s^\star_{m+j}}{\|s^\star\|_1}$ for each j, and the optimal objective is equal to $\frac{\delta}{\|s^\star\|_1}$. This establishes (3.19) and completes the proof. □

It is noteworthy that if we similarly define the dual sublevel set $\mathcal{S}_\delta$ and then Freund [22, Theorem 3.2.] shows $\delta \le r_\delta^p \cdot \max_{s\in\mathcal{S}_\delta} \|s\|_1 \le 2\delta$. A direct application of this result yields $\frac{1}{\|s^\star\|_1} \le \lim_{\delta\to 0} \frac{r_\delta^p}{\delta} \le \frac{2}{\|s^\star\|_1}$. Lemma 3.5 provides a slightly stronger result by precisely computing r_δ^p.

3.2.2. Proof of Lemma 3.2 In this subsection, we prove Lemma 3.2. We begin by demonstrating that the sublevel set $\mathcal{W}_\delta$ can be regarded as a primal sublevel set of an artificial problem. Using the results of Section 3.2.1, we then show how to approximate D_δ and compute r_δ, which subsequently allows us to approximate $\hat{\Phi}$ and prove Lemma 3.2.

Problem (2.3) can be transformed into the subsequent standard-form problem

$$\max_{s\in\mathbb{R}^n} \ q^\top(c-s) \quad \text{s.t. } Qs = Qc,\ s \ge 0 \tag{3.25}$$

for any $Q \in \mathbb{R}^{(n-m)\times n}$ whose rows are linearly independent and orthogonal to the rows of A so that $\text{Null}(Q) = \text{Im}(A^\top)$. This equivalence holds because $\text{Im}(A^\top) + c$ in (2.3) is identical to $\{s : Qs = Qc\}$ in (3.25). For problem (3.25), the optimal basis is $\bar{\Theta} = [n] \setminus [m]$, the complement of Θ. Although multiple choices of Q exist, the matrix $Q_{\bar{\Theta}}^{-1} Q_\Theta$ is always equal to $-(B^{-1}N)^\top$.

LEMMA 3.6. *Suppose that Assumption 2.1 holds. The matrix $Q_{\bar{\Theta}}^{-1} Q_\Theta$ is equal to $-(B^{-1}N)^\top$.*

Proof. Given that $\text{Null}(Q) = \text{Im}(A^\top)$, we have $QA^\top = 0$, i.e.,

$$0 = QA^\top = \begin{pmatrix} Q_\Theta & Q_{\bar{\Theta}} \end{pmatrix} \begin{pmatrix} A_\Theta^\top \\ A_{\bar{\Theta}}^\top \end{pmatrix} = Q_\Theta A_\Theta^\top + Q_{\bar{\Theta}} A_{\bar{\Theta}}^\top = Q_\Theta B^\top + Q_{\bar{\Theta}} N^\top .$$

Since the optimal bases B and $Q_{\bar{\Theta}}$ are of full rank, we obtain $Q_{\bar{\Theta}}^{-1} Q_\Theta = -N^\top (B^\top)^{-1} = -(B^{-1}N)^\top$. □

Overall, the primal problem (1.1) and the dual problem (3.25) can be combined and reformulated as an equivalent standard-form LP problem in the product space of x and s:

$$\min_{w=(x,s)\in\mathbb{R}^{2n}} \begin{pmatrix} c \\ q \end{pmatrix}^\top w \quad \text{s.t.} \begin{pmatrix} A & 0 \\ 0 & Q \end{pmatrix} w = \begin{pmatrix} b \\ Qc \end{pmatrix},\ w \geq 0 . \tag{3.26}$$

The above (3.26) is also in standard form, and satisfies Assumption 2.1. Furthermore, the duality gap $\text{Gap}(x, s)$ of any (x, s) is the same as the objective error for (3.26) because

$$\text{Gap}(x, s) = c^\top x - q^\top (c - s) = \begin{pmatrix} c \\ q \end{pmatrix}^\top w - q^\top c = \begin{pmatrix} c \\ q \end{pmatrix}^\top (w - w^\star) \tag{3.27}$$

where the last equality follows from $0 = \text{Gap}(w^\star) = c^\top x^\star - q^\top (c - s^\star) = \begin{pmatrix} c \\ q \end{pmatrix}^\top w^\star - q^\top c$. The right-hand side of (3.27) is the objective error $\text{E}_{\text{obj}}(w)$ of w, defined by $\text{E}_{\text{obj}}(w) := \begin{pmatrix} c \\ q \end{pmatrix}^\top (w - w^\star)$. Consequently, (3.27) implies that the δ-sublevel set $\mathcal{W}_\delta$ is identical to the primal δ-sublevel set of (3.26). Therefore, utilizing the results in Section 3.2.1, we can directly approximately compute D_δ and r_δ, by treating them as the condition numbers of the primal sublevel set of (3.26).

LEMMA 3.7. *Suppose that Assumption 2.1 holds. There exists a positive $\bar{\delta}$ such that for any $0 \leq \delta \leq \bar{\delta}$, it holds that*

$$\hat{D}_\delta \leq D_\delta \leq 2\hat{D}_\delta \quad \text{and} \quad r_\delta = \frac{\delta}{\|x^\star\|_1 + \|s^\star\|_1} , \tag{3.28}$$

where

$$\hat{D}_\delta := \delta \cdot \max \left\{ \max_{1\leq j\leq n-m} \frac{\sqrt{\left\|(B^{-1}N)_{\cdot,j}\right\|^2 + 1}}{s^\star_{m+j}}, \max_{1\leq i\leq m} \frac{\sqrt{\left\|(B^{-1}N)_{i,\cdot}\right\|^2 + 1}}{x^\star_i} \right\} . \tag{3.29}$$

Proof. Let H denote the constraint matrix of (3.26) for simplicity of notation. We use Ω to represent the indices of the optimal basis of (3.25), which is $\Theta \cup (n + \bar{\Theta})$. Here $n + \bar{\Theta}$ denotes $\{n + j : j \in \bar{\Theta}\}$. Similarly, the complementary set is $\bar{\Omega} = \bar{\Theta} \cup (n + \Theta)$. We use $\Omega(i)$ to denote the i-th smallest index component of Ω and use $\bar{\Omega}(j)$ to denote the j-th smallest index component of $\bar{\Omega}$. Both Ω and $\bar{\Omega}$ contain exactly n components, and the optimal basis H_Ω is given by $\begin{pmatrix} A_\Theta & 0 \\ 0 & Q_{\bar{\Theta}} \end{pmatrix}$, or equivalently $\begin{pmatrix} B & 0 \\ 0 & Q_{\bar{\Theta}} \end{pmatrix}$. By an approach similar to that used in Section 2.1 for deriving the dual problem, the dual problem of (3.26) is symmetric with (3.26) and also has a unique optimal solution $\tilde{w}^\star$ equal to $(s^\star, x^\star)$.

We now prove the first half of (3.28) using Lemma 3.4. The term $\sqrt{\left\|B^{-1}N_{\cdot,j}\right\|^2 + 1}$ in (3.16) is $\sqrt{\left\|H_\Omega^{-1} H_{\cdot,\bar{\Omega}(j)}\right\|^2 + 1}$ for problem (3.26). And $s^\star_{\bar{\Theta}}$ in (3.16) is $\tilde{w}_{\bar{\Omega}}$ in problem (3.26). Therefore, Lemma 3.4 implies

$$\bar{D}_\delta \leq D_\delta \leq 2\bar{D}_\delta, \quad \text{where } \bar{D}_\delta := \delta \cdot \max_{j\in[n]} \frac{\sqrt{\left\|H_\Omega^{-1} H_{\cdot,\bar{\Omega}(j)}\right\|^2 + 1}}{\tilde{w}_{\bar{\Omega}(j)}} . \tag{3.30}$$

To compute the value of $\bar{D}_\delta$, we consider two cases based on the structure of $\bar{\Omega} = \bar{\Theta} \cup (n + \Theta)$. When $\bar{\Omega}(j) \in \bar{\Theta}$, we have $\tilde{w}^\star_{\bar{\Omega}(j)} = s^\star_{\bar{\Omega}(j)}$, and $H_\Omega^{-1} H_{\cdot,\bar{\Omega}(j)} = H_\Omega^{-1} \begin{pmatrix} A_{\cdot,\bar{\Omega}(j)} \\ 0 \end{pmatrix} = \begin{pmatrix} B^{-1} A_{\cdot,\bar{\Omega}(j)} \\ 0 \end{pmatrix}$. When $\bar{\Omega}(j) \in n + \Theta$, we have $\tilde{w}^\star_{\bar{\Omega}(j)} = x^\star_{\bar{\Omega}(j)-n}$, and $H_\Omega^{-1} H_{\cdot,\bar{\Omega}(j)} = H_\Omega^{-1} \begin{pmatrix} 0 \\ Q_{\cdot,\bar{\Omega}(j)-n} \end{pmatrix} = \begin{pmatrix} 0 \\ Q_\Theta^{-1} Q_{\cdot,\bar{\Omega}(j)-n} \end{pmatrix} = \begin{pmatrix} 0 \\ -\left((B^{-1}N)^\top\right)_{\cdot,\bar{\Omega}(j)-n} \end{pmatrix}$, where the last inequality uses Lemma 3.6. Therefore,

$$\bar{D}_\delta = \delta \cdot \max \left\{ \max_{\bar{\Omega}(j) \in \bar{\Theta}} \frac{\sqrt{\left\| \begin{pmatrix} B^{-1} A_{\cdot,\bar{\Omega}(j)} \\ 0 \end{pmatrix} \right\|^2 + 1}}{s^\star_{\bar{\Omega}(j)}}, \max_{\bar{\Omega}(j) \in n+\Theta} \frac{\sqrt{\left\| \begin{pmatrix} 0 \\ -\left((B^{-1}N)^\top\right)_{\cdot,\bar{\Omega}(j)-n} \end{pmatrix} \right\|^2 + 1}}{x^\star_{\bar{\Omega}(j)-n}} \right\}$$

$$\overset{(2.8)}{=} \delta \cdot \max \left\{ \max_{1 \le j \le n-m} \frac{\sqrt{\left\| (B^{-1}N)_{\cdot,j} \right\|^2 + 1}}{s^\star_{m+j}}, \max_{1 \le i \le m} \frac{\sqrt{\left\| (B^{-1}N)_{i,\cdot} \right\|^2 + 1}}{x^\star_i} \right\} = \hat{D}_\delta \ .$$

Finally, substituting $\bar{D}_\delta = \hat{D}_\delta$ back to (3.30) proves the first half of (3.28).

As for the second half of (3.28), note that $\tilde{w}^\star = (s^\star, x^\star)$, so by Lemma 3.5, when δ is sufficiently small we have $r_\delta = \frac{\delta}{\|\tilde{w}^\star\|_1} = \frac{\delta}{\|x^\star\|_1 + \|s^\star\|_1}$. □

Finally, we are ready to prove Lemma 3.2.

Proof of Lemma 3.2. Lemma 3.7 establishes that $\hat{D}_\delta \le D_\delta \le 2\hat{D}_\delta$ for sufficiently small δ, so we can deduce $\lim_{\delta \searrow 0} \frac{\hat{D}_\delta}{r_\delta} \le \lim_{\delta \searrow 0} \frac{D_\delta}{r_\delta} \le 2 \cdot \lim_{\delta \searrow 0} \frac{\hat{D}_\delta}{r_\delta}$. Note that $\hat{\Phi} = \lim_{\delta \searrow 0} \frac{D_\delta}{r_\delta}$ as defined in (3.7). Substituting the values of $\hat{D}_\delta$ and r_δ in Lemma 3.7 into (3.1) yields $\Phi = \lim_{\delta \searrow 0} \frac{\hat{D}_\delta}{r_\delta}$. Therefore, $\Phi \le \hat{\Phi} \le 2\Phi$. □

4. Finite-Time Optimal Basis Identification and Fast Local Convergence In this section, we investigate the two-stage performance of rPDHG. It is frequently observed in practice that the behavior of rPDHG transitions to faster local linear convergence in a neighborhood of the optimal solution, in which the support sets of all iterates keep consistent. Lu and Yang [47] study this phenomenon for vanilla PDHG. In the first stage, PDHG converges in a sublinear rate until identifying the active set of the converging solution. In the second stage, PDHG turns to faster local linear convergence. Recently, Lu and Yang [45] propose the restarted Halpern PDHG (rHPDHG), a variant of rPDHG, and prove its two-stage performance behavior. However, the iteration bounds of the two stages proven above (for both PDHG and rHPDHG) are not accessible because they both depend on the Hoffman constant of a linear system that is determined by the converging solution. The Hoffman constant is hard to analyze, challenging to compute, and may be too conservative. And computing the trajectory of the algorithm may also be difficult.

Based on the new iteration bounds obtained in Section 3 for LPs with unique optima, this section will show an accessible refined convergence guarantee of rPDHG that avoids Hoffman constants entirely. Although using the additional assumption of unique optimum, the new iteration bounds for the two stages both have closed-form expressions and are thus straightforward to analyze and compute once an optimal solution/basis is available:

- In Stage I, rPDHG identifies the optimal basis within at most $O(\kappa\Phi \cdot \ln(\kappa\Phi))$ iterations.
- In Stage II, having identified the optimal basis Θ, the behavior of rPDHG transitions to faster local linear convergence that is no longer related to Φ. In this stage, components of index in Θ are sufficiently bounded away from 0, while all other components equal 0.

The following theorem summarizes the iteration bounds of the two stages:

THEOREM 4.1. *Suppose Assumption 2.1 holds and $Ac = 0$. Let Algorithm 1 (rPDHG) run with $\tau = \frac{1}{2\kappa}$, $\sigma = \frac{1}{2\lambda_{\max}\lambda_{\min}}$ and $\beta := 1/e$ to solve the LP. Let T_1 be the total number of OnePDHG iterations required to obtain N_1 such that for all $N \ge N_1$ the positive components of $x^{N,0}$ exactly correspond to the optimal basis. Then,*

$$T_1 \le T_{basis} \ , \quad \textit{in which} \quad T_{basis} := O\left(\kappa\Phi \cdot \ln\left(\kappa\Phi\right)\right). \tag{4.1}$$

Furthermore, let T_2 be the total number of OnePDHG iterations required to obtain the first N_2 for which $w^{N_2,0}$ is ε-optimal. Then,

$$T_2 \le T_{basis} + T_{local}, \quad \text{in which} \quad T_{local} := O\left(\|B^{-1}\|\|A\| \cdot \max\left\{0,\ \ln\left(\frac{\min_{1\le i\le n}\left\{x_i^\star + s_i^\star\right\}}{\varepsilon}\right)\right\}\right). \tag{4.2}$$

Theorem 4.1 demonstrates that it takes at most T_{basis} iterations for rPDHG to identify the optimal basis (independent of ε), after which it requires at most additional T_{local} iterations to achieve ε-optimality. The above iteration bounds are not aimed at improving upon the iteration bound in Theorem 3.1, the main value of this result is that both T_{basis} and T_{local} are accessible, and they do not contain any Hoffman constant and are straightforward to compute and analyze if the optimal solution is known. Furthermore, T_{local} is independent of Φ, which is frequently considerably larger. This provides a partial explanation for why rPDHG often becomes significantly faster in Stage II compared to Stage I.

The first-stage iteration bound T_{basis} exhibits a linear relationship with $\frac{\|w^\star\|_1}{\xi}$ (in the expression of Φ), where we use ξ to denote the smallest nonzero of $x^\star$ and $s^\star$, written as follows:

$$\xi := \min_{1\le i\le n}\left\{x_i^\star + s_i^\star\right\}\ . \tag{4.3}$$

The first-stage iteration bounds of the vanilla PDHG (without restarts) and the restarted Halpern PDHG proven by Lu and Yang [45, 47] also depend on $\frac{\|w^\star\|_1}{\xi}$ (in a different norm), but they are further multiplied by a Hoffman constant of a linear system. An empirical comparison with this characterization will be presented in Section 6. A similar dependence on $\frac{\|w^\star\|}{\xi}$ and certain norms of $B^{-1}A$ is also observed in the complexity analysis of finite-termination results for IPMs, such as Anstreicher et al. [3], Potra [58]. In those cases of IPMs, this dependence appears only within logarithmic terms, but it comes with higher per-iteration cost.

Once the optimal basis is identified, the method's behavior automatically transitions into the second stage without changing the algorithm or additional "crossover" operations. Notably, the coefficient $\|B^{-1}\|\|A\|$ is solely determined by the optimal basis and the constraint matrix, which could be upper bounded by a constant that is only determined by the constraint matrix. A similar complexity bound that only depends on the constraint matrix also holds for IPMs, proven by Vavasis and Ye [61]. Note that a smaller ξ may slightly decrease T_{local} but simultaneously significantly increase T_{basis} because T_{basis} is linear in $\frac{\|w^\star\|_1}{\xi}$.

The significance of $\kappa\Phi$ and $\|B^{-1}\|\|A\|$ in the two-stage performance will be empirically confirmed in Section 6. The remainder of this section will show the proof sketch of Theorem 4.1 and some important lemmas.

Proof sketch of Theorem 4.1 and outline of the rest of this section. In the rest of this section, we outline some key steps in the proof of Theorem 4.1. The proof proceeds in three steps.

(i) Global linear convergence with adaptive restart condition. First of all, we recall the proof idea of the global linear convergence of rPDHG in Section 4.1. Under a certain "sharpness" condition with a parameter $\mathcal{L}$ (formally defined in (4.6)) and the β-restart condition, a certain stationarity measure (the normalized duality gap) is forced to decrease by a fixed factor at every outer iteration, within at most $O(\mathcal{L}/\beta)$ OnePDHG steps. This implies that rPDHG converges linearly to any arbitrary "ε-close" neighborhood of the optimal solution within $\tilde{O}(\mathcal{L}\cdot\log(1/\varepsilon))$ OnePDHG steps, where the $\tilde{O}$ hides absolute constants and logarithmic terms of $\mathcal{L}$ and $\|w^\star\|$.

(ii) Local linear sharpness and local linear convergence. While the global linear convergence relies on a sharpness condition with a constant $\mathcal{L}$ that holds globally, the sharpness condition also holds with a potentially much smaller local constant $\mathcal{L}_{\text{loc}}$ in a neighborhood of $z^\star$, which leads to potentially faster local linear convergence. In Section 4.3, Lemma 4.1 shows that once the optimal basis has been identified and the primal variables are bounded sufficiently away from zero, a local sharpness condition holds with a local constant $\mathcal{L}_{\text{loc}} = O(\|B^{-1}\|\|A\|)$. Plugging this $\mathcal{L}_{\text{loc}}$ into the same generic rPDHG convergence theorem yields the Stage-II iteration bound T_{local} in (4.2).

(iii) Finite-time basis identification and neighborhood entrance. In Section 4.3, Lemma 4.2 shows that the condition in Step (ii) is satisfied once an iterate enters a sufficiently small neighborhood of $z^\star$. Because of

the unique optimal basis assumption, if an iterate is close enough to the optimal solution, complementary slackness forces the primal iterate of the next iteration to have the same support set as $x^\star$ after one OnePDHG step. Furthermore, due to the "nonexpansiveness" property of PDHG, once an outer iterate enters this neighborhood, all subsequent outer iterates and averaged inner iterates inherit the condition in Step (ii).

Combining these three steps gives the two-stage bound claimed in Theorem 4.1. The above three-step proof scheme is standard in the literature of the two-stage convergence, see e.g., Lu and Yang [45], while we prove the accessible bounds in each step, which avoid any Hoffman constant. The proof of the lemmas in the following subsection and the complete proof of Theorem 4.1 are all deferred to Appendix B.

4.1. Global linear convergence with adaptive restart condition This subsection recalls the global linear convergence of rPDHG with the β-restart condition. It is built on a metric of stationarity, called the "normalized duality gap," proposed by Applegate et al. [6].

Definition 4.1 (Normalized duality gap, Applegate et al. [6]). For any $z = (x, y) \in \mathbb{R}^{m+n}$ and a certain $r > 0$, the normalized duality gap of the saddle-point problem (2.2) is then defined as

$$\rho(r; z) := \frac{1}{r} \sup_{\hat{z} \in B(r;z)} \left[L(x, \hat{y}) - L(\hat{x}, y) \right] \tag{4.4}$$

in which $B(r; z) := \left\{ \hat{z} := (\hat{x}, \hat{y}) : \hat{x} \in \mathbb{R}^n_+ \text{ and } \|\hat{z} - z\|_M \le r \right\}$.
The normalized duality gap $\rho(r; z)$ is also a valid measure of the optimality and feasibility errors of z, and it can be efficiently computed or approximated in strongly polynomial time as shown by Applegate et al. [6]. Furthermore, let the k-th iterate of PDHG be denoted as z^k, and let the average of the first k iterates be $\bar{z}^k := \frac{1}{k} \sum_{i=1}^k z^i$. The normalized duality gap $\rho(\|z^0 - \bar{z}^k\|_M; \bar{z}^k)$ of the average iterate $\bar{z}^k$ converges to 0 sublinearly. See Applegate et al. [6], Xiong and Freund [65] for more details.

Then the β-restart condition holds if and only if $n = 0$ and $k = 1$, or

$$\rho(\|\bar{z}^{n,k} - z^{n,0}\|_M; \bar{z}^{n,k}) \le \beta \cdot \rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0}) \,, \tag{4.5}$$

for a chosen value of $\beta \in (0, 1)$. In words, the restart triggers once the normalized duality gap at $\bar{z}^{n,k}$, evaluated with radius $\|\bar{z}^{n,k} - z^{n,0}\|_M$, is at most a factor β times the normalized duality gap at $z^{n,0}$, evaluated with radius $\|z^{n,0} - z^{n-1,0}\|_M$.

For LP problems, it can be shown that there exists a constant $\mathcal{L} \ge 0$ such that the following condition holds for all $n \ge 1$:

$$\mathrm{Dist}_M(z^{n,0}, \mathcal{Z}^\star) \le \rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0}) \cdot \mathcal{L} \,. \tag{4.6}$$

This condition indicates the M-distance to the optimal solutions is upper bounded by the normalized duality gap multiplied by the fixed constant $\mathcal{L}$. Applegate et al. [6] refer to this as the "sharpness" property of the normalized duality gap. Under this condition, each inner loop requires at most $\left\lceil \frac{8\mathcal{L}}{\beta} \right\rceil$ iterations to achieve sufficient decrease in the normalized duality gap of Applegate et al. [6, Theorem 2] (see Appendix A for a formal statement and proof). Therefore, n outer loops contain at most $O\left(\frac{n\mathcal{L}}{\beta}\right)$ OnePDHG iterations, while reducing the normalized duality gap to a β^n fraction of its initial value. This leads to a linear convergence rate dependent on $\mathcal{L}$. Indeed, Lemma 3.13 of Xiong and Freund [67] demonstrates that (4.6) always holds with $\mathcal{L} = O(\kappa\hat{\Phi})$, which then leads to the global convergence rate in Lemma 3.1.

These results are the cornerstone of Step (i) of the proof of Theorem 4.1: the β-restarts drive both the normalized duality gap and the M-distance to the solution set down at a global linear rate governed by $O(\mathcal{L})$.

4.2. Local linear sharpness and local linear convergence The linear convergence of rPDHG relies on the condition (4.6). If (4.6) holds for all $n \ge 1$ with a constant $\mathcal{L} \ge 0$, then the global linear convergence is established by showing the number of iterations required for each inner loop does not exceed $O(\frac{\mathcal{L}}{\beta})$. We will now demonstrate the existence of a close neighborhood of the optimal solution $z^\star$ within which (4.6) holds with a potentially much smaller $\mathcal{L}_{\text{loc}}$ than the global $\mathcal{L}$, so rPDHG requires a smaller number of iterations for the inner loops and therefore have a faster local linear convergence rate. We demonstrate that the condition

(4.6) indeed holds with an alternative $\mathcal{L}_{\text{loc}}$ for the iterates $\bar{z} = (\bar{x}, \bar{y})$ if (i) each component of $\bar{x}_\Theta$ is sufficiently bounded away from 0 and (ii) each component of $\bar{x}_{\bar{\Theta}}$ equals zero. For simplicity, we assume that the optimal basis is $\{1, 2, \ldots, m\}$. We will use the following step-size dependent constant:

$$c_{\tau,\sigma} := \max\left\{\frac{1}{\sqrt{\sigma}\lambda_{\min}},\ \frac{1}{\sqrt{\tau}}\right\} . \tag{4.7}$$

When $\tau = \frac{1}{2\kappa}$ and $\sigma = \frac{1}{2\lambda_{\max}\lambda_{\min}}$ (the step-sizes used by Theorem 4.1), we have $c_{\tau,\sigma} = \sqrt{2\kappa}$.

LEMMA 4.1. *Under Assumption 2.1, for any $\bar{z} = (\bar{x}, \bar{y})$ and $r > 0$ such that*

$$(i)\ \ \bar{x}_i \geq r\sqrt{\tau} \text{ for } i \in [m], \quad \text{and} \quad (ii)\ \bar{x}_{m+j} = 0 \text{ for } j \in [n-m] , \tag{4.8}$$

it holds that

$$\|\bar{z} - z^\star\|_M \leq \sqrt{2} c_{\tau,\sigma} \|B^{-1}\| \sqrt{\frac{1}{\sigma} + \frac{\|A\|^2}{\tau}} \cdot \rho(r; \bar{z}) . \tag{4.9}$$

When the step-sizes are carefully chosen, (4.9) in Lemma 4.1 can be further simplified:

REMARK 4.1. With the choice of step-size $\tau = \frac{1}{2\kappa}$ and $\sigma = \frac{1}{2\lambda_{\max}\lambda_{\min}}$, (4.9) becomes

$$\|\bar{z} - z^\star\|_M \leq 4\|B^{-1}\|\|A\| \cdot \rho(r; \bar{z}) . \tag{4.10}$$

The above results are the foundation of Step (ii) of the proof sketch for Theorem 4.1: there exists a neighborhood of $z^\star$ within which the sharpness condition (4.6) holds with $\mathcal{L}_{\text{loc}} = O(\|B^{-1}\|\|A\|)$, leading to potentially faster local linear convergence.

4.3. Finite-time basis identification and neighborhood entrance We now show that the two conditions in Lemma 4.1 are automatically satisfied by all average iterations $\bar{z}^{N,k} = (\bar{x}^{N,k}, \bar{y}^{N,k}) = \frac{1}{k}\sum_{i=1}^{k} z^{N,i}$ once a previous outer loop iteration $z^{n,0}$ (with $n \leq N$) is sufficiently close to the optimal solution $z^\star$ under the M-norm distance. Lemma 4.2 below formalizes Step (iii) of the proof of Theorem 4.1. It quantifies how close an outer iterate $z^{n,0}$ needs to be to $z^\star$ in the M-norm in order for the subsequent averaged iterates $\bar{z}^{N,k}$ to identify the optimal basis and enter the neighborhood where Lemma 4.1 applies.

LEMMA 4.2. *Under Assumption 2.1, suppose that Algorithm 1 (rPDHG) has any β-restart condition, and let the step-sizes satisfy* (2.7) *strictly. If there exists an outer-loop index t such that*

$$\left\|z^{t,0} - z^\star\right\|_M \leq \bar{\varepsilon} := \frac{\sqrt{1 - \sqrt{\tau\sigma}\|A\|}}{3} \cdot \min\left\{\frac{1}{\sqrt{\tau}}, \sqrt{\tau}\right\} \cdot \left(\min_{1 \leq i \leq n} \left\{x_i^\star + s_i^\star\right\}\right) , \tag{4.11}$$

then for any $N \geq t$ and $k \geq 1$, the following conditions hold:

$$(i)\ \ \bar{x}_i^{N,k} \geq \sqrt{\tau}\left\|\bar{z}^{N,k} - z^{N,0}\right\|_M \ \text{ for } i \in [m] , \quad \text{and} \quad (ii)\ \bar{x}_{m+j}^{N,k} = 0 \text{ for } j \in [n-m] . \tag{4.12}$$

Moreover, the positive components of $\bar{x}^{N,k}$ correspond exactly to the optimal basis.

With the above lemma, directly using Lemma 3.1 yields the finite-time basis identification and neighborhood entrance. This lemma is the key of Step (iii) of the proof of Theorem 4.1.

Overall, Lemmas 4.1 and 4.2 provide the foundation of the proof of Theorem 4.1. In Stage I, rPDHG converges to a neighborhood of the optimal solution such that condition (4.11) of Lemma 4.2 is satisfied. The number of iterations in this stage is determined by the linear convergence rate established in Lemma 3.1. In Stage II, rPDHG converges to the optimal solution with accelerated local linear convergence, owing to the potentially smaller $\mathcal{L}$ provided by Lemma 4.1. The complete proofs of Lemmas 4.1 and 4.2 and Theorem 4.1 are all deferred to Appendix B.

5. Relationship of Φ with Stability under Data Perturbations, Proximity to Multiple Optima, and LP Sharpness The previous sections established new accessible iteration bounds for rPDHG in terms of the geometric quantity Φ, which admits a closed-form expression. In this section, we relate Φ to a measure of stability under data perturbations, and we use this relationship to derive a computational guarantee stated directly in terms of these quantities. This measure of stability is also equivalent to two other condition measures (i) proximity to multiple optima, and (ii) the LP sharpness of the instance. These interpretations help explain rPDHG's sensitivity to small data perturbations, and reveal connections to other works; for example, Lu and Yang [45, 47] use a parameter of "near-degeneracy" to bound the duration of the Stage I of PDHG, and Xiong and Freund [65, 66] provide computational guarantees of rPDHG using the LP sharpness.

We start by defining two quantities of primal and dual stabilities. For the original problem (1.1), let the perturbed problem be as follows:

$$\min\ \tilde{c}^\top x \quad \text{s.t. } Ax = \tilde{b}\ ,\ x \geq 0 \tag{5.1}$$

where $\tilde{c}$ and $\tilde{b}$ might be the perturbed versions of c and b respectively. When Assumption 2.1 holds for (1.1), we define ζ_p and ζ_d as follows:

$$\zeta_p := \inf\left\{\|\Delta c\| :\ \Theta \text{ is not the unique optimal basis for (5.1) with } \tilde{c} = c + \Delta c \text{ and } \tilde{b} = b\right\}\ , \tag{5.2}$$

$$\zeta_d := \inf\left\{\|\Delta b\|_{(AA^\top)^{-1}} :\ \Theta \text{ is not the unique optimal basis for (5.1) with } \tilde{c} = c \text{ and } \tilde{b} = b + \Delta b\right\}\ . \tag{5.3}$$

The ζ_p and ζ_d denote the size of the smallest perturbation on the cost vector c and the right-hand side vector b, respectively, such that the optimal basis becomes different. In other words, the larger they are, the more stable the optimal basis is under data perturbations on b and c. Here ζ_d uses the $(AA^\top)^{-1}$-norm instead of the Euclidean norm because later we will show that ζ_d can be defined in the symmetric way to ζ_p on the symmetric form (2.4).

More importantly, Φ has a close relationship with ζ_p and ζ_d, leading to a new computational guarantee using ζ_p and ζ_d. Below is the result of this section.

THEOREM 5.1. *Suppose Assumption 2.1 holds. The following relationship holds for Φ, ζ_p, and ζ_d:*

$$\Phi = \frac{\|x^\star\|_1 + \|s^\star\|_1}{\min\left\{\zeta_p,\ \zeta_d\right\}}\ . \tag{5.4}$$

Therefore, in the identical setting of Theorem 3.1, the total number of OnePDHG iterations required to compute an ε-optimal solution is at most

$$O\left(\kappa \cdot \frac{\|x^\star\|_1 + \|s^\star\|_1}{\min\left\{\zeta_p,\ \zeta_d\right\}} \cdot \ln\left(\frac{\kappa \frac{\|x^\star\|_1 + \|s^\star\|_1}{\min\left\{\zeta_p,\ \zeta_d\right\}} \cdot \|w^\star\|}{\varepsilon}\right)\right).$$

This theorem implies that the less stable the optimal basis is under data perturbations, the larger the value of Φ, and the more iterations rPDHG might require to compute an ε-optimal solution. Actually in Section 6 we will confirm the tightness of these bounds via experiments on LP instances. Small values of $\min\left\{\zeta_p,\ \zeta_d\right\}$ may very significantly affect the performance of rPDHG, because they stay in the denominator in the expression of Φ in (5.4). It should be noted that ζ_p and ζ_d are not intrinsic properties of the constraint matrix as they are also dependent on c and b, so ζ_p and ζ_d do not affect the Stage II iteration bound T_{local} in Theorem 4.1.

In the remainder of this section, Section 5.1 introduces the equivalence relationship between the measure of stability under data perturbations and two other condition measure, proximity to multiple optima and LP sharpness. We then highlight the implications of these relationships for rPDHG's sensitivity to small perturbations and discuss the connection with other work. Finally, Section 5.3 presents the proof of Theorem 5.1.

5.1. Stability under data perturbations, proximity to multiple optima, LP sharpness Both the primal and dual problems in the symmetric form (2.4) are instances of the following generic form of LP:

$$\min\ g^\top u \qquad \text{s.t. } u \in \mathcal{F}_{generic} := V_{generic} \cap \mathbb{R}^n_+ \tag{5.5}$$

where the feasible set $\mathcal{F}_{generic}$ is the intersection of the nonnegative orthant $\mathbb{R}^n_+$ and an affine subspace $V_{generic}$. The objective function $g^\top u$ is a linear function. We denote the optimal solution of (5.5) by $\mathcal{U}^\star$, in which $u^\star$ is an optimal solution. We let $\mathrm{OPT}(\check{g})$ denote the set of optimal solutions of the generic LP (5.5) with the objective vector equal to $\check{g}$. For example, $\mathrm{OPT}(g) = \mathcal{U}^\star$. The primal problem (1.1) is an instantiation of (5.5) with $\mathcal{F}_{generic} = \mathcal{F}_p$, $V_{generic} = V_p$, and $g = c$. Similarly, the dual problem (2.3) is another instantiation of (5.5) with $\mathcal{F}_{generic} = \mathcal{F}_d$, $V_{generic} = V_d$, and $g = q$. With this symmetric form, the stability under data perturbation ζ can be defined as follows.

Definition 5.1 (Stability under data perturbations). The stability under data perturbations is defined as

$$\zeta := \inf_{\Delta g} \left\{ \|\Delta g\| : \mathrm{OPT}(g + \Delta g) \neq \emptyset \text{ and } \mathrm{OPT}(g + \Delta g) \not\subseteq \mathrm{OPT}(g) \right\} . \tag{5.6}$$

The ζ for (1.1) and (2.3) are essentially ζ_p and ζ_d, respectively, as stated in the following remark.

Remark 5.1. Suppose Assumption 2.1 holds. The ζ for (1.1) and (2.3) is equal to ζ_p and ζ_d, respectively. The proof of Remark 5.1 is straightforward and thus deferred to Section 5.2.

Furthermore, we introduce another two measures. The first one is the proximity to multiple optima, defined as the size of the smallest perturbation that leads to multiple optimal solutions.

Definition 5.2 (Proximity to multiple optima). When $\mathrm{OPT}(g)$ is a singleton, the proximity to multiple optima is defined as

$$\eta = \min_{\Delta g} \left\{ \|\Delta g\| : |\,\mathrm{OPT}(g + \Delta g)| > 1 \right\} . \tag{5.7}$$

Since having multiple optima is equivalent to having degenerate dual optimal solutions, η can also be interpreted as the proximity to degenerate dual optima.

The second measure is LP sharpness (see Xiong and Freund [65]), which measures how quickly the objective function grows away from the optimal solution set $\mathcal{U}^\star$ (i.e., $\mathrm{OPT}(g)$) among all feasible points.

Definition 5.3 (LP sharpness). The LP sharpness of (5.5) is defined as

$$\mu := \inf_{u \in \mathcal{F}_{generic} \setminus \mathcal{U}^\star} \frac{\mathrm{Dist}(u, V_{generic} \cap \{u \in \mathbb{R}^n : g^\top u = g^\top u^\star\})}{\mathrm{Dist}(u,\ \mathcal{U}^\star)} . \tag{5.8}$$

Sharpness is a useful analytical tool (for example, see Lu and Yang [43], Yang and Lin [68]), and Applegate et al. [6] employ the sharpness of the normalized duality gap for the saddle-point problem to prove the linear convergence of rPDHG on LP. Sharpness for LP, denoted by LP sharpness, is a more natural and intuitive measure of the original LP instance.

When $\mathcal{U}^\star$ is a singleton, the above three measures are equivalent in the following sense.

Proposition 5.1. *When $\mathcal{U}^\star$ is a singleton, the following relationship holds:*

$$\frac{1}{\left\|P_{\vec{V}_{generic}}(g)\right\|} \cdot \zeta = \frac{1}{\left\|P_{\vec{V}_{generic}}(g)\right\|} \cdot \eta = \mu . \tag{5.9}$$

Here ζ is normalized by the norm of $P_{\vec{V}_{generic}}(g)$, indicating that ζ and η are equivalent to μ in a relative sense. This normalization arises because μ is a purely geometric concept that remains invariant under positive scaling of g. We use the norm of $P_{\vec{V}_{generic}}(g)$ rather than g because the complementary part $P_{\vec{V}^\perp_{generic}}(g)$ does not have any impact on the optimal solution and the smallest perturbation Δg must lie in $\vec{V}_{generic}$. The equivalence between μ and ζ is already proven by Xiong and Freund [65] so it only suffices to prove the equivalence between ζ and η, which is not hard to show. We defer the complete proof of Proposition 5.1 to

Section 5.2. We use μ_p and μ_d to denote the μ for the primal problem (1.1) and the dual problem (2.3), for which $\left\|P_{\vec{V}_p}(c)\right\|$ and $\|q\|$ correspond to $\left\|P_{\vec{V}_{generic}}(g)\right\|$ (Fact 2.1).

Interpreting the sensitivity of rPDHG to perturbations and connections to other work. Theorem 5.1 and the above equivalence relationships of ζ_p, ζ_d with the other two condition measures also provide new insights into the performance of rPDHG. Here we detail them as follows.

It is often observed that rPDHG has good performance in some LP instances with multiple optimal solutions, but a minor data perturbation results in a substantial degradation in rPDHG's performance on the perturbed problem (see Section 6 for examples). Now we have some insight and a partial explanation for this phenomenon. According to Proposition 5.1 and Remark 5.1, ζ_p and ζ_d are equal to the proximity to multiple optima for (1.1) and (2.3). The perturbed problem still stays in close proximity to the original unperturbed problem, so its $\min\{\zeta_p, \zeta_d\}$ is at most as large as the magnitude of the perturbation. Furthermore, $\min\{\zeta_p, \zeta_d\}$ lies in the denominator of the iteration bound of Theorems 5.1 (and Stage I iteration bound of Theorem 4.1), so a small perturbation may significantly increase the iteration bound. This explanation is complementary with Lu and Yang [47] which study the vanilla PDHG and explain the phenomenon by the size of the region for local fast convergence. Our result shows that perturbations significantly affect the iteration bound for Stage I of rPDHG as well. This effect of perturbations on rPDHG will be confirmed by computational experiments in Section 6.

Moreover, Theorem 5.1 together with Remark 5.1 and Proposition 5.1 also provide new computational guarantees that use κ, the size of the optimal solutions, and the LP sharpness μ_p and μ_d. They are simpler and more intuitive guarantees than those in Xiong and Freund [65], because the latter also involves the limiting error ratio. This simplification is due to the unique optimum assumption. Notably, our iteration bounds are strictly better than the iteration bounds in Xiong and Freund [66, Corollary 4.3], which also use this additional assumption but exhibit quadratic dependence on the reciprocals of μ_p and μ_d.

5.2. Proof of Remark 5.1 and Proposition 5.1

We first prove Remark 5.1:

Proof of Remark 5.1. For ζ of (1.1), $\text{OPT}(c+\Delta c) \nsubseteq \text{OPT}(c)$ if and only if Θ is no longer an optimal basis or the optimal basis is no longer unique for the perturbed problem. This proves ζ_p is equal to ζ of (1.1).

Similarly, ζ of (2.3) is defined as the Euclidean norm of the smallest perturbation Δq on the objective function so that $\bar{\Theta}$ is no longer the unique optimal basis for (2.3). Since all feasible solutions are in the affine subspace $c + \text{Im}(A^\top)$, the smallest perturbation Δq must lie in $\text{Im}(A^\top)$ and there exists $\Delta b_0 \in \mathbb{R}^m$ such that $\Delta q = A^\top(AA^\top)^{-1}\Delta b_0$. When applying the right-hand side perturbation Δb_0 on (1.1), the corresponding dual problem (2.3) has the objective vector $A^\top(AA^\top)^{-1}(b+\Delta b_0)$ that is exactly equal to $q + \Delta q$. In this case $\|\Delta q\| = \|A^\top(AA^\top)^{-1}\Delta b\| = \|\Delta b\|_{(AA^\top)^{-1}}$. Therefore, the smallest Euclidean norm of the objective vector perturbation Δq on (2.3) such that $\bar{\Theta}$ is no longer the unique optimal basis is equal to the smallest $(AA^\top)^{-1}$-norm of right-hand side perturbation Δb on (1.1) such that Θ is no longer the unique optimal basis. This proves ζ_d is equal to ζ of (2.3). □

We then prove Proposition 5.1. Before that, we show a useful result that will be frequently used later.

LEMMA 5.1 **(Theorem 5.1 of Xiong and Freund [65])**. *LP sharpness is equivalent to stability under data perturbations through the relation:* $\mu = \zeta \cdot \frac{1}{\left\|P_{\vec{V}_{generic}}(g)\right\|}$.

With Lemma 5.1 we can prove Proposition 5.1.

Proof of Proposition 5.1. First of all, according to Lemma 5.1, it suffices to prove first equality in (5.9). In the case that $\text{OPT}(g)$ (i.e., $\mathcal{U}^\star = \{u^\star\}$) is a singleton, ζ is the smallest magnitude of the perturbation that leads to multiple optimal solutions at the threshold at which a new solution is added to $\text{OPT}(g+\Delta g)$ while $u^\star$ remains optimal. Therefore, ζ is equal to η. This finishes the proof. □

5.3. Proof of Theorem 5.1

The key to proving Theorem 5.1 is the following lemma:

LEMMA 5.2. *Suppose that Assumption 2.1 holds. The ζ_p and ζ_d have the following expression:*

$$\zeta_p = \min_{1\le j\le n-m} \frac{s^\star_{m+j}}{\sqrt{\left\|(B^{-1}N)_{\cdot,j}\right\|^2+1}} \text{ and } \zeta_d = \min_{1\le i\le m} \frac{x^\star_i}{\sqrt{\left\|(B^{-1}N)_{i,\cdot}\right\|^2+1}}. \tag{5.10}$$

Before proving it, we recall how to compute the LP sharpness μ of the generic LP (5.5) by computing the smallest sharpness along all of the edges emanating from the optimal solutions.

LEMMA 5.3 (**A restatement of Theorem 5.5 of Xiong and Freund [65]**). *Suppose that Assumption 2.1 holds. Let $\mathcal{U}^\star = \{u^\star\}$ and the directions of the edges emanating from $\mathcal{U}^\star$ be $v^1, v^2, \ldots, v^{n-m}$. Then for any given $\bar{\varepsilon} > 0$, the LP sharpness μ is characterized as follows:*

$$\mu = \min_{1 \le j \le n-m} \frac{\mathrm{Dist}(u^\star + \bar{\varepsilon} \cdot v^j, V_{generic} \cap \{u \in \mathbb{R}^n : g^\top u = g^\top u^\star\})}{\|\left(u^\star + \bar{\varepsilon} \cdot v^j\right) - u^\star\|} \tag{5.11}$$

Now we are ready to prove Lemma 5.2.

Proof of Lemma 5.2. Let μ_p and μ_d denote the LP sharpness for (1.1) and (2.3) respectively. We first compute μ_p and then $\zeta_p = \left\|P_{\vec{V}_p}(c)\right\| \mu_p$ (Lemma 5.1). For (1.1), $u^\star = x^\star$, $V_{generic} = V_p$ and $\{u \in \mathbb{R}^n : g^\top u = g^\top u^\star\} = \left\{x : \mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x) = 0\right\}$ in the generic LP (5.5), with the directions of connected edges given by $\{u^j : 1 \le j \le n-m\}$ as defined in (3.11). Furthermore, the following equalities hold:

$$\begin{aligned} &\mathrm{Dist}\left(x^\star + \bar{\varepsilon} \cdot u^j, V_p \cap \left\{x : \mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x) = 0\right\}\right) = \frac{\mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x^\star + \bar{\varepsilon} \cdot u^j)}{\left\|P_{\vec{V}_p}(c)\right\|} \\ &= \frac{(s^\star)^\top (x^\star + \bar{\varepsilon} \cdot u^j) - (s^\star)^\top x^\star}{\left\|P_{\vec{V}_p}(c)\right\|} = \bar{\varepsilon} \cdot \frac{(s^\star)^\top u^j}{\left\|P_{\vec{V}_p}(c)\right\|} = \bar{\varepsilon} \cdot \frac{s^\star_{m+j} u^j_{m+j}}{\left\|P_{\vec{V}_p}(c)\right\|} = \bar{\varepsilon} \cdot \frac{s^\star_{m+j}}{\left\|P_{\vec{V}_p}(c)\right\|} \end{aligned} \tag{5.12}$$

for all $j \in [n-m]$. Here the second and third equalities use the result that $\mathrm{E}^{\mathrm{p}}_{\mathrm{obj}}(x) = \mathrm{Gap}(x, s^\star) = (s^\star)^\top x - (s^\star)^\top x^\star$ for any $x \in V_p$. The fourth equality holds because $s^\star_{[m]} = 0$ and $x^j_{[n]\setminus([m]\cup\{m+j\})} = 0$. The final equality uses $u^j_{m+j} = 1$. In addition, for all $j \in [n-m]$ we have

$$\|(x^\star + \bar{\varepsilon} \cdot u^j) - x^\star\| = \bar{\varepsilon} \cdot \|u^j\| = \bar{\varepsilon} \cdot \sqrt{\left\|(B^{-1}N)_{\cdot,j}\right\|^2 + 1}\,. \tag{5.13}$$

Substituting (5.12) and (5.13) into the definition (5.11) of μ_p yields the expression:

$$\mu_p = \frac{1}{\left\|P_{\vec{V}_p}(c)\right\|} \cdot \min_{1 \le j \le n-m} \frac{s^\star_{m+j}}{\sqrt{\left\|(B^{-1}N)_{\cdot,j}\right\|^2 + 1}}\,.$$

Finally, substituting it back to $\zeta_p = \left\|P_{\vec{V}_p}(c)\right\| \mu_p$ proves the first half of (5.10).

Next we compute μ_d and then $\zeta_d = \left\|P_{\vec{V}_d}(q)\right\| \mu_d$ (Lemma 5.1). For (2.3), we repeat the above process on (3.25), and then we can symmetrically obtain $\mu_d = \frac{1}{\left\|P_{\vec{V}_d}(q)\right\|} \cdot \min_{1 \le i \le m} \frac{x^\star_i}{\sqrt{\left\|(Q_{\bar{\Theta}}^{-1} Q_\Theta)_{\cdot,i}\right\|^2 + 1}}$. By Lemma 3.6, $Q_{\bar{\Theta}}^{-1} Q_\Theta = -(B^{-1}N)^\top$, and thus $\left\|(Q_{\bar{\Theta}}^{-1} Q_\Theta)_{\cdot,i}\right\| = \|(B^{-1}N)_{i,\cdot}\|$. Combined with $\zeta_d = \left\|P_{\vec{V}_d}(q)\right\| \mu_d$, this completes the proof. □

With Lemma 5.2, we can now prove Theorem 5.1.

Proof of Theorem 5.1. Directly substituting the expressions of ζ_p and ζ_d into the expression (3.1) of Φ completes the proof. □

6. Experimental Confirmation This section shows how our new theories align with the practical performance through computational evaluations. Section 6.1 confirms the reciprocal relationship between rPDHG complexity and the magnitude of perturbations, providing experimental evidence for our findings

in Section 5. Section 6.2 confirms the significance of $\kappa\Phi$ and $\|B^{-1}\|\|A\|$ in the two-stage performance of rPDHG, supporting our results in Section 4.

We implement rPDHG (Algorithm 1) on standard-form LPs, adhering precisely to the setting of Theorem 3.1, with one exception: the normalized duality gap uses the $\tilde{M}$-norm that is defined by $\|(x,y)\|_{\tilde{M}} := \sqrt{\frac{1}{\tau}\|x\|^2 + \frac{1}{\sigma}\|y\|^2}$ where $\tilde{M} := \begin{pmatrix} \frac{1}{\tau} I_n & \\ & \frac{1}{\sigma} I_m \end{pmatrix}$. This alternative is proven by Applegate et al. [6], Xiong and Freund [67] equivalent to the original normalized duality gap but significantly more computationally efficient. It is also widely adopted in practice, such as Applegate et al. [4], Lu and Yang [46**?**].

6.1. Effects of Data Perturbations To evaluate rPDHG's sensitivity to perturbations in the objective vector c, we construct a family of standard-form LP instances (1.1) with data $\left(A^1, b^1, c^1\right)$, where

$$A^1 = [1,1,1]\,,\ b^1 = 2,\ \text{ and } c^1 = c^1_\gamma := [2,-1,-1] + \left[0, -\frac{\gamma}{2}, \frac{\gamma}{2}\right]$$

for parameter $\gamma \geq 0$. This LP family, denoted by LP^1_γ, is designed to illustrate the effect of perturbations $[0, -\frac{\gamma}{2}, \frac{\gamma}{2}]$ on the objective vector c^1_0. When $\gamma = 0$, LP^1_γ has multiple optimal solutions along the line segment connecting $(0,2,0)$ and $(0,0,2)$. For $\gamma > 0$, the problem has a unique optimal solution at $(0,2,0)$. Since ζ_p is equivalent to the proximity to multiple optima, we have $\zeta_p \leq O(\gamma)$ when $\gamma > 0$. As for κ, it is always equal to 1 for LP^1_γ instances.

For the family of problems LP^1_γ, the values of Φ for different γ values are as follows:

γ	1e0	1e-1	1e-2	1e-3	1e-4
Φ of LP^1_γ	6.4e0	4.5e1	4.3e2	4.2e3	4.2e4

These values of Φ are clearly in a reciprocal relationship with the value of γ. This observation aligns with (5.4) of Theorem 5.1. Figure 1a shows the convergence performance of rPDHG on LP^1_γ instances for $\gamma \in \{0, 0.02, 0.005, 0.001\}$. The horizontal axis reports the number of iterations, while the vertical axis reports the relative error, defined as: $\mathcal{E}_r(x,y) := \frac{\|Ax^+ - b\|}{1+\|b\|} + \frac{\|(A^\top y - c)^+\|}{1+\|c\|} + \frac{|c^\top x^+ - b^\top y|}{1+|c^\top x^+| + |b^\top y|}$ for iterates (x,y). We use $\mathcal{E}_r(x,y)$ because it is easy to compute, and applicable when the problem has multiple optima. It is also a widely used standard tolerance (also used in Applegate et al. [4], Lu et al. [50]). The results clearly demonstrate that as $\gamma \searrow 0$, the number of iterations (of Stage I in particular) increases significantly, exhibiting an approximately reciprocal relationship with γ. Notably, the results also indicate that neither γ nor Φ influences the local convergence rate in Stage II. This finding is consistent with Theorem 4.1, which asserts that the local convergence rate is solely determined by $\|B^{-1}\|\|A\|$.

FIGURE 1. Convergence performance of rPDHG on two families of LP instances.

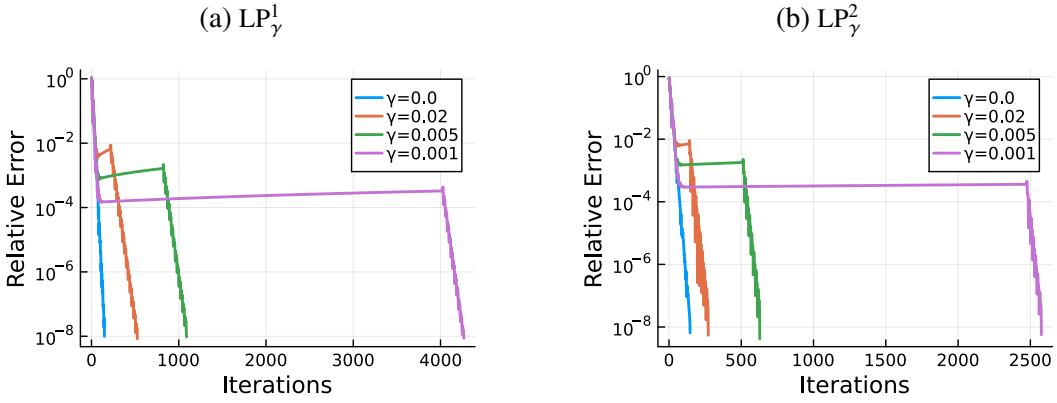

We further construct another family of standard-form LP instances (1.1) with data $\left(A^2, b^2, c^2\right)$, where

$$A^2 = \begin{bmatrix} 1 & 1 & -1 \\ 1 & 0 & 1 \end{bmatrix},\ c^2 = [-0.5, 1, 0.5],\ \text{ and } b^2 = b^2_\gamma := [1,1] + [\gamma, 2\gamma]$$

for parameter $\gamma \geq 0$. This LP family, denoted by LP^2_γ, is designed to illustrate the effect of perturbations $[\gamma, 2\gamma]$ on the right-hand side vector b^2_0. The value of κ is always equal to 1.22 for all LP^2_γ instances. When $\gamma = 0$, LP^2_γ has a unique optimal primal solution $[1, 0, 0]$. This solution is degenerate, implying multiple dual optimal solutions for LP^2_0. When $\gamma > 0$, the problem only has a unique dual optimal solution. Since ζ_d is also equivalent to the proximity to multiple dual optimal solutions, we have $\zeta_d \leq O(\gamma)$ when $\gamma > 0$.

For the family of problems LP^2_γ, the values of Φ for different γ values are as follows:

γ	1e0	1e-1	1e-2	1e-3	1e-4
Φ of LP^2_γ	6.7e0	3.4e1	3.4e2	3.4e3	3.4e4

These values of Φ still exhibit a clear reciprocal relationship with γ. Figure 1b shows the convergence performance of rPDHG for LP^2_γ instances for $\gamma \in \{0, 0.02, 0.005, 0.001\}$. Although the perturbations are now on the right-hand side vector b, our observations are nearly symmetric to those for LP^1_γ, which match the predictions of Theorems 5.1 and 4.1 again.

6.2. Two-stage Performance of rPDHG This subsection tests how well $\kappa\Phi$ and $\|B^{-1}\|\|A\|$ explain the practical two-stage performance of rPDHG. We consider (i) random LPs generated by Todd's model and (ii) LP relaxations from the MIPLIB 2017 library (see Gleixner et al. [26]). We also compare against the quantity R/δ proposed by Lu and Yang [45, 47] to characterize Stage-I performance of PDHG variants. Note that R/δ is only a portion of their iteration bounds that is relatively easy to compute and analyze, and its formal definition will follow later.

The random LP instances are generated according to Todd's random LP model:

Definition 6.1 (Random linear program). Let $u \sim \mathcal{N}(0, 1)$. Let the entries of the matrix A be independent and identically distributed (i.i.d.) copies of u. The primal and dual solutions $\hat{x}$ and $\hat{s}$ are generated as follows:

$$\hat{x}_\Theta \in \mathbb{R}^m_+, \quad \hat{s}_{\bar{\Theta}} \in \mathbb{R}^{n-m}_+, \quad \hat{x}_{\bar{\Theta}} = 0, \quad \hat{s}_\Theta = 0, \tag{6.1}$$

in which $\Theta = \{1, 2, \ldots, m\}$ and the components of $\hat{x}_\Theta$ and $\hat{s}_{\bar{\Theta}}$ are i.i.d. copies of $|u|$. The right-hand side b is generated by $b = A\hat{x}$, and the cost vector c is generated by $c = \hat{s}$. (Optionally, the cost vector $\bar{c}$ with the smallest norm is generated by $\bar{c} := \hat{s} + A^\top \hat{y}$, where $\hat{y} = \arg\min_{y \in \mathbb{R}^m} \|\hat{s} + A^\top \hat{y}\|$.)

The random LP in Definition 6.1 is Model 1 of Todd [59]. Variants of this model have been analyzed by Anstreicher et al. [2, 3], Ye [70] to elucidate the average performance of interior-point methods. One can observe that $\hat{x}$ and $\hat{s}$ are the optimal primal-dual solution because they are feasible and satisfy the complementary slackness condition $\hat{x}^\top \hat{s} = 0$. Components of $\hat{x}_\Theta$ and $\hat{s}_{\bar{\Theta}}$ are all nonzero almost surely, and the LP instance has a unique optimum with optimal basis Θ almost surely. Since the optimal basis and the optimal solution are known prior to solving the problem, $\kappa\Phi$ and $\|B^{-1}\|\|A\|$ used in the iteration bound results can be easily computed. Furthermore, replacing the cost vector c with the smallest-norm cost vector $\bar{c}$ does not influence the optimality and degeneracy of $\hat{x}$ and $\hat{s}$, but $\bar{c}$ lies in $\text{Null}(A)$, i.e., $A\bar{c} = 0$. To keep consistent with the theoretical results, we use $\bar{c}$ in the experiments.

We ran rPDHG on 100 randomly generated LP problems according to Definition 6.1 with $m = 50$ and $n = 100$. We use LP_i to denote the i-th instance, where $i \in \{1, 2, \ldots, 100\}$. Let $(x^{n,k}, y^{n,k})$ denote the iterates and define $s^{n,k} := c - A^\top y^{n,k}$. We deem an instance solved once $\|(x^{n,k}, s^{n,k}) - (x^\star, s^\star)\|_2 \leq 10^{-4}$, where $(x^\star, s^\star)$ is the optimal primal solution and dual slack vector. After solving the problem, we define the number of Stage I (optimal basis identification) iterations to be first iteration index after which the support set of $x^{n,k}$ matches that of $x^\star$ and remains unchanged thereafter. And we define the number of Stage II iterations (local convergence) to be the rest of the iterations.

According to Theorem 3.1, the overall number of iterations (in Stage I and Stage II) should be upper bounded by $O\big(\kappa\Phi \cdot \ln(\kappa\Phi \cdot \frac{\|w^\star\|}{\varepsilon})\big)$. Furthermore, Theorem 4.1 predicts that the number of iterations in Stage I and Stage II should be upper bounded by $O\big(\kappa\Phi \cdot \ln(\kappa\Phi)\big)$ and $O\big(\|B^{-1}\|\|A\| \cdot \ln(\frac{\xi}{\varepsilon})\big)$ respectively.

Figure 2a shows a scatter plot of the empirical number of Stage-I iterations versus $\kappa\Phi \ln(\kappa\Phi)$. We observe a clear association between these two quantities. Fitting a linear regression model (on the log scale used in the plot), the R^2 is equal to 0.6518, indicating that more than half of the variation in $\log_{10}$ (actual iteration number) is accounted for by this quantity. Furthermore, because Φ is equivalent to $\hat{\Phi}$ (Lemma 3.2), it also offers

empirical support for the role of the level set geometry condition number $\hat{\Phi}$, as proposed by Xiong and Freund [67].

FIGURE 2. Scatter plots of the number of OnePDHG iterations in Stage I (optimal basis identification) and Stage II (local convergence) versus $\kappa\Phi\ln(\kappa\Phi)$, $\|B^{-1}\|\|A\|$, and R/δ on random LP instances.

(a) Stage I (b) Stage II (c) Stage I

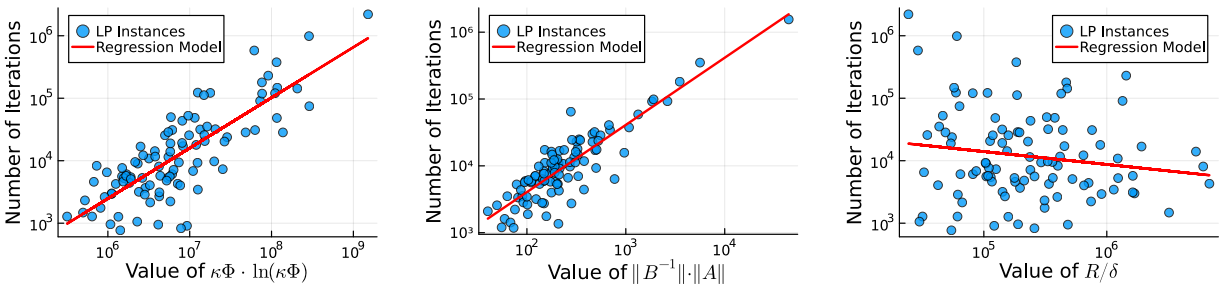

Similarly, we also validate the role of $\|B^{-1}\|\|A\|$ in the Stage-II performance of rPDHG through Figures 2b. One can also see a clear correlation between $\|B^{-1}\|\|A\|$ and the number of empirical iterations that rPDHG spent in the Stage II. The red line in Figure 2b is the linear regression model and its empirical R^2 value is equal to 0.7764.

Finally, Lu and Yang [45, 47] propose the quantity R/δ to characterize the Stage-I performance of PDHG variants (vanilla PDHG and restarted Halpern PDHG), where $R = 2\|z^{0,0} - z^\star\| + 2\|z^\star\| + 1$ and $\delta = \min\{\min_{i\in[n]:s_i^\star>0} \frac{s_i^\star}{\|A\|}, \min_{j\in[m]:x_j^\star>0} x_j^\star\}$ for the converging optimal solution $(x^\star, s^\star)$. Because their definition uses a different primal–dual step-size ratio τ/σ, we compute R/δ on a rescaled instance to remove this effect (see Applegate et al. [6] for the equivalence between step-size tuning and rescaling of b and c). Figure 2c compares R/δ with the observed Stage-I iteration counts of OnePDHG. For these randomly generated instances, the correlation is weak. We conjecture this is due to three reasons: (i) there are other quantities (e.g., a Hoffman constant) in addition to R/δ that is also in the iteration bound of PDHG (see Lu and Yang [45, 47]) (ii) the randomly generated LP instances are usually not ill-conditioned enough for R/δ to dominate, and (iii) the implemented rPDHG method is slightly different from the methods that Lu and Yang [45, 47] analyze.

In addition to randomly generated LPs, we test rPDHG on some LP relaxations from the MIPLIB 2017 library (see Gleixner et al. [26]). We first applied the presolver PaPILO (see Gleixner et al. [25]) to all instances, which detects inconsistent bounds, removes empty rows/columns, and eliminates fixed variables, among other reductions. We then converted each instance to the standard form (1.1) and applied Ruiz rescaling and Pock–Chambolle rescaling, as suggested by Applegate et al. [4]. Since many instances do not satisfy Assumption 2.1, we perturb the data to obtain unique optimal solutions almost surely: we replace c by $c + 0.01\|c\|_2 \cdot \frac{v}{\|v\|_2}$, where $v \in \mathbb{R}^n$ has i.i.d. entries distributed as $|u|$ with $u \sim \mathcal{N}(0,1)$, and we replace b by $b + 0.01\|b\|_2 \cdot \frac{A\bar{v}}{\|A\bar{v}\|_2}$, where $\bar{v}$ is an independent copy of v. These perturbations are scaled to be small relative to $\|b\|_2$ and $\|c\|_2$, but large enough to break degeneracy. We ran the same algorithmic setting as in the random-LP experiments, with a maximum of 10^7 iterations. To ensure solving most problems and avoid the tiny instances, we consider the instances whose n is between 100 and 1,000.

Figure 3 reports the results in these MIPLIB instances. Some instances did not meet the stopping criterion within 10^7 iterations, and some did not enter Stage II. These runs are indicated by dashed lines in the plots. Similar to Figure 2, we observe clear correlations between Stage I iterations and $\kappa\Phi\ln(\kappa\Phi)$ (Figure 3a) and between Stage II iterations and $\|B^{-1}\|\|A\|$ (Figure 3b). The regression models are fitted based only on the instances not on dashed lines. The R^2 values in these two figures are equal to 0.3627 and 0.4474, respectively. These values are lower than those for the random instances, in part because iteration counts are truncated. Finally, now one may observe a correlation between Stage I iteration counts and R/δ in Figure 3c. This is

FIGURE 3. Scatter plots of the number of OnePDHG iterations in Stage I (optimal basis identification) and Stage II (local convergence) versus $\kappa\Phi\ln(\kappa\Phi)$, $\|B^{-1}\|\|A\|$, and R/δ on selected MIPLIB 2017 LP relaxation instances.

(a) Stage I (b) Stage II (c) Stage I

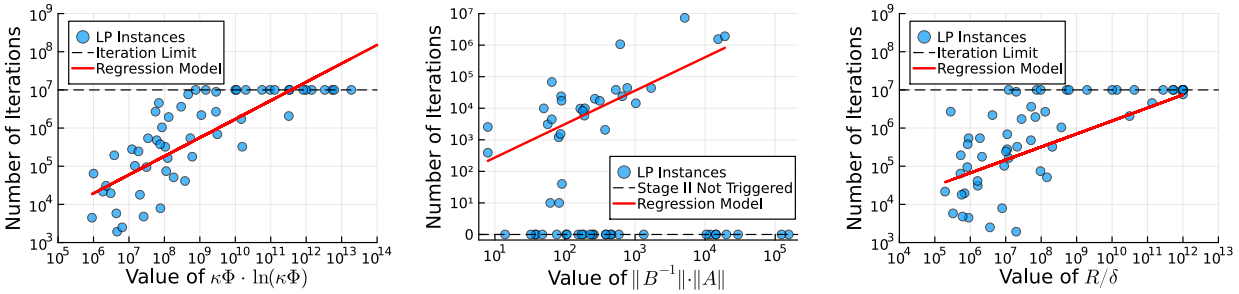

probably because some real-world instances are ill-conditioned enough for R/δ to play a more pronounced role.

7. Optimized Reweighting and Step-Size Ratio Having shown the new accessible iteration bounds, a natural question arises. *What new practical insights of the performance of rPDHG can these new bounds enable?* After the release of this paper, Xiong [64] shows that rPDHG has a high-probability polynomial iteration bound on LP instances whose data follows certain random distributions, via directly analyzing the new bounds presented in this paper on the random LP instances.

In this section we provide another answer, by studying how rescaling the right-hand side vector b and the cost vector c (or equivalently, changing the ratio between the primal and dual step-sizes) affects the quantity Φ that drives our computational guarantees. The specific rescaling we analyze is not intended as a practical method. Instead, it serves as an idealized benchmark that allows us to partially explain why an existing primal–dual step-size heuristic works so well in practice.

Strategic configuration of the step-size ratio, also denoted as the "primal weight", has been observed to yield significant improvements in rPDHG's convergence rate (see, e.g., Applegate et al. [4], Xiong and Freund [66]). Applegate et al. [4] reparameterize the step-sizes as $\tau = \eta/\omega$ and $\sigma = \eta\omega$, and interpret the scalar $\omega > 0$ as a *primal weight* that controls the relative scaling between the primal and dual updates. At each restart, they compare the distances traveled in the primal and dual spaces between two consecutive restart points, say Δ_x^n and Δ_y^n, and update ω by exponentially smoothing the ratio Δ_y^n/Δ_x^n on a log scale. Applegate et al. [4] claim this heuristic aims to choose ω so that the weighted distances from the algorithm iterates x and y to the optimal solutions, $\sigma\|x - x^\star\|$ and $\tau\|y - y^\star\|$, are approximately equal.

This step-size heuristic has no convergence guarantee and sometimes even leads to divergence in practice, but it has been used in many large-scale LP solvers (see, e.g., Applegate et al. [4], Lu et al. [41], Lu and Yang [46], Lu et al. [50]), convex quadratic program solver (see, e.g., Huang et al. [33], Lu and Yang [49]) and even other solvers (see, e.g., Lin et al. [37], Liu et al. [38]). In spite of its practical success, there is little theoretical understanding of why it works so well for LPs in particular. Xiong and Freund [65] verify the value of tuning step-sizes by providing a formula for an "optimized" step-size ratio, but this formula is dramatically different because it incorporates the terms μ_p and μ_d.

In this section, we examine the step-size ratio from the perspective of reweighting the primal and dual variables. We present a simple expression of the "optimized" reweights, through which we can validate the heuristic of choosing the primal weights that balance the primal and dual iterates. For reweights $\omega_1, \omega_2 > 0$, the reweighted problem is as follows:

$$\min_{x\in\mathbb{R}^n} \quad (\omega_1 \cdot c)^\top x \qquad \text{s.t.} \quad Ax = \omega_2 \cdot b \ , \ x \geq 0 \tag{7.1}$$

This reweighting can be interpreted as changing the relative scaling of the primal and dual variables. Applegate et al. [6] demonstrate that applying PDHG to the reweighted problem (7.1) with primal-dual step-sizes $(\tilde{\tau}, \tilde{\sigma})$

is equivalent to applying PDHG to the original problem (1.1) with primal-dual step-sizes (τ, σ) when the step-sizes are related by $(\tau, \sigma) = \left(\frac{\omega_1}{\omega_2}\tilde{\tau}, \frac{\omega_2}{\omega_1}\tilde{\sigma}\right)$.

For simplicity of notation, we let $\tilde{x}^\star$ and $\tilde{s}^\star$ represent the optimal solutions of the reweighted problem. Then $\tilde{x}^\star = \omega_2 \cdot x^\star$ and $\tilde{s}^\star = \omega_1 \cdot s^\star$. Thus, rPDHG can directly run on (7.1) and $(x^\star, s^\star)$ is computed by recovering from the convergent solution of rPDHG. We use Φ_{ω_1,ω_2} to denote the Φ value of (7.1), e.g. $\Phi_{1,1}$ is the value of Φ for (1.1). Although κ remains invariant under reweighting, Φ_{ω_1,ω_2} may become different than $\Phi_{1,1}$, leading to different convergence rates of rPDHG.

The theorem below shows that, once $\frac{\omega_1}{\omega_2} = \frac{\|x^\star\|_1}{\|s^\star\|_1}$, then (ω_1, ω_2) is actually a "nearly optimal" reweighting within a factor of 2.

THEOREM 7.1. *Suppose that Assumption 2.1 holds. When $\hat{\omega}_1, \hat{\omega}_2 > 0$, $\frac{\hat{\omega}_1}{\hat{\omega}_2} = \frac{\|x^\star\|_1}{\|s^\star\|_1}$ if and only if $\|\tilde{x}^\star\|_1 = \|\tilde{s}^\star\|_1$. With this choice of reweights, the corresponding Φ value satisfies*

$$\Phi_{\hat{\omega}_1,\hat{\omega}_2} = 2 \cdot \max\left\{\frac{\|x^\star\|_1}{\zeta_d}, \frac{\|s^\star\|_1}{\zeta_p}\right\} \le 2\left(\min_{\omega_1,\omega_2>0} \Phi_{\omega_1,\omega_2}\right). \tag{7.2}$$

As an immediate consequence of Theorems 3.1 and 7.1, applying rPDHG to the reweighted problem (7.1) so that $\|\tilde{x}^\star\|_1 = \|\tilde{s}^\star\|_1$ yields a linear convergence bound in which the coefficient Φ is replaced by $O\left(\min_{\omega_1,\omega_2>0} \Phi_{\omega_1,\omega_2}\right)$. The precise statement and the proof of Theorem 7.1 are given in Appendix C.

Comparing $\Phi_{\hat{\omega}_1,\hat{\omega}_2}$ and the value of Φ in Theorem 5.1, $\Phi_{\hat{\omega}_1,\hat{\omega}_2}$ is $2 \cdot \max\left\{\frac{\|x^\star\|_1}{\zeta_d}, \frac{\|s^\star\|_1}{\zeta_p}\right\}$, which is potentially much smaller than $\Phi = \frac{\|x^\star\|_1 + \|s^\star\|_1}{\min\{\zeta_p, \zeta_d\}}$ when running rPDHG for the original problem. This distinction implies that when $\frac{\|x^\star\|_1}{\zeta_p}$ and $\frac{\|s^\star\|_1}{\zeta_d}$ significantly exceed $\frac{\|x^\star\|_1}{\zeta_d}$ and $\frac{\|s^\star\|_1}{\zeta_p}$, adjusting the reweights (or the equivalent step-size ratio) may have a huge benefit for rPDHG. This observation is also consistent with the experimental validation presented in Figure 3 of Xiong and Freund [65].

According to Theorem 7.1, the optimized reweighting $(\hat{\omega}_1, \hat{\omega}_2)$ balances the ℓ_1-norms of the optimal solutions of the reweighted problem, i.e., $\|\tilde{x}^\star\|_1 = \|\tilde{s}^\star\|_1$. Since $x^\star$ and $s^\star$ are unknown *a priori*, this reweighting is not directly implementable as a practical algorithmic rule. Nevertheless, it provides intuition for the empirical effectiveness of the widely used "primal-weight adjustment."

As discussed by Applegate et al. [6], tuning the primal weight so that the weighted distances $\sigma\|x - x^\star\|$ and $\tau\|y - y^\star\|$ are of comparable magnitude can be interpreted as running rPDHG in an implicitly reweighted problem in which primal and dual progress are balanced. In other words, the primal-weight adjustment seeks to balance the distances from the current iterate to the primal and dual optima in the reweighted space. This suggests the following interpretation of its practical effectiveness: *Balancing primal and dual distances to optimal solutions via reweighting may improve the practical performance of rPDHG, in part because it tends to induce reweightings that also balance the magnitudes of the optimal solutions.*

Finally, note that in practical implementations of rPDHG for LPs, the primal weight adjustment is always used together with other heuristics (e.g., preconditioning the constraint matrix A to reduce κ of A). It would be hard to totally isolate the effect of the primal weight adjustment alone in practice.

Appendix A: Proofs of Section 3

A.1. Technical lemmas about the normalized duality gap and β-restart condition Note that Section 4.1 has reviewed the basic information of the β-restart condition and the normalized duality gap, which were omitted in Section 3 but will be heavily used throughout the proofs. Here we formally present the sublinear convergence result of the normalized duality gap summarized by Applegate et al. [6], using an equivalent result presented by Xiong and Freund [65]:

LEMMA A.1 **(Corollary 2.4 of Xiong and Freund [65])**. *Suppose that σ, τ satisfy* (2.7)*. Then for any $z^0 := (x^0, y^0)$ with $x^0 \in \mathbb{R}^n_+$, the following inequality holds for all $k \ge 1$:*

$$\rho(\|\bar{z}^k - z^0\|_M; \bar{z}^k) \le \frac{8\,\mathrm{Dist}_M(z^0, \mathcal{Z}^\star)}{k}. \tag{A.1}$$

Then an upper bound on the number of iterations for inner loops in rPDHG is as follows. The following lemma is essentially part of Applegate et al. [6, Theorem 2].

LEMMA A.2. *Suppose that* (4.6) *holds for* $z^{n,0}$ *and* $z^{n-1,0}$ *with* $\mathcal{L} > 0$. *Whenever* $k \geq \frac{8\mathcal{L}}{\beta}$, *sufficient decrease has been made on the normalized duality gap, i.e., the restart condition* (4.5) *is satisfied.*

Proof. For $k \geq 1$, Lemma A.1 implies:

$$\rho(\|\bar{z}^{n,k} - z^{n,0}\|_M; \bar{z}^{n,k}) \leq \frac{8\,\mathrm{Dist}_M(z^{n,0}, \mathcal{Z}^\star)}{k}\,. \tag{A.2}$$

If $\rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0}) = 0$, then $z^{n,0}$ is a saddle point of (2.2) and thus $z^{n,0} \in \mathcal{Z}^\star$. In this case, $z^{n,k} = z^{n,0}$ for all $k \geq 1$, and any $k \geq 1$ satisfies (4.5).

If $\rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0}) \neq 0$, dividing both sides of (A.2) by $\rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0})$ yields:

$$\frac{\rho(\|\bar{z}^{n,k} - z^{n,0}\|_M; \bar{z}^{n,k})}{\rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0})} \leq \frac{8}{k} \cdot \frac{\mathrm{Dist}_M(z^{n,0}, \mathcal{Z}^\star)}{\rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0})} \leq \frac{8}{k} \cdot \mathcal{L} \tag{A.3}$$

where the last inequality follows from (4.6). Therefore, when $k \geq \frac{8\mathcal{L}}{\beta}$, the right-hand side is no larger than β, i.e., $\frac{8}{k} \cdot \mathcal{L} \leq \beta$, and the restart condition (4.5) is satisfied. □

A.2. Proof of Proposition 3.1

Proof of Proposition 3.1. For the second term of multiplication in the right-hand side of (3.1), we have

$$\begin{aligned}
&\max\left\{\max_{1\leq j\leq n-m} \frac{\sqrt{\left\|(B^{-1}N)_{\cdot,j}\right\|^2+1}}{s^\star_{m+j}}, \max_{1\leq i\leq m} \frac{\sqrt{\left\|(B^{-1}N)_{i,\cdot}\right\|^2+1}}{x^\star_i}\right\} \\
\leq\ & \frac{\max\left\{\max_{1\leq j\leq n-m}\sqrt{\left\|(B^{-1}N)_{\cdot,j}\right\|^2+1}, \max_{1\leq i\leq m}\sqrt{\left\|(B^{-1}N)_{i,\cdot}\right\|^2+1}\right\}}{\min\left\{\min_{1\leq i\leq m} x^\star_i, \min_{1\leq j\leq n-m} s^\star_{m+j}\right\}} \\
=\ & \frac{\max\left\{\sqrt{\|B^{-1}N\|_{1,2}^2+1}, \sqrt{\|B^{-1}N\|_{2,\infty}^2+1}\right\}}{\min_{1\leq i\leq n}\left\{x^\star_i + s^\star_i\right\}}
\end{aligned} \tag{A.4}$$

where the equality holds because $x^\star$ and $s^\star$ are strictly complementary. Furthermore, because $\|\cdot\|_{1,2}$ and $\|\cdot\|_{2,\infty}$ norms are upper bounded by the spectral norm (denoted by $\|\cdot\|$), we have the following inequalities:

$$\begin{aligned}
&\max\left\{\sqrt{\|B^{-1}N\|_{1,2}^2+1}, \sqrt{\|B^{-1}N\|_{2,\infty}^2+1}\right\} \leq \sqrt{\|B^{-1}N\|^2+1} = \sqrt{\sigma^+_{\max}(B^{-1}NN^\top B^{-\top})+1} \\
&= \sqrt{\sigma^+_{\max}(B^{-1}(NN^\top + BB^\top)B^{-\top})} = \sqrt{\sigma^+_{\max}(B^{-1}AA^\top B^{-\top})} = \|B^{-1}A\|\,.
\end{aligned}$$

Finally, by complementary slackness, $\|x^\star\|_1 + \|s^\star\|_1 = \|x^\star + s^\star\|_1$. Applying these inequalities to the definition of Φ in (3.1) completes the proof. □

A.3. Proof of Lemma 3.1

In this subsection, we prove Lemma 3.1. We begin with the following lemma.

LEMMA A.3. *Suppose that* $Ac = 0$. *Algorithm 1 (rPDHG) is run starting from* $z^{0,0} = (x^{0,0}, y^{0,0}) = (0, 0)$, *and the step-sizes* σ *and* τ *satisfy* (2.7). *Then for all* $n \geq 1$, *it holds that*

$$\mathrm{Dist}_M(z^{n,0}, \mathcal{Z}^\star) \leq \sqrt{2}c_{\tau,\sigma} \cdot \mathrm{Dist}(w^{n,0}, \mathcal{W}^\star) \leq (3\sqrt{2}+4)c^2_{\tau,\sigma} \cdot \hat{\Phi} \cdot \rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0})\,. \tag{A.5}$$

In other words, condition (4.6) *holds with* $\mathcal{L} = (3\sqrt{2}+4)c^2_{\tau,\sigma}\hat{\Phi}$.

This lemma is Lemma 3.13 of Xiong and Freund [67] by taking limits as δ approaches 0 on both sides.

LEMMA A.4 (**Proposition 3.7 of Xiong and Freund [67]**). *Suppose that* $Ac = 0$ *and* $z^{0,0} = (0,0)$. *Then* $\|w^{0,0} - w^\star\| = \|(0,c) - w^\star\| \le \|w^\star\|$.

Proof of Lemma 3.1. By Assumption 2.1, let $\mathcal{Z}^\star$ be $\{z^\star\}$, then $\text{Dist}_M(z^{n,0}, \mathcal{Z}^\star) = \|z^{n,0} - z^\star\|_M$, and Lemma A.3 states:

$$\|z^{n,0} - z^\star\|_M \le \sqrt{2} c_{\tau,\sigma} \cdot \|w^{n,0} - w^\star\| \le (3\sqrt{2}+4) c_{\tau,\sigma}^2 \cdot \hat{\Phi} \cdot \rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0}) \tag{A.6}$$

Substituting the step-sizes, we have $c_{\tau,\sigma} = \sqrt{2\kappa}$, $(3\sqrt{2}+4)c_{\tau,\sigma}^2 \approx 16.4853\kappa \le 16.5\kappa$, and

$$\|w^{n,0} - w^\star\| \le (3\sqrt{2}+4)\sqrt{\kappa}\hat{\Phi} \cdot \rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0}) \le 8.25\sqrt{\kappa}\hat{\Phi} \cdot \rho(\|z^{n,0} - z^{n-1,0}\|_M; z^{n,0}) \ . \tag{A.7}$$

Let $\bar{T}_\varepsilon$ denote the total number of OnePDHG iterations required to obtain the first outer iteration N that satisfies $\rho(\|z^{N,0} - z^{N-1,0}\|_M; z^{N,0}) \le \varepsilon$. Theorem 3.5 of Xiong and Freund [67] and (A.6) then guarantee:

$$\bar{T}_{\bar{\varepsilon}} \ \le \ 23 \cdot 16.5\kappa\hat{\Phi} \cdot \ln\left(\frac{23\|z^{0,0} - z^\star\|_M)}{\varepsilon}\right) \le 380\kappa\hat{\Phi} \cdot \ln\left(\frac{23\|z^{0,0} - z^\star\|_M}{\varepsilon}\right) \ . \tag{A.8}$$

Note that (A.7) ensures that when the normalized duality gap is sufficiently small, the distance to the optimal solution is correspondingly small. Therefore,

$$T \le \bar{T}_{\frac{\varepsilon}{8.25\sqrt{\kappa}\hat{\Phi}}} \le 380\kappa\hat{\Phi} \cdot \ln\left(\frac{189.75\sqrt{\kappa}\hat{\Phi} \cdot \|z^{0,0} - z^\star\|_M}{\varepsilon}\right) \le 380\kappa\hat{\Phi} \cdot \ln\left(\frac{379.5\kappa\hat{\Phi} \cdot \|w^{0,0} - w^\star\|}{\varepsilon}\right) \tag{A.9}$$

where the last inequality follows from (B.2) and $c_{\tau,\sigma} = \sqrt{2\kappa}$. Finally, by Lemma A.4, we have $\|w^{0,0} - w^\star\| \le \|w^\star\|$. Consequently, (A.9) leads to (3.6) of Lemma 3.1. □

Furthermore, similar results also exist for the goal of obtaining an iterate $z^{N,0}$ satisfying $\|z^{N,0} - z^\star\|_M \le \varepsilon$, which will be useful later in Appendix B.

REMARK A.1. Under the same conditions as Lemma 3.1, let $\tilde{T}$ denote the total number of OnePDHG iterations required to obtain the first outer iteration N that satisfies both $\rho(\|z^{N-1,0} - z^{N,0}\|_M; z^{N,0}) \le \frac{\varepsilon}{16.5\kappa\hat{\Phi}}$ and $\|z^{N,0} - z^\star\|_M \le \varepsilon$. Then, $\tilde{T} \le 380\kappa\hat{\Phi} \cdot \ln\left(\frac{760\kappa^{1.5}\hat{\Phi} \cdot \|w^\star\|}{\varepsilon}\right)$.

The proof of Remark A.1 is almost identical to that of Lemma 3.1, except (A.9) is replaced by the inequality derived from $\tilde{T} \le \bar{T}_{\frac{\varepsilon}{16.5\kappa\hat{\Phi}}}$ (due to (A.6)).

Appendix B: Proofs of Section 4

B.1. Proofs of Lemmas 4.1 and 4.2 First of all, we defined the $\tilde{M}$-norm

$$\|(x,y)\|_{\tilde{M}} := \sqrt{\frac{1}{\tau}\|x\|^2 + \frac{1}{\sigma}\|y\|^2} \quad \text{where} \quad \tilde{M} := \begin{pmatrix} \frac{1}{\tau} I_n & \\ & \frac{1}{\sigma} I_m \end{pmatrix}. \tag{B.1}$$

When τ and σ are sufficiently small, the M-norm and $\tilde{M}$-norm are equivalent up to well-specified constants related to τ and σ (Proposition 2.8 of Xiong and Freund [65]). For any point $z := (x,y) \in \mathbb{R}^{n+m}$, and $w := (x, c - A^\top y)$, it holds that

$$\sqrt{1 - \sqrt{\tau\sigma}\lambda_{\max}} \cdot \|z\|_{\tilde{M}} \le \|z\|_M \le \sqrt{2} \cdot \|z\|_{\tilde{M}} \le \sqrt{2} c_{\tau,\sigma} \cdot \|w\| \ . \tag{B.2}$$

For example, if $\tau = \frac{1}{2\kappa}$ and $\sigma = \frac{1}{2\lambda_{\max}\lambda_{\min}}$, then (B.2) becomes $\frac{\sqrt{2}}{2}\|z\|_{\tilde{M}} \le \|z\|_M \le \sqrt{2}\|z\|_{\tilde{M}} \le 2\sqrt{\kappa}\|w\|$. This result will be extensively used later.

Proof of Lemma 4.1. Let $\bar{s} = c - A^\top \bar{y}$ and $\bar{w} = (\bar{x}, \bar{s})$. From the definition of $\rho(r;\cdot)$ we have:

$$L(\bar{x}, y) - L(x, \bar{y}) \le r\rho(r;\bar{z}) \text{ for any } z = (x, y) \in B(r;\bar{z}) \tag{B.3}$$

where recall that $B(r;\bar{z})$ is defined in Definition 4.1.

Firstly, we prove that

$$\|\bar{x} - x^\star\| \le \frac{\|B^{-1}\|}{\sqrt{\sigma}} \cdot \rho(r;\bar{z}) \ . \tag{B.4}$$

As the optimal basis is unique, $x^\star$ is represented by its basic and nonbasic parts: $x^\star_{[m]} = B^{-1}b$ and $x^\star_{[n]\setminus[m]} = 0$. Consequently, due to (4.8),

$$\left\|\bar{x}_{[m]} - x^\star_{[m]}\right\| = \left\|\bar{x}_{[m]} - B^{-1}b\right\| \le \|B^{-1}\| \cdot \left\|B\bar{x}_{[m]} - b\right\| \text{ and } \left\|\bar{x}_{[n]\setminus[m]} - x^\star_{[n]\setminus[m]}\right\| = 0 \ . \tag{B.5}$$

Let $u = b - B\bar{x}_{[m]}$ and define $y := \bar{y} + \sqrt{\sigma} r \cdot u/\|u\|$. Let $z := (\bar{x}, y)$, and then $z \in B(r;\bar{z})$. Thus, from (B.3) we obtain

$$r\rho(r;\bar{z}) \ge L(\bar{x}, y) - L(\bar{x}, \bar{y}) = (b - A\bar{x})^\top (y - \bar{y}) = (b - B\bar{x}_{[m]})^\top (y - \bar{y}) = \sqrt{\sigma} r \|u\|, \tag{B.6}$$

implying $\|u\| = \|b - B\bar{x}_{[m]}\| \le \frac{\rho(r;\bar{z})}{\sqrt{\sigma}}$. Substituting this result back into (B.5) yields (B.4).

Secondly, we prove that

$$\|\bar{y} - y^\star\| \le \frac{\|B^{-1}\|}{\sqrt{\tau}} \cdot \rho(r;\bar{z}) \ . \tag{B.7}$$

Given that the optimal basis is $[m]$, we have $B^\top y^\star = c_{[m]}$, and $y^\star = (B^\top)^{-1} c_{[m]}$. Consequently,

$$\left\|\bar{y} - y^\star\right\| = \left\|\bar{y} - (B^\top)^{-1} c_{[m]}\right\| \le \left\|(B^\top)^{-1}\right\| \cdot \left\|B^\top \bar{y} - c_{[m]}\right\| = \left\|B^{-1}\right\| \cdot \left\|B^\top \bar{y} - c_{[m]}\right\| \ . \tag{B.8}$$

Let $v = c_{[m]} - B^\top \bar{y}$ and define x as follows: $x_{[m]} := \bar{x}_{[m]} - \sqrt{\tau} r \cdot \frac{v}{\|v\|}$ and $x_{[n]\setminus[m]} := 0$. Note that due to the condition $\bar{x}_i \ge r\sqrt{\tau}$ for all $i \in [m]$ in (4.8), x remains in $\mathbb{R}^n_+$. Now, let $z := (x, \bar{y})$, and then $z \in B(r;\bar{z})$. Thus, from (B.3), we derive:

$$r\rho(r;\bar{z}) \ge L(\bar{x}, \bar{y}) - L(x, \bar{y}) = (c - A^\top \bar{y})^\top (\bar{x} - x) = (c_{[m]} - B^\top \bar{y})^\top (\bar{x}_{[m]} - x_{[m]}) = \sqrt{\tau} r \|v\| \ , \tag{B.9}$$

implying $\|v\| = \|B^\top \bar{y} - c_{[m]}\| \le \frac{\rho(r;\bar{z})}{\sqrt{\tau}}$. Substituting this result back into (B.8) yields (B.7).

Finally, we can assert that

$$\|\bar{w} - w^\star\|^2 = \|\bar{x} - x^\star\|^2 + \|\bar{s} - s^\star\|^2 \le \|\bar{x} - x^\star\|^2 + \|A\|^2 \|\bar{y} - y^\star\|^2 \overset{(B.4)(B.7)}{\le} \left(\frac{\|B^{-1}\|^2}{\sigma} + \frac{\|B^{-1}\|^2 \|A\|^2}{\tau}\right) \rho(r;\bar{z})^2 \tag{B.10}$$

Due to (B.2), $\|\bar{z} - z^\star\|_M \le \sqrt{2} c_{\tau,\sigma} \cdot \|\bar{w} - w^\star\|$. Applying it on the left-hand side of (B.10) proves (4.9). Thus, the proof is complete. □

Before proving Lemma 4.2, we recall the nonexpansive property of rPDHG proven by Applegate et al. [6]. We use the more convenient format presented by Xiong and Freund [67].

LEMMA B.1 (**Lemma 2.2 of Xiong and Freund [67]**). *For any $n, k \ge 0$, it holds that $\|\bar{z}^{n,k} - z^\star\|_M \le \left\|z^{n,0} - z^\star\right\|_M$. For any n_1, n_2 so that $n_2 \ge n_1 \ge 0$, it holds that $\|z^{n_2,0} - z^\star\|_M \le \|z^{n_1,0} - z^\star\|_M$.*

Proof of Lemma 4.2. The proof contains two steps. In the first step, we prove that when

$$\left\|z^{t,0} - z^\star\right\|_M \le \frac{\sqrt{1 - \sqrt{\tau\sigma}\|A\|}}{3\sqrt{\tau}} \cdot \xi \ , \tag{B.11}$$

then for any $N \ge t$ and $k \ge 1$, item (i) of (4.12) holds.

We begin by establishing a lower bound for $\min_{1\le i\le m}\bar{x}_i^{N,k}$, the left-hand side of item (i):

$$\min_{1\le i\le m}\bar{x}_i^{N,k}\ge\left(\min_{1\le i\le m}x_i^\star\right)-\left\|\bar{x}_{[m]}^{N,k}-x_{[m]}^\star\right\|\ge\xi-\left\|\bar{x}_{[m]}^{N,k}-x_{[m]}^\star\right\|\ . \tag{B.12}$$

The second term on the right-hand side of (B.12) can be bounded as follows:

$$\left\|\bar{x}_{[m]}^{N,k}-x_{[m]}^\star\right\|\le\left\|\bar{x}^{N,k}-x^\star\right\|\overset{\text{(B.1)}}{\le}\sqrt{\tau}\left\|\bar{z}^{N,k}-z^\star\right\|_{\tilde{M}}\overset{\text{(B.2)}}{\le}\frac{\sqrt{\tau}\left\|\bar{z}^{N,k}-z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}}\overset{\text{Lemma B.1}}{\le}\frac{\sqrt{\tau}\left\|z^{t,0}-z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} \tag{B.13}$$

for any $t\le N$. In addition, due to the nonexpansive property, the right-hand side of item (i) is upper bounded by:

$$\left\|\bar{z}^{N,k}-z^{N,0}\right\|_M\le\left\|\bar{z}^{N,k}-z^\star\right\|_M+\left\|z^\star-z^{N,0}\right\|_M\le2\left\|z^\star-z^{t,0}\right\|_M\le\frac{2\left\|z^{t,0}-z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} \tag{B.14}$$

where the last inequality holds because $\sqrt{1-\sqrt{\tau\sigma}\|A\|}\le1$. Finally, we have

$$\begin{aligned}&\left(\min_{1\le i\le m}\bar{x}_i^{N,k}\right)-\sqrt{\tau}\left\|\bar{z}^{N,k}-z^{N,0}\right\|_M\overset{\text{(B.12),(B.13)}}{\ge}\xi-\frac{\sqrt{\tau}\left\|z^{t,0}-z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}}-\sqrt{\tau}\left\|\bar{z}^{N,k}-z^{N,0}\right\|_M\\&\overset{\text{(B.14)}}{\ge}\xi-\frac{\sqrt{\tau}\left\|z^{t,0}-z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}}-2\sqrt{\tau}\left\|z^\star-z^{t,0}\right\|_M\ge\xi-\frac{3\sqrt{\tau}\left\|z^{t,0}-z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}}\ ,\end{aligned} \tag{B.15}$$

in which the last term is nonnegative when (B.11) holds. This shows that when $\left\|z^{t,0}-z^\star\right\|_M$ is small enough and satisfies (B.11), the left-hand side of (B.15) is nonnegative, and item (i) of (4.12) holds. This completes the first step of the proof.

In the second step, we prove that when

$$\left\|z^{t,0}-z^\star\right\|_M\le\frac{\sqrt{\tau}\sqrt{1-\sqrt{\tau\sigma}\|A\|}}{2}\cdot\xi, \tag{B.16}$$

then for any $N\ge t$ and $k\ge1$ item (ii) of (4.12) holds.

Let α denote $\left(\min_{1\le j\le n-m,k\ge0}s_{m+j}^{N,k}\right)$, then according to (2.6) (in Line 5 of Algorithm 1),

$$x_{[n]\setminus[m]}^{N,k}=\left(x_{[n]\setminus[m]}^{N,k-1}-\tau s^{N,k-1}\right)^+\le\left(x_{[n]\setminus[m]}^{N,k-1}-\tau\alpha\right)^+$$

and applying this inequality recursively as k decreases to 0 yields

$$x_{[n]\setminus[m]}^{N,k}\le\left(x_{[n]\setminus[m]}^{N,k-1}-\tau\alpha\right)^+\le\left(x_{[n]\setminus[m]}^{N,k-2}-2\tau\alpha\right)^+\le\cdots\le\left(x_{[n]\setminus[m]}^{N,0}-k\tau\alpha\right)^+\ . \tag{B.17}$$

Therefore, if $\alpha\ge0$ and

$$\left(\max_{1\le j\le n-m}x_{m+j}^{N,0}\right)\le\tau\cdot\left(\min_{1\le j\le n-m,k\ge0}s_{m+j}^{N,k}\right)\ , \tag{B.18}$$

then $x_{[n]\setminus[m]}^{N,k}=0$ for all $k\ge1$ and thus $\bar{x}_{[n]\setminus[m]}^{N,k}=\frac{1}{k}\sum_{i=1}^k x_{[n]\setminus[m]}^{N,i}=0$ for all $k\ge1$. Intuitively, since Assumption 2.1 holds, and $x^{N,0}$ and $s^{N,k}$ converge to the optimal solution, $x_{[n]\setminus[m]}^{N,0}$ should converge to 0 and $s_{[n]\setminus[m]}^{N,k}$ should stay away from 0. In the rest of the proof we will establish a lower bound of $\left(\min_{1\le j\le n-m,k\ge0}s_{m+j}^{N,k}\right)$ and an upper bound of $\left(\max_{1\le j\le n-m}x_{m+j}^{N,0}\right)$.

Similar to the first step, for the dual iteration $y^{N,k}$ (and $s^{N,k} = c - A^\top y^{N,k}$), we have

$$\begin{aligned}\left(\min_{1\le j\le n-m} s^{N,k}_{m+j}\right) &\ge \left(\min_{1\le j\le n-m} s^{\star}_{m+j}\right) - \left\|s^{N,k}_{[n]\setminus[m]} - s^{\star}_{[n]\setminus[m]}\right\| \\ &= \xi - \left\|A^\top_{[n]\setminus[m]}(y^{N,k} - y^\star)\right\| \ge \xi - \left\|A_{[n]\setminus[m]}\right\| \cdot \left\|y^{N,k} - y^\star\right\| .\end{aligned} \tag{B.19}$$

As for $\left\|y^{N,k} - y^\star\right\|$ in the last term of the above inequality, we derive:

$$\left\|y^{N,k} - y^\star\right\| \overset{\text{(B.1)}}{\le} \sqrt{\sigma}\left\|z^{N,k} - z^\star\right\|_{\tilde{M}} \overset{\text{(B.2)}}{\le} \frac{\sqrt{\sigma}\left\|z^{N,k} - z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} \overset{\text{Lemma B.1}}{\le} \frac{\sqrt{\sigma}\left\|z^{t,0} - z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} . \tag{B.20}$$

Combining (B.19) and (B.20) yields a valid lower bound of $\left(\min_{1\le j\le n-m} s^{N,k}_{m+j}\right)$ for all $k \ge 0$:

$$\left(\min_{1\le j\le n-m} s^{N,k}_{m+j}\right) \ge \xi - \sqrt{\sigma}\|A_{[n]\setminus[m]}\| \cdot \frac{\left\|z^{t,0} - z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} \ge \xi - \frac{1}{\sqrt{\tau}} \cdot \frac{\left\|z^{t,0} - z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} , \tag{B.21}$$

where the second inequality holds because $\sqrt{\tau\sigma}\|A_{[n]\setminus[m]}\| \le \sqrt{\tau\sigma}\|A\| \le 1$ (due to the step-size requirement (2.7)). The above (B.21) presents a lower bound of $\left(\min_{1\le j\le n-m, k\ge 0} s^{N,k}_{m+j}\right)$.

On the other hand, we also have the following upper bound of $\left(\max_{1\le j\le n-m} x^{N,0}_{m+j}\right)$ for $N \ge t$:

$$\left(\max_{1\le j\le n-m} x^{N,0}_{m+j}\right) \le \left\|x^{N,0}_{[n]\setminus[m]}\right\| = \left\|x^{N,0}_{[n]\setminus[m]} - x^{\star}_{[n]\setminus[m]}\right\| \le \|x^{N,0} - x^\star\| \overset{\text{(B.13)}}{\le} \frac{\sqrt{\tau}\left\|z^{t,0} - z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} . \tag{B.22}$$

Finally, applying the upper and lower bounds (B.21) and (B.22), we find that (B.18) holds when

$$\frac{\sqrt{\tau}\left\|z^{t,0} - z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} \le \tau\xi - \sqrt{\tau} \cdot \frac{\left\|z^{t,0} - z^\star\right\|_M}{\sqrt{1-\sqrt{\tau\sigma}\|A\|}} ,$$

which is equivalent to (B.16). This completes the second step of the proof.

Finally, when both (B.11) and (B.16) hold, which is satisfied by (4.11), then for any $N \ge t$ and $k \ge 1$, both item (i) and (ii) of (4.12) hold. Furthermore, (B.13) and (B.12) ensure $\bar{x}^{N,k}_{\Theta} > 0$, while the item (ii) ensures $\bar{x}^{N,k}_{\bar{\Theta}} = 0$. Therefore, the positive components of $\bar{x}^{N,k}$ correspond exactly to the optimal basis. This completes the proof. □

B.2. Proof of Theorem 4.1

Proof of Theorem 4.1. We first prove (4.1). According to Lemma 4.2, once N_0 satisfies $\|z^{N_0,0} - z^\star\|_M \le \bar{\varepsilon}$, which is equivalent to:

$$\|z^{N_0,0} - z^\star\|_M \le \bar{\varepsilon} = \frac{\sqrt{2}}{6} \cdot \frac{1}{\sqrt{2\kappa}} \cdot \xi = \frac{\xi}{6\sqrt{\kappa}} , \tag{B.23}$$

then for all $N > N_0$, we have:

$$(i)\ \ x^{N,0}_i \ge \sqrt{\tau}\|z^{N,0} - z^{N-1,0}\|_M \ \text{ for } i \in [m] , \quad \text{and} \quad (ii)\ x^{N,0}_{m+j} = 0 \text{ for } j \in [n-m] , \tag{B.24}$$

and the positive components of $x^{N,0}$ correspond exactly to the optimal basis. Therefore, $N_1 \le N_0 + 1$ and T_1 is bounded above by the number of OnePDHG iterations required to obtain $z^{N_0+1,0}$.

According to Remark A.1, the number of OnePDHG iterations needed to obtain such a $z^{N_0,0}$ is upper bounded by the number of iterations $\tilde{T}$ required to obtain $z^{\tilde{N}_0,0}$ such that $\rho(\|z^{\tilde{N}_0-1,0} - z^{\tilde{N}_0,0}\|_M; z^{\tilde{N}_0,0}) \le \frac{\bar{\varepsilon}}{16.5\kappa\hat{\Phi}}$, and

$$\tilde{T} \le 380\kappa\hat{\Phi} \cdot \ln\left(\frac{4560\kappa^2\hat{\Phi}\|w^\star\|}{\xi}\right) . \tag{B.25}$$

Furthermore, Lemma A.3 implies that, with step-sizes of Theorem 3.1, for all n:

$$\|z^{n,0} - z^\star\|_M \le 2\sqrt{\kappa}\|w^{n,0} - w^\star\| \le (6\sqrt{2}+8)\kappa\hat{\Phi}\cdot\rho(\|z^{n-1,0} - z^{n,0}\|_M; z^{n,0}) \ . \tag{B.26}$$

Lemma A.2 guarantees that the number of additional OnePDHG iterations before obtaining the next outer loop iteration $z^{\tilde{N}_0+1,0}$ is at most $\left\lceil \frac{8\cdot(6\sqrt{2}+8)\kappa\hat{\Phi}}{\beta} \right\rceil$. Overall, since $\tilde{N}_0 \ge N_0$, the number of OnePDHG iterations required before obtaining $z^{\tilde{N}_0+1,0}$ is at most:

$$\tilde{T} + \left\lceil \frac{8\cdot(6\sqrt{2}+8)\kappa\hat{\Phi}}{\beta} \right\rceil \overset{\text{(B.25)}}{\le} 380\kappa\hat{\Phi}\cdot\ln\left(\frac{4560\kappa^2\hat{\Phi}\|w^\star\|}{\xi}\right) + \left\lceil \frac{8\cdot(6\sqrt{2}+8)\kappa\hat{\Phi}}{\beta} \right\rceil ,$$

which we use T_{basis} to denote. Because Φ is equivalent to $\hat{\Phi}$ as demonstrated by Lemma 3.2 and $\frac{\|w^\star\|}{\xi} \le \Phi$ from the definition, T_{basis} reduces to $O(\kappa\Phi\ln(\kappa\Phi))$ in (4.1).

Next, we prove (4.2). We first study how large $(N_2 - \tilde{N}_0 - 1)$ could be. We mainly consider the case $N_2 > \tilde{N}_0 + 1$; otherwise $N_2 \le \tilde{N}_0 + 1$ and $T_2 \le T_{basis}$. In this case $N_2 > \tilde{N}_0 + 1$, because N_2 is the first iteration such that $\|w^{N_2,0} - w^\star\| \le \varepsilon$, for the previous iteration we have $\|w^{N_2-1,0} - w^\star\| > \varepsilon$. By (B.26), this implies:

$$\rho\left(\|z^{N_2-1,0} - z^{N_2-2,0}\|_M; z^{N_2-1,0}\right) > \frac{\sqrt{\kappa}\varepsilon}{(3\sqrt{2}+4)\kappa\hat{\Phi}} \ . \tag{B.27}$$

On the other hand, as shown in the definition of $\tilde{N}_0$,

$$\rho\left(\|z^{\tilde{N}_0,0} - z^{\tilde{N}_0+1,0}\|_M; z^{\tilde{N}_0+1,0}\right) \le \beta\cdot\frac{\bar{\varepsilon}}{16.5\kappa\hat{\Phi}} = \frac{1}{e}\cdot\frac{1}{16.5\kappa\hat{\Phi}}\cdot\frac{\xi}{6\sqrt{\kappa}} \le \frac{\xi}{99e\cdot\kappa^{1.5}\hat{\Phi}} \le \frac{\xi}{269\cdot\kappa^{1.5}\hat{\Phi}} \ . \tag{B.28}$$

Furthermore, due to the restart condition:

$$\rho\left(\|z^{N_2-1,0} - z^{N_2-2,0}\|_M; z^{N_2-1,0}\right) \le \beta^{N_2-\tilde{N}_0-2}\cdot\rho(\|z^{\tilde{N}_0,0} - z^{\tilde{N}_0+1,0}\|_M; z^{\tilde{N}_0+1,0}) \ . \tag{B.29}$$

Substituting (B.27) and (B.28) into (B.29) yields an upper bound of $(N_2 - \tilde{N}_0 - 1)$:

$$\frac{\sqrt{\kappa}\varepsilon}{(3\sqrt{2}+4)\kappa\hat{\Phi}} < \left(\frac{1}{e}\right)^{N_2-\tilde{N}_0-2}\cdot\frac{\xi}{269\cdot\kappa^{1.5}\hat{\Phi}} \Rightarrow N_2 - \tilde{N}_0 - 1 \le \ln\left(\frac{e\cdot(3\sqrt{2}+4)\kappa\hat{\Phi}}{269\kappa^2\hat{\Phi}}\right) + \ln\left(\frac{\xi}{\varepsilon}\right) \le \ln\left(\frac{\xi}{\varepsilon}\right) \ . \tag{B.30}$$

Here the final inequality is due to $\kappa \ge 1$ and $\frac{e\cdot(3\sqrt{2}+4)\kappa\hat{\Phi}}{269\kappa^2\hat{\Phi}} \le 1$. Since we have assumed $N_2 > \tilde{N}_0 + 1$, the upper bound of $(N_2 - \tilde{N}_0 - 1)$ in (B.30) has to be strictly positive. Once $\ln\left(\frac{\xi}{\varepsilon}\right) \le 0$, then it is no longer in the case $N_2 > \tilde{N}_0 + 1$ and as previously stated before we already have $N_2 \le \tilde{N}_0 + 1$ and $T_2 \le T_{basis}$. In conclusion, we can assert that $N_2 - \tilde{N}_0 - 1 \le \max\left\{0, \ln\left(\frac{\xi}{\varepsilon}\right)\right\}$.

We now turn our attention to the number of OnePDHG iterations between $z^{N_0+1,0}$ and $z^{N_2,0}$. Remark 4.1 ensures that for all $N \ge \tilde{N}_0 + 1$:

$$\|z^{N,0} - z^\star\|_M \le 4\|B^{-1}\|\|A\|\cdot\rho\left(\|z^{N,0} - z^{N-1,0}\|_M; z^{N,0}\right) \ . \tag{B.31}$$

Then due to Lemma A.2, the number of iterations in inner loops is at most $\left\lceil \frac{8\cdot4\|B^{-1}\|\|A\|}{1/e} \right\rceil$. Therefore, we can bound the overall number of iterations T_2 as follows:

$$T_2 \le T_{basis} + (N_2 - \tilde{N}_0 - 1)\cdot\left\lceil \frac{8\cdot4\|B^{-1}\|\|A\|}{1/e} \right\rceil \tag{B.32}$$

Substituting the upper bound $N_2 - \tilde{N}_0 - 1 \le \max\left\{0, \ln\left(\frac{\xi}{\varepsilon}\right)\right\}$ into (B.32) finishes the proof. □

Appendix C: Appendix of Section 7

Proof of Theorem 7.1. Applying the definition of Φ to the reweighted problem (7.1) and using Lemma 5.2 yield the formula of Φ_{ω_1,ω_2}:

$$\Phi_{\omega_1,\omega_2} = \left(\omega_2\|x^\star\|_1 + \omega_1\|s^\star\|_1\right) \cdot \max\left\{\frac{1}{\omega_1\zeta_p}, \frac{1}{\omega_2\zeta_d}\right\}. \tag{C.1}$$

Using the condition $\frac{\hat{\omega}_1}{\hat{\omega}_2} = \frac{\|x^\star\|_1}{\|s^\star\|_1}$, $\Phi_{\hat{\omega}_1,\hat{\omega}_2} = 2 \cdot \max\left\{\frac{\|x^\star\|_1}{\zeta_d}, \frac{\|s^\star\|_1}{\zeta_p}\right\}$, which is exactly the first equality of (7.2). On the other hand, (C.1) implies that $\Phi_{\omega_1,\omega_2} \geq \max\left\{\frac{\|x^\star\|_1}{\zeta_d}, \frac{\|s^\star\|_1}{\zeta_p}\right\}$ for all ω_1, ω_2. This proves the inequality of (7.2). □

Remark 3.4 of Xiong and Freund [65] proposes the "optimized" step-sizes

$$\tau = \frac{\mu_d \cdot \text{Dist}(0, V_p)}{2\kappa\mu_p \cdot \text{Dist}(0, V_d)} \quad \text{and} \quad \sigma = \frac{\mu_p \cdot \text{Dist}(0, V_d)}{2\lambda_{\max}\lambda_{\min}\mu_d \cdot \text{Dist}(0, V_p)} \tag{C.2}$$

where μ_p and μ_d denote the LP sharpness of the primal and dual problems, as defined in Definition 5.3. As demonstrated by Applegate et al. [6], setting the step-size ratio (C.2) is equivalent to setting primal weights $(\breve{\omega}_1, \breve{\omega}_2)$ such that $\frac{\breve{\omega}_1}{\breve{\omega}_2} = \frac{\mu_d \cdot \text{Dist}(0,V_p)}{\mu_p \cdot \text{Dist}(0,V_d)}$ with the standard step-sizes τ, σ of Theorem 3.1. Furthermore, due to Lemmas 5.1 and 5.2, $\frac{\breve{\omega}_1}{\breve{\omega}_2} = \frac{\zeta_d}{\zeta_p}$. Now, substituting $(\omega_1, \omega_2) = (\breve{\omega}_1, \breve{\omega}_2)$ into (C.1) yields $\Phi_{\breve{\omega}_1,\breve{\omega}_2} = \frac{\|x^\star\|_1}{\zeta_d} + \frac{\|s^\star\|_1}{\zeta_p}$. This is equivalent to $\Phi_{\hat{\omega}_1,\hat{\omega}_2}$ (when $\frac{\hat{\omega}_1}{\hat{\omega}_2} = \frac{\|x^\star\|_1}{\|s^\star\|_1}$) up to a constant of 2, so Remark 3.4 of Xiong and Freund [65] is also approximately optimizing the geometric measure $\hat{\Phi}_{\omega_1,\omega_2}$ via reweighting.

Acknowledgments

The author is grateful to Robert M. Freund for the valuable suggestions while preparing this manuscript. The author also thanks Yinyu Ye, Haihao Lu, Tianyi Lin, Javier Pena and Xin Jiang for the helpful discussions. The author also thanks the anonymous referees and editors for helpful comments and suggestions that improved the quality of the manuscript.

References

[1] Alizadeh F, Haeberly JPA, Overton ML (1998) Primal-dual interior-point methods for semidefinite programming: Convergence rates, stability and numerical results. *SIAM Journal on Optimization* 8(3):746–768.

[2] Anstreicher KM, Ji J, Potra FA, Ye Y (1993) Average performance of a self–dual interior point algorithm for linear programming. *Complexity in Numerical Optimization*, 1–15 (World Scientific).

[3] Anstreicher KM, Ji J, Potra FA, Ye Y (1999) Probabilistic analysis of an infeasible-interior-point algorithm for linear programming. *Mathematics of Operations Research* 24(1):176–192.

[4] Applegate D, Díaz M, Hinder O, Lu H, Lubin M, O'Donoghue B, Schudy W (2021) Practical large-scale linear programming using primal-dual hybrid gradient. *Proceedings of the Advances in Neural Information Processing Systems 34*, volume 34, 20243–20257.

[5] Applegate D, Díaz M, Lu H, Lubin M (2024) Infeasibility detection with primal-dual hybrid gradient for large-scale linear programming. *SIAM Journal on Optimization* 34(1):459–484.

[6] Applegate D, Hinder O, Lu H, Lubin M (2023) Faster first-order primal-dual methods for linear programming using restarts and sharpness. *Mathematical Programming* 201(1-2):133–184.

[7] Basu K, Ghoting A, Mazumder R, Pan Y (2020) Eclipse: An extreme-scale linear program solver for web-applications. *Proceedings of the 37th International Conference on Machine Learning*, 704–714 (PMLR).

[8] Bertsimas D, Tsitsiklis JN (1997) *Introduction to linear optimization*, volume 6 (Athena scientific Belmont, MA).

[9] Borgwardt KH (1987) *The simplex method: A probabilistic analysis*, volume 1 (Springer Nature).

[10] Bowman EH (1956) Production scheduling by the transportation method of linear programming. *Operations Research* 4(1):100–103.

[11] Cardinal Operations LLC (2025) *Logging — Cardinal Optimizer (COPT) User Guide, Ver 8.0*. URL `https://guide.coap.online/copt/en-doc/logging.html#first-order-method-pdlp-logging-for-gpu-solver`, Accessed: 2025-12-16.

[12] Chambolle A, Pock T (2011) A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision* 40:120–145.

[13] Charnes A, Cooper WW (1954) The stepping stone method of explaining linear programming calculations in transportation problems. *Management Science* 1(1):49–69.

[14] Chen K, Sun D, Yuan Y, Zhang G, Zhao X (2025) HPR-LP: An implementation of an HPR method for solving linear programming: K. chen et al. *Mathematical Programming Computation* 1–28.

[15] Cormen TH, Leiserson CE, Rivest RL, Stein C (2022) *Introduction to algorithms* (MIT Press), 4th edition.

[16] Dantzig GB (2002) Linear programming. *Operations Research* 50(1):42–47.

[17] De Rosa A, Khajavirad A (2024) On the power of linear programming for K-means clustering. *arXiv preprint arXiv:2402.01061* .

[18] Deng Q, Feng Q, Gao W, Ge D, Jiang B, Jiang Y, Liu J, Liu T, Xue C, Ye Y, Zhang C (2025) An enhanced alternating direction method of multipliers-based interior point method for linear and conic optimization. *INFORMS Journal on Computing* 37(2):338–359.

[19] Drusvyatskiy D, Lewis AS (2011) Generic nondegeneracy in convex optimization. *Proceedings of the American Mathematical Society* 2519–2527.

[20] Esser E, Zhang X, Chan TF (2010) A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM Journal on Imaging Sciences* 3(4):1015–1046.

[21] FICO (2025) FICO Xpress optimizer documentation: Solution methods (hybrid gradient method). URL `https://www.fico.com/fico-xpress-optimization/docs/latest/solver/optimizer/HTML/chapter4.html`, Accessed: 2025-12-16.

[22] Freund RM (2003) On the primal-dual geometry of level sets in linear and conic optimization. *SIAM Journal on Optimization* 13(4):1004–1013.

[23] Freund RM, Vera JR (1999) Condition-based complexity of convex optimization in conic linear form via the ellipsoid algorithm. *SIAM Journal on Optimization* 10(1):155–176.

[24] Ge D, Hu H, Huangfu Q, Liu J, Liu T, Lu H, Yang J, Ye Y, Zhang C (2024) cuPDLP-C. `https://github.com/COPT-Public/cuPDLP-C`. Accessed: 2024-09-11.

[25] Gleixner A, Gottwald L, Hoen A (2023) PaPILO: A parallel presolving library for integer and linear optimization with multiprecision support. *INFORMS Journal on Computing* 35(6):1329–1341.

[26] Gleixner A, Hendel G, Gamrath G, Achterberg T, Bastubbe M, Berthold T, Christophel P, Jarck K, Koch T, Linderoth J, Lübbecke M, Mittelmann HD, Ozyurt D, Ralphs TK, Salvagnin D, Shinano Y (2021) MIPLIB 2017: data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation* 13(3):443–490.

[27] Greene WH (2017) *Econometric analysis* (Pearson), 7th edition.

[28] Güler O, Ye Y (1993) Convergence behavior of interior-point algorithms. *Mathematical Programming* 60(1-3):215–228.

[29] Gurobi Optimization, LLC (2025) Gurobi 13.0 FAQ: Primal-Dual Hybrid Gradient (PDHG) for very large linear programs. URL `https://www.gurobi.com/faqs/gurobi-13-0/`, Accessed: 2025-12-16.

[30] Hanssmann F, Hess SW (1960) A linear programming approach to production and employment scheduling. *Management Science* 1(1):46–51.

[31] Hinder O (2023) Worst-case analysis of restarted primal-dual hybrid gradient on totally unimodular linear programs. *arXiv preprint arXiv:2309.03988* .

[32] Hough M, Vavasis SA (2024) A primal-dual Frank-Wolfe algorithm for linear programming. *arXiv preprint arXiv:2402.18514* .

[33] Huang Y, Zhang W, Li H, Ge D, Liu H, Ye Y (2025) A restarted primal-dual hybrid conjugate gradient method for large-scale quadratic programming. *INFORMS Journal on Computing* .

[34] Koch T, Berthold T, Pedersen J, Vanaret C (2022) Progress in mathematical programming solvers from 2001 to 2020. *EURO Journal on Computational Optimization* 10:100031.

[35] Li B, Yang L, Chen Y, Wang S, Mao H, Chen Q, Ma Y, Wang A, Ding T, Tang J, Sun R (2024) Pdhg-unrolled learning-to-optimize method for large-scale linear programming. *International Conference on Machine Learning*, 29164–29180 (PMLR).

[36] Lin T, Ma S, Ye Y, Zhang S (2021) An ADMM-based interior-point method for large-scale linear programming. *Optimization Methods and Software* 36(2-3):389–424.

[37] Lin Z, Xiong Z, Ge D, Ye Y (2025) PDCS: A primal-dual large-scale conic programming solver with GPU enhancements. *arXiv preprint arXiv:2505.00311* .

[38] Liu H, Huang Y, Li H, Ge D, Ye Y (2025) PDHCG: A scalable first-order method for large-scale competitive market equilibrium computation. *arXiv preprint arXiv:2506.06258* .

[39] Liu XW, Dai YH, Huang YK (2022) A primal-dual majorization-minimization method for large-scale linear programs. *arXiv preprint arXiv:2208.03672* .

[40] Lu H, Applegate D (2024) Scaling up linear programming with PDLP. `https://research.google/blog/scaling-up-linear-programming-with-pdlp/`. Accessed: 2024-09-26.

[41] Lu H, Peng Z, Yang J (2025) cuPDLPx: A further enhanced GPU-based first-order solver for linear programming. *arXiv preprint arXiv:2507.14051* .

[42] Lu H, Simester D, Zhu Y (2025) Optimizing scalable targeted marketing policies with constraints. *Marketing Science* .

[43] Lu H, Yang J (2022) On the infimal sub-differential size of primal-dual hybrid gradient method. *arXiv preprint arXiv:2206.12061* .

[44] Lu H, Yang J (2024) PDOT: A practical primal-dual algorithm and a GPU-based solver for optimal transport. *arXiv preprint arXiv:2407.19689* .

[45] Lu H, Yang J (2024) Restarted Halpern PDHG for linear programming. *arXiv preprint arXiv:2407.16144* .

[46] Lu H, Yang J (2025) cupdlp. jl: A gpu implementation of restarted primal-dual hybrid gradient for linear programming in julia. *Operations Research* 73(6):3440–3452.

[47] Lu H, Yang J (2025) On the geometry and refined rate of primal–dual hybrid gradient for linear programming: On the geometry and refined rate of primal–dual... *Mathematical Programming* 212(1):349–387.

[48] Lu H, Yang J (2025) A practical and optimal first-order method for large-scale convex quadratic programming. *Mathematical Programming* 1–38.

[49] Lu H, Yang J (2025) A practical and optimal first-order method for large-scale convex quadratic programming: H. lu, j. yang. *Mathematical Programming* 1–38.

[50] Lu H, Yang J, Hu H, Huangfu Q, Liu J, Liu T, Ye Y, Zhang C, Ge D (2023) cuPDLP-C: A strengthened implementation of cuPDLP for linear programming by C language. *arXiv preprint arXiv:2312.14832* .

[51] Mehrotra S, Ye Y (1993) Finding an interior point in the optimal face of linear programs. *Mathematical Programming* 62(1-3):497–515.

[52] Mirrokni V (2023) 2022 & beyond: Algorithmic advances. `https://ai.googleblog.com/2023/02/google-research-2022-beyond-algorithmic.html`. Accessed: 2024-09-11.

[53] NVIDIA Corporation (2025) LP features — NVIDIA cuOpt user guide (25.12). URL `https://docs.nvidia.com/cuopt/user-guide/latest/lp-features.html`, Accessed: 2025-12-16.

[54] O'Donoghue B (2021) Operator splitting for a homogeneous embedding of the linear complementarity problem. *SIAM Journal on Optimization* 31(3):1999–2023.

[55] O'Donoghue B, Chu E, Parikh N, Boyd S (2016) Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications* 169:1042–1068.

[56] Pena J, Vera JC, Zuluaga LF (2021) New characterizations of Hoffman constants for systems of linear constraints. *Mathematical Programming* 187:79–109.

[57] Pock T, Cremers D, Bischof H, Chambolle A (2009) An algorithm for minimizing the Mumford–Shah functional. *Proceedings of the 2009 International Conference on Computer Vision*, 1133–1140 (IEEE).

[58] Potra FA (1994) A quadratically convergent predictor—corrector method for solving linear programs from infeasible starting points. *Mathematical Programming* 67(1):383–406.

[59] Todd MJ (1991) Probabilistic models for linear programming. *Mathematics of Operations Research* 16(4):671–693.

[60] Todd MJ, Ye Y (1990) A centered projective algorithm for linear programming. *Mathematics of Operations Research* 15(3):508–529.

[61] Vavasis SA, Ye Y (1996) A primal-dual interior point method whose running time depends only on the constraint matrix. *Mathematical Programming* 74(1):79–120.

[62] Wagner M, Meller J, Elber R (2004) Large-scale linear programming techniques for the design of protein folding potentials. *Mathematical Programming* 101:301–318.

[63] Wang H, Ghosal P, Mazumder R (2023) Linear programming using diagonal linear networks. *arXiv preprint arXiv:2310.02535* .

[64] Xiong Z (2025) High-probability polynomial-time complexity of restarted PDHG for linear programming. *arXiv preprint arXiv:2501.00728* .

[65] Xiong Z, Freund RM (2023) Computational guarantees for restarted PDHG for LP based on "limiting error ratios" and LP sharpness. *arXiv preprint arXiv:2312.14774* .

[66] Xiong Z, Freund RM (2023) On the relation between LP sharpness and limiting error ratio and complexity implications for restarted PDHG. *arXiv preprint arXiv:2312.13773* .

[67] Xiong Z, Freund RM (2024) The role of level-set geometry on the performance of PDHG for conic linear optimization. *arXiv preprint arXiv:2406.01942* .

[68] Yang T, Lin Q (2018) Rsg: Beating subgradient method without smoothness and strong convexity. *The Journal of Machine Learning Research* 19(1):236–268.

[69] Ye Y (1992) On the finite convergence of interior-point algorithms for linear programming. *Mathematical Programming* 57(1):325–335.

[70] Ye Y (1994) Toward probabilistic analysis of interior-point algorithms for linear programming. *Mathematics of Operations Research* 19(1):38–52.

[71] Ye Y (2011) The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate. *Mathematics of Operations Research* 36(4):593–603.