

Infeasibility detection with primal-dual hybrid gradient for large-scale linear programming

David Applegate* Mateo Díaz† Haihao Lu‡ Miles Lubin§

Abstract

We study the problem of detecting infeasibility of large-scale linear programming problems using the primal-dual hybrid gradient method (PDHG) of Chambolle and Pock (2011). The literature on PDHG has mostly focused on settings where the problem at hand is assumed to be feasible. When the problem is not feasible, the iterates of the algorithm do not converge. In this scenario, we show that the iterates diverge at a controlled rate towards a well-defined ray. The direction of this ray is known as the infimal displacement vector v . The first contribution of our work is to prove that this vector recovers certificates of primal and dual infeasibility whenever they exist. Based on this fact, we propose a simple way to extract approximate infeasibility certificates from the iterates of PDHG. We study three different sequences that converge to the infimal displacement vector: the difference of iterates, the normalized iterates, and the normalized average. All of them are easy to compute, and thus the approach is suitable for large-scale problems. Our second contribution is to establish tight convergence rates for these sequences. We demonstrate that the normalized iterates and the normalized average achieve a convergence rate of $O\left(\frac{1}{k}\right)$, improving over the known rate of $O\left(\frac{1}{\sqrt{k}}\right)$. This rate is general and applies to any fixed-point iteration of a nonexpansive operator. Thus, it is a result of independent interest since it covers a broad family of algorithms, including, for example, ADMM, and can be applied settings beyond linear programming, such as quadratic and semidefinite programming. Further, in the case of linear programming we show that, under nondegeneracy assumptions, the iterates of PDHG identify the active set of an auxiliary feasible problem in finite time, which ensures that the difference of iterates exhibits eventual linear convergence to the infimal displacement vector. Numerical experiments on *Netlib* dataset of infeasible LP instances verify our theoretical results.

1 Introduction

First-order methods (FOMs) have been extensively studied by the optimization community since the late 2000s, following a long period where interior-point methods dominated research in continuous optimization. FOMs, which use only gradient information, are appealing for their simplicity and low computational overhead, in particular when solving large-scale optimization problems that arise in machine learning and data science. These methods have matured in many aspects (see, e.g., the recent textbook by Beck [10]) and are known to be useful for obtaining moderately accurate solutions to convex and non-convex optimization problems in a reasonable amount of time. Despite this progress, FOMs have made only modest inroads into linear programming (LP), a fundamental problem in mathematical optimization.

FOMs applied to LP provide relatively simple methods whose most expensive operations are matrix-vector multiplications with the (typically sparse) constraint matrix. Such matrix-vector products are amenable to scale efficiently given increasingly ubiquitous computing resources like multi-threaded CPUs, GPUs [63], or distributed clusters [24]. In contrast, interior-point and simplex-based methods that dominate current practice are limited in how they use available computing resources because they depend on matrix inversion. To mark

*Google Research (dapplegate@google.com)

†Cornell University Center for Applied Mathematics (md825@cornell.edu). Work done while interning at Google Research.

‡University of Chicago Booth School of Business and Google Research (haihao.lu@chicagobooth.edu)

§Google Research (mlubin@google.com)

Algorithm 1: Primal-dual hybrid gradient**Data:** $x_0 \in \mathbf{R}^d$ **Step k :** ($k \geq 0$)Update $x^{k+1} \leftarrow \mathbf{prox}_{\eta f}(x^k - \eta A^\top y^k)$,Update $y^{k+1} \leftarrow \mathbf{prox}_{\tau h}(y^k + \tau A(2x^{k+1} - x^k))$.

this distinction, Nesterov [50] defines methods that use at most matrix-vector products as capable of handling *large-scale* problems and methods that use matrix inversion as handling *medium-scale* problems. In the context of LP, these definitions of scale perhaps belie the reliability and practical efficiency of interior-point and simplex methods, but nevertheless the contrast in the computing requirements of the algorithms is an important one. Even though such computational aspects are outside the scope of this paper, it is this practical potential to efficiently solve large-scale LP that motivates the theoretical developments in this work.

While FOMs are typically studied in more general settings, the underlying assumptions and convergence rates in these settings do not necessarily hold or may not be tight for the special case of LP. Of particular relevance to this work, theory for FOMs is often developed under the assumption that an optimal solution exists, whereas LP solvers need to be able to detect infeasibility (i.e., when no optimal solution exists) and compute corresponding certificates. Infeasibility detection and computation of certificates are an essential aspect of solving LP, not only to provide feedback on modeling errors but also for algorithms that directly exploit LP certificates like Benders decomposition and branch-and-cut [1].

This work addresses the question of how to detect infeasibility in LP using the *Primal-Dual Hybrid Gradient method* (PDHG). PDHG is a popular first-order method introduced by Chambolle and Pock [15] to solve *convex-concave minimax problems*, that is, problems of the form

$$\min_{x \in \mathbf{R}^n} \max_{y \in \mathbf{R}^m} \langle Ax, y \rangle + g(x) - h(y) \quad (1)$$

where $g : \mathbf{R}^n \rightarrow \mathbf{R} \cup \{\infty\}$ and $h : \mathbf{R}^m \rightarrow \mathbf{R} \cup \{\infty\}$ are proper lower semicontinuous convex functions and $A \in \mathbf{R}^{m \times n}$. LP can be recast as a minimax problem through duality, and hence PDHG is applicable. The method consists of alternating updates between the primal and dual variables, see Algorithm 1, where \mathbf{prox} is the proximal operator (see the definition in (4)). In particular, when instantiated for LP, these updates correspond to matrix-vector products and projections onto simple sets (such as the positive orthant). In contrast with other methods like the *Alternating Direction Method of Multipliers* (ADMM), PDHG does not require projections onto linear subspaces, which involve matrix inversions by direct or indirect methods.

The behavior of PDHG for *feasible* problems (i.e., problems that have an optimal solution) has been studied in depth under several regularity assumptions. In their seminal work, Chambolle and Pock [15] show that the algorithm converges at a rate of $O(1/k)$ given appropriate choices for the step sizes η and τ . However, the situation for infeasible problems remains largely unstudied.

While it is relatively straightforward to formulate always-feasible auxiliary problems that can be used to detect infeasibility, for example, by penalizing violations of primal and dual constraints, this approach is unappealing for two reasons: First is the aesthetic interest of having a single algorithm that robustly handles all possible input [32]. Second is the practical interest in effectively using available computing resources, as solving such auxiliary problems would approximately double the necessary work. Instead, we aim to use *one* execution of PDHG and ask the following question:

Do the PDHG iterates encode information about infeasibility?

We answer this question in the affirmative. We show that if the primal (and/or dual) problem is infeasible, the iterates of PDHG recover primal (and/or dual) infeasibility certificates. Moreover, we completely characterize the behavior of the iterates under different infeasibility settings. Before diving into our main contributions, let us present an illustrative example. Recall that for a primal-dual LP pair, there exist three exhaustive and mutually exclusive possibilities: (1) both primal and dual are feasible, (2) both primal and dual are

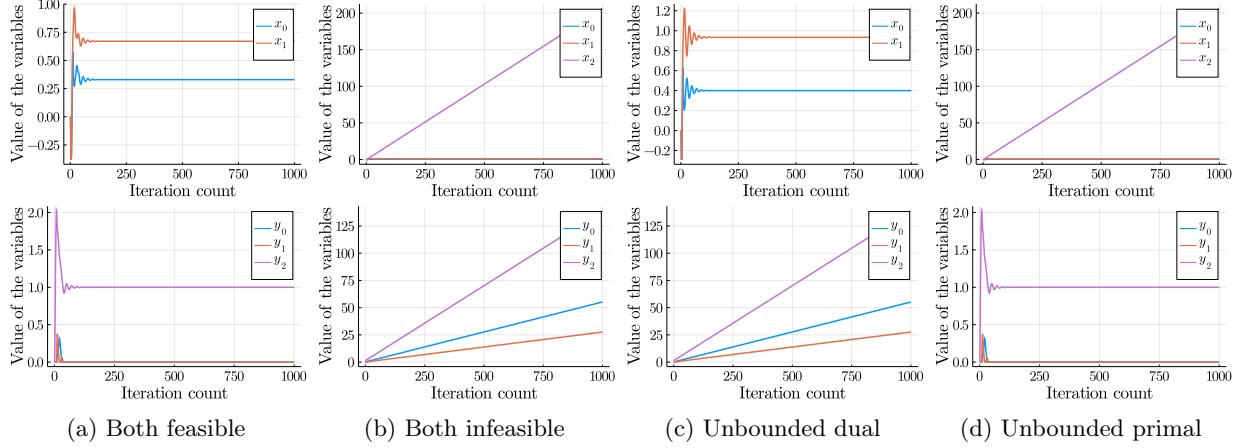


Figure 1: Four different settings depicted in Example 1. Every subplot shows the component-wise value of the iterates against the iteration count. The first and the second rows correspond to the primal and dual iterates, respectively.

infeasible, and (3) one of the two problems is unbounded, and consequently, the other problem is infeasible. Small numerical experiments reveal that the behavior of PDHG is different depending on the setting.

Example 1. Consider the LP problem with constants $\alpha, \beta \in \mathbf{R}$:

$$\begin{aligned}
 & \text{minimize} && x_0 + x_1 - \alpha x_2 \\
 & \text{subject to} && x_0 + 2x_1 \leq 2 \\
 & && 3x_0 + x_1 \leq 2 \\
 & && x_0 + x_1 \geq \beta.
 \end{aligned}$$

Figure 1 displays four choices of α and β

1. **Both feasible.** Set $\alpha = 0$ and $\beta = 1$, then both primal and dual problems are feasible. In this case, both the primal and dual variables converge to a solution.
2. **Both infeasible.** Set $\alpha = 1$ and $\beta = 2$, then both primal and dual are infeasible. We observe that both primal and dual iterates diverge at a rate proportional to the number of iterations.
3. **Unbounded dual.** Set $\alpha = 0$ and $\beta = 2$, then the dual problem is unbounded and, thus, the primal problem is infeasible and the dual is feasible. Then the dual iterates diverge, and, interestingly, the primal iterates converge.
4. **Unbounded primal.** Set $\alpha = 1$ and $\beta = 1$, then the primal problem is unbounded and, thus, the dual problem is infeasible and the primal is feasible. Then the primal iterates diverge, and the dual iterates converge.

From the experiments, we see that the iterates have a very stable asymptotic behavior. In particular, if the primal is feasible, then the dual variables converge, and analogously if the dual is feasible, then the primal iterates converge. Similarly, whenever the primal is infeasible, the dual iterates diverge at a controlled linear rate and vice-versa. Such behavior has not been previously observed or characterized in the literature.

1.1 Main contributions

For notational convenience, we use $z = (x, y)$ as the primal-dual pair, and $\bar{z}^k := \frac{1}{k} \sum_{j=1}^k z^j$ as the average of iterates. We propose to detect infeasibility using three sequences:

$$\text{(Difference of iterates)} \quad d_k = (z^{k+1} - z^k), \quad (2a)$$

$$\text{(Normalized iterates)} \quad \frac{z^k}{k}, \quad (2b)$$

$$\text{(Normalized average iterates)} \quad \frac{2}{k+1} \bar{z}^k. \quad (2c)$$

Our proposal to detect infeasibility is as follows:

Use these three sequences' primal and dual components as candidates for dual and primal infeasibility certificates. The algorithm should periodically check if any of these iterates satisfy the conditions that define an infeasibility certificate within numerical tolerances. If at any point this happens, it should conclude that the problem is (primal or dual) infeasible.

The overhead cost of extracting the certificates is negligible, making it suitable for large-scale problems. Most of the content of this work justifies this strategy theoretically.

Operator theory shows that all three of these sequences converge to a point v known as the *infimal displacement vector*. Section 2 will give a formal definition of this and other relevant concepts. We list our contributions assuming, for now, the existence of such a vector v .

Sublinear convergence rate to the infimal displacement vector (Section 3). It is natural to wonder how fast the three sequences (2) converge to the infimal displacement vector. We study this question through the lenses of general operators and fixed-point iteration. To the best of our knowledge, the only known result in this vein ensures a rate of $O\left(\frac{1}{\sqrt{k}}\right)$ for the difference of iterates (2a), which is known to be tight [39, 19]. In contrast, we show that two other sequences, the normalized iterates (2b) and the normalized average iterates (2c), converge at a faster rate of $O\left(\frac{1}{k}\right)$ in this same situation. Furthermore, this faster sublinear result generalizes to *any fixed-point iteration of a nonexpansive operator*, not only the firmly nonexpansive operators studied in [39]. Specifically, it also applies to many popular first-order methods, including but not limited to PDHG, ADMM [56], and Mirror-prox [49], and it extends to other settings beyond LP such as quadratic convex programming and semidefinite programming. Furthermore, we show that this result is tight for PDHG; i.e., there exist instances with a convergence rate lower bounded by $\Omega\left(\frac{1}{k}\right)$. This result suggests that current ADMM-based codes like OSQP [61] that use exclusively the difference of iterates to detect infeasibility should additionally consider the normalized iterates and normalized average iterates.

Characterization of the iterates for infeasible problems (Section 4). We characterize the behavior of PDHG for all the LP feasibility scenarios (see Table 1). In particular, we show that if the primal (or dual) iterates diverge, then the iterates diverge in the direction of a ray, where the direction of the ray recovers certificates of dual (or primal) infeasibility. Such direction turns out to be the infimal displacement vector (v_x, v_y) . This justifies using the sequences (2) as infeasibility certificate candidates. Furthermore, we show that when the primal problem is feasible, then the dual iterates, without any normalization, converge to some y^* that is closely related to v . An analogous result holds if the dual is feasible. This describes the dynamics of PDHG for unbounded problems. The next table summarizes our findings for the four possible cases.

Dual \ Primal	Feasible	Infeasible
Feasible	x^k, y^k both converge	x^k diverges, y^k converges
Infeasible	x^k converges, y^k diverges	x^k, y^k both diverge

Table 1: Behavior of PDHG for solving under different feasibility assumptions.

Eventual linear convergence for nondegenerate problems (Section 5). In the process of characterizing the dynamics of PDHG, we show that the iterates (x^k, y^k) always converge to a unique ray $\{(x^*, y^*) + \lambda v \mid \lambda \in \mathbf{R}_+\}$. We show that under a non-degeneracy condition (a direct extension of the strict complementary condition to the infeasible LPs), the iterates (x^k, y^k) fix their active set after finitely many iterations. In turn, this leads to the eventual linear convergence of the difference of iterates (2a). Formally, we show that there exists $K \geq 0$ such that for all sufficiently large $k \geq K$ we have

$$\|d^k - v\| \leq O(\gamma^{k-K}) \quad \text{for some } \gamma \in (0, 1) .$$

We further show that even after the active set is fixed, the normalized iterates and normalized average do not exhibit faster convergence. Thus, it is strictly better to use the difference of iterates to detect infeasibility in this regime.

Computational experiments (Section 6). We verify our theoretical results by presenting numerical experiments displaying the efficacy of the different certificate candidates (2). Specifically, our experiments show that using all three sequences in (2) is beneficial. On the one hand, if the active set’s finite time identification occurs relatively quickly, then the differences of iterates (2a) exhibit faster convergence. On the other hand, for some problems, identifying the active set might not happen in a reasonable amount of time. In this case, both the normalized iterates (2b) recovers approximate infeasibility certificates much more efficiently than the differences.

1.2 Related work

Primal-dual hybrid gradient. Chambolle and Pock [16] review PDHG among other methods and describe its applications in computer vision. O’Connor and Vandenberghe [51] show that PDHG is in fact a particular application of Douglas-Rachford Splitting (DRS) [20, 30, 27, 38].

FOMs for LP. Lan et al. [33] and Renegar [58] develop FOMs for LP, considered as a special case of semidefinite programming, with $O(\frac{1}{k})$ convergence rates. Gilpin et al. [29] obtain a restarted FOM for LP with a linear convergence rate. These analyses assume an optimal solution exists. Pock and Chambolle [55] apply PDHG with diagonal preconditioning to LP on a small number of test instances. They note that on small-scale problems, interior-point methods clearly dominate, while their method outperforms MATLAB’s LP solver on one larger LP motivated by a computer vision application. Most recently, Basu et al. [5] apply accelerated gradient descent to a specialized LP instance, obtaining solutions to industrial problems with up to 10^{12} variables.

Infeasibility detection. Classically, the primal simplex method for LP detects primal infeasibility while solving a “phase-one” auxiliary problem for an initial feasible basis and detects dual infeasibility based on conditions when computing a step size (i.e., the ratio test) [42]. Infeasibility certificates are extracted from the iterates of interior-point methods without substantial extra work [62]. Infeasibility detection is only the first step of diagnosing the cause of the infeasibility in an LP model [17].

Most research on infeasibility detection capabilities for FOMs for convex optimization has focused on ADMM or equivalently Douglas-Rachford Splitting. Eckstein and Bertsekas [23] show that when no solution exists, then the iterates diverge. Recent practical successes motivated further research in this direction, characterizing the asymptotic behavior of the iterates under additional assumptions. For example, the line of work [7, 8, 48] studies Douglas-Rachford applied to problems that look for a point at the intersection of two non-intersecting convex sets. On the other hand, Raghunathan and Di Cairano [57] investigate the asymptotic dynamics of ADMM for convex quadratic problems when the matrices involved are full rank.

Banjak et al. [4] show that the infimal displacement vector of ADMM recovers certificates of infeasibility for convex quadratic problems with conic constraints. Based on this, they proposed to use the difference of iterates to test infeasible. Complementing this work, [39] establishes a $O(\frac{1}{\sqrt{k}})$ convergence rate for the difference of iterates of any algorithm that induces a firmly nonexpansive operator and introduced a scheme that utilizes multiple runs of ADMM to detect infeasibility. This type of scheme aims to handle pathological scenarios that do not occur in LP.

O’Donoghue et al. [53] propose to apply ADMM to a homogeneous self-dual embedding of a convex conic problem¹. A nice byproduct of this approach is that it automatically produces infeasibility certificates. Subsequent work [52] extends this approach to Linear Complementarity Problems, which cover quadratic convex losses with conic constraints.

To our knowledge, the only work analyzing the behavior of PDHG on potentially infeasible instances is by Malitsky [40], which considers linearly constrained problems. This analysis applies only to linear equality constraints, not to linear inequalities present in LP.

Finite time identifiability. Finite time identifiability has a long history in the field of optimization. This phenomenon is first documented for the projected gradient descent method [22, 14, 13, 12]. Soon after it is studied for other methods, such as the Proximal Point Method [25] and Projected subgradient descent [26], among others [2]. Identifiability is also exploited as tool for algorithmic design for the so called “ \mathcal{UV} -algorithms” [46]. Recent works [47, 36, 37] study finite time identification for popular FOMs. In particular, Liang et al. [37] show that the iterates of PDHG identify the active constraints in finite time, provided the limit point is nondegenerate. All of these works assume the underlying problem is feasible. The significant number of algorithms exhibiting this behavior motivated researchers to develop general theory (even beyond the realms of optimization) [64, 44, 45, 34, 21, 35]. We refer the interested reader to [35] for an elegant geometrical definition that generalizes most notions of nondegeneracy.

1.3 Outline

The paper is structured as follows. Section 2 presents all the necessary background. In Section 3, we show a convergence rate of $O(1/k)$ for the normalized iterates and normalized average generated by the fixed-point iteration of a nonexpansive operator. Then, Section 4 shows a complete characterization of the behavior of PDHG under different infeasibility assumptions. In Section 5, we study a condition that ensures finite time identifiability of the active set. We show that under this condition, the difference of iterates exhibits eventual linear convergence. We present numerical experiments that complement the theoretical results in Section 6. We close the paper with some conclusions and open questions in Section 7.

1.4 Notation

We use the symbols \mathbf{N} and \mathbf{R} to denote the natural numbers and reals. Our results take place in a finite dimensional spaces \mathbf{R}^d . We denote the cone of nonnegative vectors as \mathbf{R}_+^d . We endow \mathbf{R}^d with the standard dot product $\langle x, y \rangle = x^\top y$ and norm $\|x\|_2^2 := \langle x, x \rangle$. We use I to denote both the identity map and the identity matrix in $\mathbf{R}^{d \times d}$. Given a matrix $M \in \mathbf{R}^{d \times d}$ we use $\lambda_1(M), \dots, \lambda_d(M)$ to denote its eigenvalues. We define the spectral radius of a matrix M as $\rho(M) = \max_{j \in [d]} |\lambda_j(M)|$. Similarly, for a matrix $M \in \mathbf{R}^{m \times d}$, with $d \leq m$, we use $\sigma_1(M) \geq \dots \geq \sigma_d(M)$ for its singular values and we let $\sigma_{\max}(M)$ and $\sigma_{\min}(M)$ denote the maximum and minimum nonzero singular values. We define the operator norm as $\|M\|_2 = \sigma_{\max}(M)$. For any matrix M we denote its pseudo-inverse as M^\dagger . A symmetric matrix $M \in \mathbf{R}^{d \times d}$ is positive definite if its minimum eigenvalue is positive, $\min_{j \in [d]} \lambda_j(M) > 0$. Every positive definite matrix M defines an inner product and norm given by $\langle x, y \rangle_M = x^\top M y$ and $\|x\|_M^2 = \langle x, x \rangle_M$, respectively. Frequently, we will make use of the symbol $\|\cdot\|$ to denote an arbitrary matrix norm in \mathbf{R}^n . Often, $\|\cdot\|$ will denote a norm with respect to which an operator is (firmly) nonexpansive; see next section for a formal definition.

Let $C \subseteq \mathbf{R}^d$ be an arbitrary set. The symbol $\text{cl}(C)$ denotes the closure of C . We use $\text{proj}_C(\cdot)$ to denote the Euclidean projection mapping onto C ; this map is well-defined provided C is closed and convex. We use $(z)_+$ as a shorthand for $\text{proj}_{\mathbf{R}_+^d}(z)$. We define the indicator function of C as $\iota_C : \mathbf{R}^d \rightarrow \mathbf{R} \cup \{\infty\}$ given by $\iota_C(z) = 0$ if $z \in C$ and $\iota_C(z) = \infty$ otherwise. Given a map $T : \mathbf{R}^d \rightarrow \mathbf{R}^d$, its range is defined as $\text{range}(T) = \{T(z) \mid z \in \mathbf{R}^d\}$. Given two mappings $T_1, T_2 : \mathbf{R}^d \rightarrow \mathbf{R}^d$, we denote their composition as $T_1 \circ T_2$, that is $T_1 \circ T_2(z) = T_1(T_2(z))$. Since we study primal and dual problems, we use $z = (x, y) \in \mathbf{R}^{n+m}$ as a placeholder for primal and dual variables. We will sometimes refer to a vector $v \in \mathbf{R}^{n+m}$ and use v_x and v_y

¹Linear objective with conic constraints.

to denote its primal and dual components. We use superscripts to denote iteration counts, consequently z^k is the k th iterate. We use $\text{supp}(x) := \{i \in [d] \mid x_i \neq 0\}$ to denote the support of the vector x .

2 Preliminaries

Convex analysis [59, 11, 6]. For any closed convex function $f : \mathbf{R}^d \rightarrow \mathbf{R} \cup \{\infty\}$, we define the subdifferential set of f at $\bar{x} \in \mathbf{R}^m$, denoted by $\partial f(\bar{x})$, as the set of all $g \in \mathbf{R}^d$ such that

$$f(x) \geq f(\bar{x}) + \langle g, x - \bar{x} \rangle \quad \text{for all } x \in \mathbf{R}^d .$$

This is a generalization of the gradient mapping for functions that lack differentiability. The subdifferential set characterizes the set of global minima. In fact, it is easy to show that x^* is a minimizer of f if, and only if, $0 \in \partial f(x^*)$. For any closed convex set $C \subseteq \mathbf{R}^d$, we define its *normal cone* at $\bar{x} \in C$ to be

$$N_C(\bar{x}) := \{g \in \mathbf{R}^d \mid \langle g, x - \bar{x} \rangle \leq 0 \quad \text{for all } x \in C\} .$$

Analogously, $N_C(\bar{x})$ can be seen as the set of all points x such that $\mathbf{proj}_C(x) = \bar{x}$. The convex subgradient is known to follow a nice set of calculus rules. In this work we will make use of two of them: let $f_1 : \mathbf{R}^d \rightarrow \mathbf{R}$ and $f_2 : \mathbf{R}^d \rightarrow \mathbf{R} \cup \{\infty\}$ be closed convex functions and C be a closed convex set, then

$$\partial(f_1 + f_2)(x) = \partial f_1(x) + \partial f_2(x) \quad \text{and} \quad \partial \iota_C(x) = N_C(x) \quad \text{for all } x \in \mathbf{R}^d . \quad (3)$$

Given any convex function f , we define its *proximal operator* $\mathbf{prox}_f : \mathbf{R}^d \rightarrow \mathbf{R}^d$ as the unique solution of

$$\mathbf{prox}_f(\bar{x}) := \underset{x}{\operatorname{argmin}} f(x) + \frac{1}{2} \|x - \bar{x}\|_2^2 . \quad (4)$$

Using optimality conditions one can prove that

$$x \in (I + \partial f)(x^+) \iff x^+ = \mathbf{prox}_f(x) . \quad (5)$$

Linear programming [18, 43]. LP problems can be parameterized using multiple equivalent forms. For our theoretical results we focus on the *standard form* of LP:

$$\begin{aligned} & \text{minimize} && c^\top x \\ & \text{subject to} && Ax = b \quad (\in \mathbf{R}^m) \\ & && x \geq 0 \quad (\in \mathbf{R}^n) , \end{aligned} \quad (\mathbf{P})$$

where $A \in \mathbf{R}^{m \times n}$, $b \in \mathbf{R}^m$, and $c \in \mathbf{R}^n$ are given. The dual of this problem is given by

$$\begin{aligned} & \text{maximize} && b^\top x \\ & \text{subject to} && A^\top y \leq c \quad (\in \mathbf{R}^n) . \end{aligned} \quad (\mathbf{D})$$

Although the proofs in this paper are tailored to this form, the techniques we use should extend easily to any other form.

Farkas' Lemma states that a feasible solution of (\mathbf{P}) exists if, and only if, the following set is empty

$$\{y \in \mathbf{R}^m \mid b^\top y < 0 \text{ and } A^\top y \geq 0\} . \quad (6)$$

We call the elements of this set *certificates of primal infeasibility*, as their existence guarantees that the primal problem is infeasible. Analogously, the certificates of infeasibility for the dual problem (\mathbf{D}) are

$$\{x \in \mathbf{R}^n \mid c^\top x < 0, Ax = 0 \text{ and } x \geq 0\} . \quad (7)$$

Primal-dual hybrid gradient. Chambolle and Pock [15] establish convergence to a saddle point at a rate of $O(1/k)$ provided that a *saddle exists* and $\eta\tau\|A\|_2^2 < 1$.² The primal-dual problems **(P)**-**(D)** can be recast as a *convex-concave saddle point problem*. In particular we choose $g(x) = c^\top x + \iota_{\{x \geq 0\}}(x)$ and $h(y) = b^\top y$. In this case the proximal updates can be computed in closed form. In fact, a PDHG update reduces to

$$\begin{aligned} x^+ &= \mathbf{proj}_{\mathbf{R}_+^n}(x - \eta A^\top y - \eta c) \\ y^+ &= y + \tau A(2x^+ - x) - \tau b. \end{aligned} \quad (8)$$

Observe that the most complex operations in the update formula are matrix-vector products, and all other operations are separable by component.

An update of PDHG (Algorithm 1) can be equivalently defined with a differential inclusion of the form

$$M \begin{bmatrix} x^k - x^{k+1} \\ y^k - y^{k+1} \end{bmatrix} \in \begin{bmatrix} \partial g(x^{k+1}) \\ \partial h(y^{k+1}) \end{bmatrix} + \begin{bmatrix} A^\top y^{k+1} \\ -Ax^{k+1} \end{bmatrix} \quad \text{with} \quad M := \begin{bmatrix} \frac{1}{\eta} I_n & -A^\top \\ -A & \frac{1}{\tau} I_m \end{bmatrix}, \quad (9)$$

this follows from (5). We will later leverage this inclusion in our proofs.

Operators and the fixed-point iteration. We will find it useful to think of iterative algorithms from an operator viewpoint. Given an arbitrary map $T : \mathbf{R}^d \rightarrow \mathbf{R}^d$, the corresponding fixed-point iteration is defined as

$$z^{k+1} = T(z^k). \quad (10)$$

Most first-order methods can be described in this form. The primal-dual hybrid gradient method can be encoded as $T(x, y) = (x^+, y^+)$ where the output pair is defined in (8). When looking at an algorithm from this perspective, we transform the problem of finding a solution of the optimization problem to that of finding a fixed-point of the operator, i.e., $z^* = T(z^*)$. This idea has proven fruitful for proving optimal converge rates for a variety of algorithms [19].

Here we make a minimal assumption that is sufficient to analyze PDHG in the infeasible case. An operator T is said to be *nonexpansive* if it is 1-Lipschitz continuous with respect to a matrix norm $\|\cdot\|$, meaning that for any $z_1, z_2 \in \mathbf{R}^d$ we have

$$\|T(z_1) - T(z_2)\| \leq \|z_1 - z_2\|. \quad (11)$$

Nonexpansiveness does not ensure the convergence of iterates in the feasible case. Yet a slightly stronger condition does. An operator T is *firmly nonexpansive* if it satisfies

$$\|T(z_1) - T(z_2)\|^2 \leq \|z_1 - z_2\|^2 - \|(T - I)(z_1) - (T - I)(z_2)\|^2 \quad \text{for all } z_1, z_2 \in \mathbf{R}^d.$$

Note that the norm here is not necessarily the Euclidean norm. All the results concerning (firmly) nonexpansiveness in this section and the next one are with respect to the norm in which these properties hold. The following is a beautiful geometrical result proved by Pazy that defines a pivotal object in our studies.

Lemma 1 (Lemma 4 in [54]). *Let T be a nonexpansive operator, then the set $\text{cl}(\text{range}(T - I))$ is convex. Consequently, there exists a unique minimum norm vector in this set:*

$$v_T := \underset{z \in \text{cl}(\text{range}(T - I))}{\text{argmin}} \frac{1}{2} \|z\|^2. \quad (12)$$

This vector is known as the *infimal displacement vector*. We drop the subscript T and make the corresponding operator clear from the context. Intuitively, v is the minimum size perturbation we should subtract from T to ensure it has a fixed point.

²More precisely, Chambolle and Pock proved this rate for the primal-dual gap of the averaged iterates.

Theorem 1 ([54] and [3]). *Let T be a nonexpansive operator and (z^k) be a sequence generated by the fixed-point iteration (10). Then, we have*

$$\lim_{k \rightarrow \infty} \frac{z^k}{k} = v . \quad (13)$$

If further T is firmly nonexpansive, then

$$\lim_{k \rightarrow \infty} z^{k+1} - z^k = v . \quad (14)$$

That is the normalized iterate converges to the infimal displacement vector when T is nonexpansive and if T is firmly nonexpansive the difference of iterates also converge. One might wonder whether the stronger condition is necessary. This turns out to be the case.

The following proposition shows two things: first, (14) is provably stronger than (13) and second, the convergence of the iterates ensures convergence of the normalized average. Recall from the previous section that we use $\bar{z}^k := \frac{1}{k} \sum_{j=1}^k z^j$ to denote the average.

Proposition 1. *Let $(z^k)_{k=1}^\infty \subseteq \mathbf{R}^d$ be an arbitrary sequence and let $v \in \mathbf{R}^d$ be a fixed vector. Then the following implications hold:*

1. *Difference convergence implies normalized iterate convergence.*

$$\lim_{k \rightarrow \infty} (z^{k+1} - z^k) = v \implies \lim_{k \rightarrow \infty} \frac{z^k}{k} = v .$$

2. *Normalized iterate convergence implies normalized average convergence.*

$$\lim_{k \rightarrow \infty} \frac{z_k}{k} \rightarrow v \implies \lim_{k \rightarrow \infty} \frac{2\bar{z}^k}{(k+1)} \rightarrow v .$$

Moreover, these implications cannot be reversed as there exist simple counterexamples in \mathbf{R} .

The proof of this proposition is technical, so it is deferred to Appendix A.1.

Naturally when concerned with practical algorithms one would like to have convergence rates for (13) and (14). As far as we know, the state-of-the-art result in this vein is due to Liu, Ryu, and Yin [39].

Theorem 2 ([39]). *Let T be a firmly nonexpansive operator and (z^k) be a sequence generated by (10). Then, for any $\varepsilon > 0$, then there exists a point z_ε such that*

(Average iterate rate).

$$\left\| v - \frac{2}{k+1}(\bar{z}^k - z^0) \right\| \leq \sqrt{\frac{2}{k+1}} \|z^0 - z_\varepsilon\| + \varepsilon ,$$

(Last iterate rate).

$$\left\| v - \frac{1}{k}(z^k - z^0) \right\| \leq \sqrt{\frac{1}{k}} \|z^0 - z_\varepsilon\| + \varepsilon ,$$

(Difference rate).

$$\min_{j \leq k} \|v - z^{j+1} - z^j\| \leq \sqrt{\frac{1}{k}} \|z^0 - z_\varepsilon\| + \varepsilon .$$

Remark 1. *In the paper [39], this result is presented only for the difference of iterates, yet a simple modification of their argument proves the other two results.*

The theorem guarantees a rate of convergence that depends on a target accuracy ε . The rate could get worse as $\varepsilon \rightarrow 0$. Indeed, z_ε could diverge as ε goes to zero, see Example 4 in Appendix B. We will see in the next section that for LP it is possible to get rates that are independent of the accuracy ε .

Since the algorithm of interest is PDHG, we might wonder whether or not its operator is firmly nonexpansive. It turns out that it is, but with respect to the norm induced by the matrix M .

Proposition 2. *If $\eta\tau\|A\|_2^2 < 1$, then the operator defined by a PDHG iteration is firmly nonexpansive with respect to the norm the $\|\cdot\|_M$ with M defined as in (9).*

Proof. This is a known result [31], yet we include a proof for the interested reader. A Schur complement argument proves that the condition $\eta\tau\|A\|_2^2 < 1$ ensures that $M \succ 0$ is positive definite. Then a direct application of Proposition 4.2 in [9] proves the result. \square

3 Sublinear convergence of nonexpansive operators

This section presents the $O(1/k)$ convergence rate of the normalized iterates and the normalized average for nonexpansive operators. This rate applies to a broader class of operators than the previously known results (restated in Theorem 2) as it does not require the operator to be firmly nonexpansive. The resulting rate applies to many popular FOMs for convex optimization, including but not limited to PDHG [15], the Alternating Method of Multipliers (ADMM) or equivalently Douglas-Rachford Splitting (DRS) [56], and Mirror-Prox [49].

Theorem 3 presents our main result in this section.

Theorem 3. *Let T be a nonexpansive operator for some norm $\|\cdot\|$ and define v to be the minimum norm element in $\text{cl}(\text{range}(T - I))$. Then, for any $\varepsilon > 0$, there exists z_ε such that the following two inequalities hold*

(Average iterate rate).

$$\left\| v - \frac{2}{(k+1)} (\bar{z}^k - z^0) \right\| \leq \frac{4}{k+1} \|z^0 - z_\varepsilon\| + \varepsilon. \quad (15)$$

(Last iterate rate).

$$\left\| v - \frac{1}{k} (z^k - z^0) \right\| \leq \frac{2}{k} \|z^0 - z_\varepsilon\| + \varepsilon. \quad (16)$$

Furthermore, if $\text{range}(T - I)$ is closed, then there exists a finite z^* such that $T(z^*) = z^* + v$ and for any such z^* and all k :

(Average iterate rate).

$$\left\| v - \frac{2}{(k+1)} (\bar{z}^k - z^0) \right\| \leq \frac{4}{k+1} \|z^0 - z^*\|. \quad (17)$$

(Last iterate rate).

$$\left\| v - \frac{1}{k} (z^k - z^0) \right\| \leq \frac{2}{k} \|z^0 - z^*\|. \quad (18)$$

Remark 2. We comment that when $\text{range}(T - I)$ is not closed, Theorem 3 may not imply a $O(1/k)$ sublinear convergence rate. In fact, as $\varepsilon \rightarrow 0$, the vector $\|z_\varepsilon\|$ could grow to infinity, see Example 4 in Appendix B for a one dimensional example. When $\text{range}(T - I)$ is closed, we obtain a $O(1/k)$ sublinear rate.

Remark 3. When $\text{range}(T - I)$ is closed, the above result together with the lower bound proved later in Theorem 6 shows that the normalized iterates and normalized average of a nonexpansive operator exhibit a $\Theta\left(\frac{1}{k}\right)$ convergence rate. It is faster than the difference of iterates, by noticing that the difference of iterates converges at rate $\Theta\left(\frac{1}{\sqrt{k}}\right)$ (see Theorem 8 of [19]).

The next Lemma is used in the proof of Theorem 3.

Lemma 2. *Suppose the assumptions of Theorem 3. Fix $\varepsilon > 0$, then there exists a point z_ε , such that the following two inequalities hold for all $k \geq 0$:*

1. $\|T^{k+1}(z_\varepsilon) - T^k(z_\varepsilon) - v\| \leq \varepsilon.$
2. $\|(T^k(z_\varepsilon) - z_\varepsilon) - kv\| \leq k\varepsilon.$

Furthermore, if $\text{range}(T - I)$ is closed, then there exists a point z^* such that $T(z^*) = z^* + v$. For all such z^* , for all $k \geq 0$:

3. $T^{k+1}(z^*) - T^k(z^*) = v.$
4. $T^k(z^*) - z^* = kv.$

Proof. Without loss of generality we assume that $\varepsilon \leq 1$. Notice $v \in \text{cl}(\text{range}(T - I))$, thus there exists z_ε such that

$$\|T(z_\varepsilon) - z_\varepsilon - v\| \leq \frac{\varepsilon^2}{\max\{1, 2(\|v\| + 1)\}}. \quad (19)$$

We start by proving the first claim. Fix an arbitrary $k \geq 0$. We will make use of two facts in the proof. Since T is a nonexpansive operator, an application of the triangle inequality yields

$$\|T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon)\| - \|v\| \leq \|T(z_\varepsilon) - z_\varepsilon\| - \|v\| \leq \|T(z_\varepsilon) - z_\varepsilon - v\| \leq \frac{\varepsilon^2}{2(\|v\| + 1)}. \quad (20)$$

Noticing v is the nearest point to zero in $W = \text{cl}(\text{range}(T - I))$ with respect to the norm $\|\cdot\|$ and the set W is convex, it follows from the optimality conditions of this problem that

$$\langle w, v \rangle \geq \|v\|^2 \quad \text{for all } w \in W. \quad (21)$$

Armed with these two facts, we derive for any arbitrary k :

$$\begin{aligned} \|T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon) - v\|^2 &= \|T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon)\|^2 - 2\langle T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon), v \rangle + \|v\|^2 \\ &\leq \|T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon)\|^2 - 2\|v\|^2 + \|v\|^2 \\ &= (\|T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon)\| + \|v\|) (\|T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon)\| - \|v\|) \\ &\leq (\|T(z_\varepsilon) - z_\varepsilon - v\| + 2\|v\|) \frac{\varepsilon^2}{2(\|v\| + 1)} \\ &\leq (\varepsilon^2 + 2\|v\|) \frac{\varepsilon^2}{2(\|v\| + 1)} \leq \varepsilon^2, \end{aligned}$$

where the first inequality utilizes (21) by noticing $T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon) \in W$, the second inequality uses (20) and the triangle inequality, the third inequality is from (19), and the last inequality uses $\varepsilon \leq 1$. This proves the first statement.

The second claim follows by induction. The base case $k = 0$ holds directly. For the inductive step, assume that the statement holds for $k - 1$. Then

$$\|(T^k(z_\varepsilon) - z_\varepsilon) - kv\| \leq \|T^k(z_\varepsilon) - T^{k-1}(z_\varepsilon) - v\| + \|T^{k-1}(z_\varepsilon) - z_\varepsilon - (k-1)v\| \leq k\varepsilon,$$

where we used the first claim and the inductive hypothesis.

Furthermore, if $\text{range}(T - I)$ is closed, the statements follow by taking $\varepsilon = 0$ in the previous proofs. \square

Proof of Theorem 3. Let z_ε be the point given by Lemma 2. We proceed to prove the first two statements.

1. It follows from $z^j = T^j(z^0)$ that

$$\begin{aligned} \left\| \frac{2}{k(k+1)} \sum_{j=1}^k (z^j - z^0) - v \right\| &= \left\| \frac{2}{k(k+1)} \sum_{j=1}^k (T^j(z^0) - z^0 - jv) \right\| \\ &= \left\| \frac{2}{k(k+1)} \sum_{j=1}^k ((T^j(z^0) - T^j(z_\varepsilon)) + (z_\varepsilon - z^0) + (T^j(z_\varepsilon) - z_\varepsilon - jv)) \right\| \\ &\leq \left\| \frac{2}{k(k+1)} \sum_{j=1}^k (T^j(z^0) - T^j(z_\varepsilon)) \right\| + \frac{2}{(k+1)} \|z^0 - z_\varepsilon\| + \varepsilon, \end{aligned}$$

where the inequality uses Lemma 2 and the triangle inequality. Applying the triangle inequality to the first term yields

$$\begin{aligned} \left\| \frac{2}{k(k+1)} \sum_{j=1}^k (T^j(z^0) - T^j(z_\varepsilon)) \right\| &\leq \frac{2}{k(k+1)} \sum_{j=1}^k \|T^j(z^0) - T^j(z_\varepsilon)\| \\ &\leq \frac{2}{k(k+1)} \sum_{j=1}^k \|z^0 - z_\varepsilon\| = \frac{2}{(k+1)} \|z^0 - z_\varepsilon\|, \end{aligned}$$

where the second inequality follows since T is nonexpansive.

2. Notice that

$$\begin{aligned} \left\| \frac{1}{k} (z^k - z^0) - v \right\| &= \left\| \frac{1}{k} ((T^k(z^0) - z^0) - (T^k(z_\varepsilon) - z_\varepsilon) + (T^k(z_\varepsilon) - z_\varepsilon - kv)) \right\| \\ &\leq \left\| \frac{1}{k} ((T^k(z^0) - z^0) - (T^k(z_\varepsilon) - z_\varepsilon)) \right\| + \varepsilon \\ &\leq \frac{1}{k} (\|T^k(z^0) - T^k(z_\varepsilon)\| + \|z^0 - z_\varepsilon\|) + \varepsilon \\ &\leq \frac{2}{k} \|z^0 - z_\varepsilon\| + \varepsilon, \end{aligned}$$

where the first inequality uses triangle inequality and Lemma 2, and the last inequality is from non-expansiveness of T .

Finally, if $\text{range}(T - I)$ is closed, we obtain the results following the same argument as above with $\varepsilon = 0$, using the results for closed $\text{range}(T - I)$ from Lemma 2. \square

A drawback of (15) and (16) in Theorem 3, as well as the results in Theorem 2, is that the constants accompanying the rates depend on ε . Nonetheless, we can bypass this issue, using (17) and (18), for problems where $\text{range}(T - I)$ is closed. The next proposition guarantees that $\text{range}(T - I)$ is indeed closed for a broad family of algorithms for solving LP.

Proposition 3. *Let $T : \mathbf{R}^d \rightarrow \mathbf{R}^d$ be an operator that can be decomposed as $T = T_k \circ \dots \circ T_1$ where T_j is either an affine mapping or a projection onto a polyhedron. Then, $\text{range}(T)$ is a finite union of polyhedra.*

Proof. The proof follows inductively. Assume that $C = \text{range}(T_j \circ \dots \circ T_1)$ is a finite union of polyhedra. Without loss of generality, we can assume that C is equal to a single polyhedron. Now we consider two cases:

Case 1. Assume that T_{j+1} is an affine transformation. This is a well-known consequence of Fourier-Motzkin elimination [43].

Case 2. Assume that T_{j+1} is a projection onto a polyhedron Q . First, we start with an intuitive sketch of the proof and then formalize it. In this case, different pieces of the polyhedron C are going to be projected to different faces of the polyhedron Q . Each one of these pieces is a polyhedron and since there are only finitely many faces of Q , the projection is a finite union of polyhedra.

More formally, any polytope Q defines a finite polyhedral partition of the space $\{P_F\}_{F \in \Delta}$ where Δ is the collection of faces of the polyhedron Q .³ Each cell P_F corresponds to the region of the space that projects onto F , that is $\mathbf{proj}_Q(P_F) = F$. Define a partition of the polyhedron C as $\{C_F\}_{F \in \Delta}$ given by $C_F = P_F \cap C$. Within each cell P_F the projection $\mathbf{proj}_Q|_{P_F}$ is an affine transform. Thus, by Case 1 we have that $\mathbf{proj}_Q(C_F)$ is a polyhedron and thus

$$T_{j+1}(C) = \mathbf{proj}_Q(C) = \bigcup_{F \in \Delta} \mathbf{proj}_Q(C_F)$$

is a finite union of polyhedra. □

As a result of Proposition 3 applied to the PDHG update for solving an LP problem, $\text{range}(T - I)$ is a polyhedron, thus closed:

Corollary 1. *Let T be the PDHG operator for an LP problem, then $\text{range}(T - I)$ is a finite union of closed polyhedra.*

Proof. Notice that the operator $T - I$ is composite of linear operators and projection operators, thus we obtain the results by using Proposition 3. □

Remark 4. *Proposition 3 shows that $\text{range}(T - I)$ is closed when T is an operator that corresponds to other first-order algorithms, such as ADMM and mirror-prox, to solve an LP problem. Further, one can extend the result to cover convex quadratic problems with polyhedral constraints.*

4 The complete behavior of PDHG for solving LP problems

In Figure 1, we saw low-dimensional examples of the dynamics of PDHG when solving LP problems in different feasibility settings. Indeed, such convergence/divergence dynamics generally hold when using PDHG to solve arbitrary LP problems. In this section, we present a complete description of the behavior of PDHG for feasible and infeasible LP problems and discuss how to recover the infeasibility certificate from the iterates of PDHG. The next theorem compiles the full characterization. The proof of this theorem will be deferred to Section 4.3.

Theorem 4. *Consider the primal (P) and dual (D) problems. Assume that $\eta\tau\|A\|_2^2 < 1$, let T be the operator induced by (8), and let $\{z^k\}_k$ be a sequence generated by the fixed-point iteration for an arbitrary starting point z^0 , i.e., $z^k = T^k(z^0)$. Then, one of the following holds:*

1. *If both primal and dual are feasible, then the iterates (x^k, y^k) converge to a primal-dual solution $z^* = (x^*, y^*)$ and $v = (T - I)(z^*) = 0$.*
2. *If both primal and dual are infeasible, then both primal and dual iterates diverge to infinity towards the direction of the infimal displacement vector $v = (v_x, v_y)$. Moreover, the primal and dual components of the infimal displacement vector v_x and v_y give certificates of dual and primal infeasibility, respectively.*
3. *If the primal is infeasible and the dual is feasible, then the dual iterates diverge to infinity in the direction of v_y , while the primal iterates converge to a vector x^* . Furthermore, the dual-component v_y is a certificate of primal infeasibility, and there exists a vector y^* such that $v = (T - I)(x^*, y^*)$.*

³I.e., each cell $P_F \subseteq \mathbf{R}^d$ is a polyhedron and $\cup_{F \in \Delta} P_F = \mathbf{R}^d$.

4. *If the primal is feasible and the dual is infeasible, then the same conclusions as in the previous item hold by swapping primal with dual.*

To show this characterization, we establish two intermediate results: first, the infimal displacement vector v is nonzero if, and only if, either the primal or dual problems are infeasible; and second, the iterates (x^k, y^k) “converge” to a well-defined ray of the form $(x^*, y^*) + \lambda v$ for $\lambda \in \mathbf{R}_+$. The first result describes the asymptotic divergent behavior of the primal (resp. dual) iterates when the dual (resp. primal) problem is infeasible. The second one, ensures the asymptotic convergence of the primal (resp. dual) iterates without any normalization when the dual (resp. primal) problem is feasible. These two intermediate results are proved in Section 4.1 and Section 4.2, respectively.

4.1 The infimal displacement vector recovers infeasibility certificates

In Section 3, we demonstrated that the differences of iterates, the normalized iterates, and the normalized average for a nonexpansive operator converge to the infimal displacement vector v . Here, we show that the infimal displacement vector v for PDHG applied to LP recovers infeasibility certificates whenever it is nonzero. First, some simple properties of v .

Lemma 3. *Consider the primal (P) and dual (D) problems. Assume that $\eta\tau\|A\|_2^2 < 1$, let T be the operator induced by (8), and let $v = (v_x, v_y)$ be the infimal displacement vector of T . Then $v_x \geq 0$, $Av_x = 0$, and $A^\top v_y \geq 0$.*

Proof. From Theorem 1,

$$\frac{1}{k}(x^k, y^k) \rightarrow v = (v_x, v_y) \quad \text{and} \quad \frac{1}{k}(z_{k+1} - z_k) \rightarrow 0. \quad (22)$$

Notice that PDHG for LP has the following iteration update in terms of a differential inclusion,

$$M \begin{bmatrix} x^k - x^{k+1} \\ y^k - y^{k+1} \end{bmatrix} \in \begin{bmatrix} N_{\mathbf{R}_+^n}(x^{k+1}) + A^\top y^{k+1} + c \\ -Ax^{k+1} + b \end{bmatrix}, \quad (23)$$

where this relation comes from (9) and (3). Dividing (23) by k and letting $k \rightarrow \infty$, we have from (22) that

$$0 \in \lim_{k \rightarrow \infty} N_{\mathbf{R}_+^n}(x^k) + A^\top \frac{1}{k} y_k \subseteq -\mathbf{R}_+^n + A^\top v_y \implies A^\top v_y \geq 0, \quad (24)$$

where we utilize the fact that $N_{\mathbf{R}_+^n}(x) \subseteq -\mathbf{R}_+^n$ for any $x \in \mathbf{R}_+^n$ and $\lim_{k \rightarrow \infty} \frac{1}{k} y_k = v_y$; and

$$0 = \lim_{k \rightarrow \infty} -\frac{1}{k} Ax^k = -Av_x \implies Av_x = 0. \quad (25)$$

Furthermore, note that $v_x \geq 0$ since $v_x = \lim_{k \rightarrow \infty} x^k/k$ and $x^k \geq 0$ for all k . □

Now we derive the main result of this section.

Proposition 4. *Consider the primal (P) and dual (D) problems. Assume that $\eta\tau\|A\|_2^2 < 1$, let T be the operator induced by (8), and let $(z^k)_{k \in \mathbf{N}}$ be a sequence generated by the fixed-point iteration for an arbitrary starting point z^0 . Then, the primal problem (P) is infeasible if and only if v_y is a nonzero vector, and in this case, v_y is an infeasibility certificate for the primal problem. Analogously, the dual problem (D) is infeasible if and only if v_x is a nonzero vector, and in this case, v_x is an infeasibility certificate for the dual problem.*

Proof. To establish the first implication in this result we have to prove that if v_y is non-zero, then v_y is an infeasibility certificate for the primal problem, namely,

$$A^\top v_y \geq 0 \text{ and } b^\top v_y < 0,$$

thus the primal problem is infeasible. Similarly if v_x is non-zero, then v_x is an infeasibility certificate for the dual problem, namely

$$Av_x = 0, v_x \geq 0, \text{ and } c^\top v_x < 0 ,$$

thus the dual problem is infeasible. We proved all the nonstrict inequalities in Lemma 3, so it suffices to show the strict ones.

First, consider the case when $v_x \neq 0$. Let $B = \{i \in [n] \mid (v_x)_i > 0\}$ and let $N = \{i \in [n] \mid (v_x)_i = 0\}$, then $B \cup N = [n]$ by noticing $v_x \geq 0$ (from Lemma 3). Given a vector x , let x_B be the vector of entries of x with indices in B ; similarly given a matrix A , let A_B be the submatrix with columns of A with indices in B . Then for any $i \in B$, we have $(v_x)_i > 0$, thus there exists some K such that $(x^k/k)_i > 0$ for all $k \geq K$, and furthermore

$$(x^{k+1})_B = (x^k)_B - \eta A_B^\top y^k - \eta c_B .$$

Taking the limit $k \rightarrow \infty$ and noticing $\lim_{k \rightarrow \infty} (x^{k+1})_B - (x^k)_B = (v_x)_B$, we obtain

$$(v_x)_B = \lim_{k \rightarrow \infty} -\eta(A_B^\top y^k + c_B) .$$

Thus it holds that

$$c^\top v_x = c_B^\top (v_x)_B = -\frac{1}{\eta} \|v_x\|_2^2 - \lim_{k \rightarrow \infty} (y^k)^\top A_B (v_x)_B = -\frac{1}{\eta} \|v_x\|_2^2 < 0 , \quad (26)$$

where the last equality uses $A_B(v_x)_B = Av_x = 0$. Combining with $v_x \geq 0$ and $Av_x = 0$ proves that v_x is a certificate of infeasibility whenever it is nonzero.

Second, consider the case when $v_y \neq 0$. By taking $k \rightarrow \infty$, we have

$$v_y = \lim_{k \rightarrow \infty} y_{k+1} - y_k = \lim_{k \rightarrow \infty} \tau A(2x^{k+1} - x^k) - \tau b = \lim_{k \rightarrow \infty} \tau A x^{k+2} - \tau b , \quad (27)$$

where the second equality uses the update rule (8), and the third equality uses $\lim_{k \rightarrow \infty} 2x^{k+1} - x^k = \lim_{k \rightarrow \infty} x^{k+1} + v_x = \lim_{k \rightarrow \infty} x^{k+2}$.

Now we claim the following two facts:

Fact 1. *There exists some K such that if $(A^\top v_y)_i > 0$ then $x_i^k = 0$ for all $k \geq K$.*

Fact 2. *The support (nonzero components) of $A^\top v_y$ satisfies $\text{supp}(A^\top v_y) \subseteq N$.*

The first fact is because if $(A^\top v_y)_i > 0$ then we have that $(A^\top y^k/k)_i \geq (A^\top v_y)_i/2 > 0$ for large enough k . Dividing (23) by k yields

$$-\frac{1}{k}(A^\top y^k)_i + \frac{1}{\eta k} (x^k - x^{k+1} - \eta c)_i \in N_{\mathbf{R}_+^n}(x_i^{k+1}) .$$

For large enough k , the second term on the left-hand side of the inclusion will be as small as $(A^\top y^k/k)_i/2$ and hence the sign of entire expression on the left-hand side will be negative. If $N_{\mathbf{R}_+}((x^{k+1})_i)$ contains a negative number, then $(x^{k+1})_i = 0$, which implies that $(x^k)_i = 0$ for large enough k .

The second fact is because for any entry i in the support of $A^\top v_y$, namely $(A^\top v_y)_i > 0$, it follows from the first part that $(x^k)_i = 0$ for all k large enough, thus $(v_x)_i = \lim_{k \rightarrow \infty} \frac{1}{k} x_i^k = 0$, which proves the second fact by the definition of the set N .

Returning to the proof of Proposition 4, notice that

$$\lim_{k \rightarrow \infty} v_y^\top A x^k = \lim_{k \rightarrow \infty} \sum_{i \in N} (A^\top v_y)_i x_i^k = 0 , \quad (28)$$

where the first equality uses Fact 2, and the second equality uses Fact 1. Therefore, it holds that

$$v_y^\top b = -\frac{1}{\tau} \|v_y\|_2^2 + \lim_{k \rightarrow \infty} v_y^\top A x^{k+2} = -\frac{1}{\tau} \|v_y\|_2^2 < 0 ,$$

where the first inequality uses (27) and the second equality is from (28). Together with (24), we know v_y is an infeasibility certificate for the primal problem.

Now we turn to the inverse direction. Recall that it follows from the closedness of the set $\text{range}(T - I)$ that there exists a pair $z^* = (x^*, y^*)$ such that $T(z^*) = z^* + v$. If the dual problem is infeasible, we will show that $v_x \neq 0$ by contradiction. Assume $v_x = 0$; then it follows from the update rule (8) that

$$x^* = \text{proj}_{\mathbf{R}_+^n}(x^* - \eta(A^\top y^* + \eta c)) ,$$

thus $A^\top y^* + \eta c \geq 0$ by noticing $x^* \geq 0$, which contradicts the assumption that the dual problem is infeasible. If the primal problem is infeasible, we will show that $v_y \neq 0$ by contradiction. Suppose $v_y = 0$, then it follows from the update rule (8) that

$$y^* = y^* + \tau A(2(x^* + v_x) - x^*) - \tau b ,$$

thus $Ax^* = b$ by noticing $Av_x = 0$ from (25). Furthermore, we know $x^* \geq 0$, thus the primal is a feasible problem, which contradicts with assumption.

This concludes the proof. □

4.2 The iterates converge to a ray

Combining facts from the previous sections we know that if both primal and dual problems are feasible then the iterates (without normalization) will converge to a solution, and when both primal and dual problems are infeasible then the normalized iterates converge to a vector (v_x, v_y) with nonzeros on both primal and dual components. Yet the techniques used to prove these results do not explain what happens when one of the problems is feasible and the other one is infeasible. In this scenario the convergence of the primal and dual iterates happen at different scales, one with normalization by $\frac{1}{k}$ and the other without it. In this section, we fill in this gap by showing that the iterates of PDHG always converge to ray with direction v , emanating from a point z^* . In turn, this allows us to connect the convergence results for the two scales.

Definition 1 (Ray). Given a starting point $z^* \in \mathbf{R}^{n+m}$ and a direction $v \in \mathbf{R}^{n+m}$, we define their ray as

$$[z^*, v] = \{z^* + \lambda v \mid \lambda \in \mathbf{R}_+\} .$$

With this definition at hand we can now state the main result of this section.

Theorem 5. Consider the primal (P) and dual (D) problems. Assume that $\eta\tau\|A\|_2^2 < 1$, let T be the operator induced by (8), and let $(z^k)_{k \in \mathbf{N}}$ be a sequence generated by the fixed-point iteration for an arbitrary starting point z^0 . Then, the iterates of PDHG converge to a ray $[z^*, v]$, in particular

$$\|z^k - z^* - kv\| \rightarrow 0 \quad \text{for some } z^* \in (T - I)^{-1}(v) .$$

To prove this result, we establish a connection between the iterates of PDHG applied to the original (possibly infeasible) problem and the iterates of PDHG applied to a feasible auxiliary LP problem. Let us start by defining this auxiliary problem. Define the index sets

$$\begin{aligned} B &= \{i \in [n] \mid (v_x)_i > 0\} , \\ N_1 &= \{i \in [n] \mid (v_x)_i = 0, (A^\top v_y)_i = 0\} , \\ N_2 &= \{i \in [n] \mid (v_x)_i = 0, (A^\top v_y)_i > 0\} . \end{aligned} \tag{29}$$

Define the operator $\tilde{T} : \mathbf{R}^{n+m} \rightarrow \mathbf{R}^{n+m}$ given by $\tilde{T}(z) := z^+$ with

$$\begin{aligned} (x^+)_B &= x_B - \eta A_B^\top y - \eta c_B - (v_x)_B \\ (x^+)_{N_1} &= \text{proj}_{\mathbf{R}_+^{|N_1|}}(x_{N_1} - \eta A_{N_1}^\top y - \eta c_{N_1}) - (v_x)_{N_1} \\ (x^+)_{N_2} &= -(v_x)_{N_2} \\ y^+ &= y + \tau A(2x^+ - x) - \tau b - v_y . \end{aligned} \tag{30}$$

In turn, this is a PDHG operator for the auxiliary LP problem:

$$\begin{aligned}
& \text{minimize} && (c_B + (v_x)_B/\eta)^\top x_B + c_{N_1}^\top x_{N_1} + c_{N_2}^\top x_{N_2} \\
& \text{subject to} && A_B x_B + A_{N_1} x_{N_1} + A_{N_2} x_{N_2} = b + \frac{v_y}{\tau} \\
& && x_{N_1} \geq 0, \ x_{N_2} = 0.
\end{aligned} \tag{31}$$

Then we claim the following connection between T and \tilde{T} .

Proposition 5. *Given an arbitrary initial solution z^0 , there exists a large enough $K \in \mathbf{N}$ such that*

$$\tilde{T}^k(z^K) = T^k(z^K) - kv \quad \text{for all } k \geq 0. \tag{32}$$

Proof. For any initial solution, we know that there exists some K such that it holds for any $k \geq K$ that

$$(x^k)_B > 0 \text{ and } (x^k)_{N_2} = 0. \tag{33}$$

The former is because $(v_x)_B > 0$, and the latter follows from Fact 1. With some abuse of notation, we let $z^0 \leftarrow z^K$, so that we may study the iterates starting at z^0 (rather than starting at z^K), for notational convenience. From Lemma 3:

$$v_x \geq 0, \quad Av_x = 0, \quad \text{and} \quad A^\top v_y \geq 0.$$

In addition, from the converse of Fact 1,

$$(A^\top v_y)_B = 0.$$

We show the stated claim by induction. Denote $z^k = T^k(z^0)$ and $\tilde{z}^k = \tilde{T}^k(z^0)$. First, (32) holds with $k = 0$. Now suppose (32) holds for k , and consider $k + 1$. Then by induction we have $\tilde{z}^{k+1} = \tilde{T}(z^k - kv)$, thus it holds by (30) that

$$\begin{aligned}
(\tilde{x}^{k+1})_B &= (x^k)_B - k(v_x)_B - \eta A_B^\top (y^k - kv_y) - \eta c_B - (v_x)_B \\
&= (x^{k+1})_B - (k+1)(v_x)_B,
\end{aligned}$$

where the second equality utilizes the update rule of PDHG by noticing $A_B^\top v_y = 0$ and $(x^{k+1})_B > 0$. For the components in N_1 we get

$$\begin{aligned}
(\tilde{x}^{k+1})_{N_1} &= \mathbf{proj}_{\mathbf{R}_+^{|N_1|}}((x_k)_{N_1} - k(v_x)_{N_1} - \eta A_{N_1}^\top (y^k - kv_y) - \eta c_{N_1}) - (v_x)_{N_1} \\
&= \mathbf{proj}_{\mathbf{R}_+^{|N_1|}}((x_k)_{N_1} - \eta A_{N_1}^\top y^k - \eta c_{N_1}) \\
&= (x^{k+1})_{N_1} - (k+1)(v_x)_{N_1},
\end{aligned}$$

where the second equality follows from $A_{N_1}^\top v_y = 0$ and $(v_x)_{N_1} = 0$, the third one utilizes $(v_x)_{N_1} = 0$ and the update rule of PDHG. Similarly, for the N_2 block

$$(\tilde{x}^{k+1})_{N_2} = -(v_x)_{N_2} = 0 = (x^{k+1})_{N_2} - (k+1)(v_x)_{N_2},$$

where the equations follow from $(x^{k+1})_{N_2} = 0$ and $(v_x)_{N_2} = 0$. Finally, for the dual iterates we have

$$\begin{aligned}
\tilde{y}^{k+1} &= y^k - kv_y + \tau A(2\tilde{x}^{k+1} - \tilde{x}^k) - \tau b - v_y \\
&= y^k + \tau A(2(x^{k+1} - (k+1)v_x) - (x^k - kv_x)) - \tau b - (k+1)v_y \\
&= y^k + \tau A(2x^{k+1} - x^k) - \tau b - (k+1)v_y \\
&= y^{k+1} - (k+1)v_y,
\end{aligned}$$

where the third equality utilizes $Av_x = 0$, and the last equality is from the update rule of PDHG. \square

Equipped with this proposition we can now prove the theorem.

Proof of Theorem 5. Since \tilde{T} is a PDHG operator, it is firmly nonexpansive with respect to $\|\cdot\|_M$. Thus, if \tilde{T} has a fixed point, then the iteration $\tilde{T}^k(z^K)$ should converge to it. To see that \tilde{T} has a fixed point, let z^* be a point such that $(T - I)(z^*) = v$ and let K be the iteration after which $\tilde{T}^k(T^K(z^*)) = T^{k+K}(z^*) - kv$, which exists thanks to Proposition 5. We claim that $T^K(z^*)$ is a fixed point of \tilde{T} . To see this, note that

$$\tilde{T}(T^K(z^*)) = T^{K+1}(z^*) - v = T^K(z^*),$$

where the last equality follows from Lemma 2.

Now, let z^0 an arbitrary initial point and recall that K is defined in (32). Now that we know that the set of fixed points of \tilde{T} is nonempty, we can define $z^* = \lim_{k \rightarrow \infty} \tilde{T}^k(T^K(z_0))$. We will prove that z^* satisfies $(T - I)(z^*) = v$. Due to Proposition 5, we know that

$$z^* = \tilde{T}(z^*) = T(z^*) - v.$$

Finally, using decomposition (32) we get

$$\|z^{k+K} - z^* - kv\| = \|\tilde{T}(z^k) - z^*\| \rightarrow 0. \quad (34)$$

The statement of theorem claimed this convergence where the coefficient accompanying v is $(k + K)$. We can get around this by setting $z^* \leftarrow z^* - Kv$, a point that also satisfies $(T - I)(z^*) = v$ thanks to Lemma 2. This establishes the result. \square

4.3 Proof of Theorem 4

Proof. As a direct result of Proposition 4, we know that if both primal and dual are feasible, then $v = 0$ and Theorem 5 ensures that PDHG converges to an optimal solution (or equivalently a fixed point of T). If primal (and/or dual) is infeasible, then the dual iterate of PDHG (and/or primal) diverges to infinity, and the diverging direction recovers primal (and/or dual) infeasibility certificate.

Thus, the only thing left to prove is the final conclusion of item 3 (and analogously item 4). Assume that the primal problem is infeasible and the dual is feasible. By Proposition 4 we know that $v_x = 0$. Then, Theorem 5 guarantees the existence of some $z^* = (x^*, y^*)$ such that $x^k \rightarrow x^* + kv_x = x^*$ and $(T - I)(z^*) = v$. The proof for the case where the primal is feasible and the dual is infeasible follows from an analogous argument. This completes the proof of the theorem. \square

5 Finite time identifiability and eventual linear convergence

In this section, we introduce a nondegeneracy condition that ensures that after a finite amount the difference of iterates converges linearly to the infimal convergence vector. To show this, we demonstrate under said condition the iterates “identify” the support of x^k , i.e., the support freezes after a finite number iterations. Finite-time identification has a long history in the analysis of iterative algorithms for feasible problems [22, 14, 13, 12, 26, 34, 37]. Roughly speaking, these algorithms’ behaviors exhibit two phases: a first one that only takes finitely many steps but suffers from slow sublinear convergence, and then a second one after the active set is identified where the convergence is significantly faster and becomes linear.

For PDHG, finite-time identifiability is known to hold for *feasible* minimax problems under suitable nondegeneracy conditions [37]. In contrast, here we study this phenomenon for infeasible LP problems. We demonstrate that *even when there is no primal (and/or dual) feasible solution, active set of the iterates with respect to an auxiliary feasible LP problem is fixed after finitely many iterations*.

Recall that the iterates of PDHG converge to a ray $[z^*, v] = \{z^* + kv \mid k \in \mathbf{N}\}$ where z^* is a solution to the feasible LP problem given by (31). Consider the constraint set defined by said auxiliary problem, that is

$$Ax = b + \frac{v_y}{\tau}, \quad x_{N_1} \geq 0 \quad \text{and} \quad x_{N_2} = 0. \quad (35)$$

Here, the active set is the set of inequality constraints that attain their extreme values, namely, $\{i \in N_1 \mid x_i^* = 0\}$. Note that when the problem is feasible, $N_1 = \{1, \dots, n\}$ and thus the constrained set defined by the auxiliary problem (35) matches that of the original problem.

Now, we introduce the nondegeneracy condition for possibly infeasible problems, which generalizes the classical identifiability theory of PDHG for feasible LP problems [34, 37]. Similar variations of the nondegeneracy condition have appeared in numerous works that deal with finite time identifiability. Further generalizations of this idea have led to conditions beyond the context of optimization, we refer the interested reader to [35] for a perspective from differential geometry.

Definition 2. A ray $[z^*, v]$ is nondegenerate if for any $i \in N_1$ (recall N_1 is defined in (29)), the pair $(x_i^*, (A^\top y^* + c)_i)$ satisfies strict complementarity with respect to the auxiliary problem (31): $x_i^* > 0$ if, and only if, $(A^\top y^* + c)_i = 0$.

Although here we chose to define nondegeneracy in terms of the extreme point z^* , this definition is independent of the point we take in the ray $[z^*, v]$. This follows easily from the fact that $(v_x)_{N_1}$ and $(A^\top v_y)_{N_1}$ are zero vectors. Additionally, notice that when the original problem is feasible, this definition reduces to the classical strict complementarity of the original problem.

The next lemma presents the finite time identifiability of PDHG for infeasible LP:

Lemma 4. Suppose $[z^*, v]$ is a nondegenerate ray. Then, every PDHG iterate sequence $z^k = (x^k, y^k)$, converging to the ray $[z^*, v]$, fixes the active set of (35) after finitely many steps. Furthermore, this ensures that the support of x^k is fixed for all large enough k .

Proof. First let us prove that the active set of (35) is identified in finite time. Notice that this is equivalent to saying that the support of $x_{N_1}^k$ freezes after finitely many iterations. Let $i \in N_1$, due to strict complementarity it is enough to consider two cases:

Case 1. Assume that $(Ay^* + c)_i > 0$, then complementary slackness implies $(x^k)_i \rightarrow 0$. By construction, we should have $\eta \cdot (A^\top y^k + c)_i > x_i^k$ for all k large enough. After this condition starts to hold, the PDHG update at i gives

$$x_i^{k+1} = (x_i^k - \eta \cdot (A^\top y^k + c)_i)_+ = 0.$$

Hence, we have $x_i^k = 0$ for all large k .

Case 2. Assume that $x_i^* > 0$, then after finitely many iterations we have $x_i^k > 0$.

Thus, the support of $x_{N_1}^k$ is identified in finite time. Now, we argue that the same happens to the support of x^k . Assume that $i \in B$, then $(v_x)_i > 0$ and consequently for all large k we have $x_i^k/k > 0$, as we wanted. Lastly, if $i \in N_2$ then $A^\top v_y > 0$ and Fact 1 guarantees that $x_i^k = 0$ for large enough k . This concludes the proof. \square

When nondegeneracy holds, PDHG eventually identifies the support of the primal iterate x^k . This simplifies the form of each iteration. Let S be the support of any x^k with k large enough. The projection to the positive orthant applied by PDHG (8) becomes a projection to the subspace $\{x \mid \text{supp}(x) = S\}$. As a consequence, one can recast each iteration (8) as an affine transformation:

$$\begin{bmatrix} x^{k+1} \\ y^{k+1} \end{bmatrix} = \underbrace{\begin{bmatrix} I & -\eta D A^\top \\ \tau A D & I - 2\tau \eta A D A^\top \end{bmatrix}}_{Q:=} \begin{bmatrix} x^k \\ y^k \end{bmatrix} - \underbrace{\begin{bmatrix} \eta D c \\ 2\tau \eta A D c + \tau b \end{bmatrix}}_{p:=} \quad (36)$$

where D is a diagonal matrix with ones on the indices (i, i) such that $i \in S$ and zeros everywhere else, and matrix Q and vector p are defined in (36).

The next theorem presents upper and lower bounds for the convergence of the three sequences (2) under the nondegeneracy condition. In particular, we show that the difference of iterates (2a) exhibit eventual linear convergence, while normalized iterates (2b) and normalized average iterates (2c) exhibit eventual sublinear convergence.

Theorem 6 (Eventual convergence rate of the three sequences under nondegeneracy). *Consider the primal (P) and dual (D) problems. Assume that $\eta\tau\|A\|_2^2 < 1$, and let $(z^k)_{k \in \mathbf{N}}$ be a sequence generated by PDHG. Suppose that the iterates $z^k = (x^k, y^k)$ converge to a nondegenerate ray. Then, the k th power of Q converges $\lim_{k \rightarrow \infty} Q^k = Q^\infty$ to a projection matrix, and there exists a finite K such that for any $k \geq 0$, the active set of x^{K+k} is fixed. Furthermore, there exist positive constants $\underline{\mu}, c_1, c_2, c_3, c_4 > 0$ such that the following holds:*

1. **Linear convergence of the differences.** *For any $\mu \in (\sqrt{1 - \eta\tau\sigma_{\min}(A)^2}, 1)$, the differences $z^{K+k+1} - z^{K+k}$ converge at a linear rate to the infimal displacement vector v , i.e.,*

$$\underline{\mu}^k \|(Q - I)z^K + (Q^\infty - I)p\|_2 \leq \|z^{K+k+1} - z^{K+k} - v\|_2 \leq \mu^k \|(Q - I)z^K - p\|_2 \quad (37)$$

for all sufficiently large k .

2. **Sublinear convergence of the iterates.** *The normalized iterates converge to v at a $\Theta(\frac{1}{k})$ rate, i.e.,*

$$\frac{c_1}{k} \|(I - Q^\infty)p\|_2 \leq \left\| \frac{1}{k} z^{K+k} - v \right\|_2 \leq \frac{c_2}{k} (\|(Q - I)^\dagger(I - Q^\infty)p - z^K\|_2 + \|z^K\|_2) \quad (38)$$

for all sufficiently large k .

3. **Sublinear convergence of the average.** *The normalized average converges to v at a $\Theta(\frac{1}{k})$ rate, i.e.,*

$$\frac{c_3}{k+1} \|(I - Q)p\|_2 \leq \left\| \frac{2}{k(k+1)} \sum_{j=1}^k z^{K+j} - v \right\|_2 \leq \frac{c_4}{k+1} (\|(Q - I)^\dagger(I - Q^\infty)p - z^K\|_2 + \|z^K\|_2) \quad (39)$$

for all sufficiently large k .

A few remarks are in order.

Remark 5. *Although equations (38) and (39) state a bound for the normalized iterates and normalized average of PDHG started from z^K , this result implies the same asymptotic bounds (with slightly worse constants) for the normalized iterates and normalized averaged started from z^0 . Thus, the result concludes that under nondegeneracy, the difference of iterates converges much faster than the iterates and average. Furthermore, such conclusion is tight, as we provide both the upper and lower bounds for each sequence.*

Remark 6. *The proof of this result shows the same rates for Bilinear games. Recall that a Bilinear game is a minimax problem of the form*

$$\min_{x \in \mathbf{R}^n} \max_{y \in \mathbf{R}^m} c^\top x + \langle Ax, y \rangle - b^\top y.$$

For these problems, the updates of PDHG (Algorithm 1) take the form of (36) with $D = I$. Thus, all the arguments in the proof of this result follow with $K = 0$. In particular, this shows that the upper bound derived in Theorem 3 is tight for Bilinear minimax problems.

Remark 7. *Every LP problem that has an optimal solution furthermore has at least one primal-dual solution that satisfies strict complementarity [43, Theorem 2.35]. Consequently, every infeasible problem has at least one nondegenerate ray. Thus there exists at least one initial point z^0 , such that if PDHG is initialized at this point, then the iterates converge to the nondegenerate ray and thus enjoy linear convergence.*

Proof of Theorem 6. We start by making a few simplifying assumptions. First we assume $D = I$. If that is not the case, we can consider a submatrix of A where we trim out the columns indexed by $\{i \in n \mid D_{ii} = 0\}$. This has no effect in the end result since for all these indices $x_i^k = 0$, and thus it does not affect the nonzero

entries of x^k nor the entries of y^k . Furthermore, without loss of generality we assume that A is diagonal. This is because, otherwise, we can decompose the matrix Q as

$$Q = \begin{bmatrix} V & 0 \\ 0 & U \end{bmatrix} \begin{bmatrix} I & \eta \Sigma^\top \\ \tau \Sigma & I - 2\eta \tau \Sigma \Sigma^\top \end{bmatrix} \begin{bmatrix} V^\top & 0 \\ 0 & U^\top \end{bmatrix} \quad (40)$$

where $A = U \Sigma V^\top$ is the SVD decomposition of A . Then we can change the primal and dual basis using V and U , which is equivalent to applying orthogonal maps to the primal and dual variables; thus it does not alter the metric nor the algorithm. Therefore, we can change the basis to enforce this assumption. Notice that the columns and rows of Q can be further permuted so that it becomes a block diagonal matrix of the form

$$Q = \begin{bmatrix} B_1 & & \\ & \ddots & \\ & & B_{\min\{n,m\}} \\ & & & I \end{bmatrix} \quad \text{where} \quad B_i = \begin{bmatrix} 1 & -\eta \sigma_i \\ \tau \sigma_i & 1 - 2\eta \tau \sigma_i^2 \end{bmatrix}, \quad (41)$$

where σ_i is the i th singular value of A . Note that when $n > m$ (or $n < m$), we might also have block corresponding to an identity of size $n - m$ (resp. $m - n$), the arguments below can be easily extended to cover the identity block (as it follows from the rationale applied when $\sigma_i = 0$) and so we assume that $n = m$. Thus, from now on Q has the form (41) with $n = m$ and hence without the last identity block.

Now, we can compute closed-form formulas for the three certificate candidates (2a)-(2c). Let K be the smallest integer after which the PDHG iteration can be written as (36). By recursively expanding, we obtain

$$\begin{aligned} z^{k+1+K} &= Q z^{k+K} - p \\ &= Q^2 z^{k-1+K} - Qp - p \\ &\vdots \\ &= Q^{k+1} z^K - \sum_{i=0}^k Q^i p. \end{aligned} \quad (42)$$

If we take the difference between two consecutive iterates, this yields

$$z^{k+1+K} - z^{k+K} = Q^k (Q - I) z^K - Q^k p = Q^k ((Q - I) z^K - p). \quad (43)$$

On the other hand, summing the first k iterates (42) gives

$$\sum_{j=1}^k z^{j+K} = \sum_{j=1}^k Q^j z^K - \sum_{j=1}^k \sum_{l=0}^{j-1} Q^l p. \quad (44)$$

We will show that Q^k converges to a matrix Q^∞ . We define the matrix Q^∞ as a block diagonal matrix

$$Q^\infty = \begin{bmatrix} B_1^\infty & & \\ & \ddots & \\ & & B_n^\infty \end{bmatrix} \quad \text{where} \quad B_i^\infty = \begin{cases} I_2 & \text{if } \sigma_i = 0 \\ 0 & \text{otherwise} \end{cases}, \quad (45)$$

where I_2 is the 2-by-2 identity matrix. Since each block is independent of each other, we can analyze Q^k by studying B_i^k . A simple calculation reveals that the i th block has two eigenvalues of the form

$$\lambda_i^\pm = (1 - \eta \tau \sigma_i^2) \pm i (\eta \tau \sigma_i^2 (1 - \eta \tau \sigma_i^2))^{\frac{1}{2}}.$$

Taking the norm, we find $\rho(B_i) = |\lambda_i^\pm| = \sqrt{1 - \eta \tau \sigma_i^2}$. Then, we have that the iterated product of the i th block B_i^k converges to B_i^∞ . To see this, consider two cases:

Case 1. Assume that $\sigma_i > 0$. By assumption $0 < 1 - \eta\tau\sigma_i^2 < 1$ hence $B_i^k \rightarrow 0 = B_i^\infty$. This follows since the spectral radius $\rho(B_i^k) = (1 - \eta\tau\sigma_i^2)^{k/2} \rightarrow 0$ as k .

Case 2. Assume that $\sigma_i = 0$. Then $B_i^k = B_i = I = B_i^\infty$.

The matrix Q^∞ turns out to be the projection onto the kernel of $Q - I$ (that is, $Q^\infty(Q - I) = 0$). We use $\Delta \subseteq [n]$ to denote the set of indices such that $\sigma_i > 0$.

Differences. We start by analyzing the differences (43). To prove the upper bound, we expand

$$\begin{aligned} \|z^{K+k+1} - z^{K+k} - v\|_2 &= \|Q^k(((Q - I)z^K - p) - Q^\infty((Q - I)z^K - p))\|_2 \\ &\leq \|Q^k - Q^\infty\|_2 \|(Q - I)z^K - p\|_2 \\ &= \max_{i \in \Delta} \|B_i^k\|_2 \|(Q - I)z^K - p\|_2 \\ &\leq \mu^k \|(Q - I)z^K - p\|_2, \end{aligned}$$

where the first equality comes from taking the limit $k \rightarrow \infty$ of (43), the first inequality used the fact that $\|Wz\|_2 \leq \|W\|_2 \|z\|_2$ for any matrix W and vector z , and the last inequality follows since $\rho(B_i) = \lim \|B_i^k\|_2^{\frac{1}{k}}$ and hence for any $\mu \in (\rho(B_i), 1)$ we have that $\|B_i^k\|_2 \leq \mu^k$ for all large enough k .

Now we turn our attention to the lower bound. Given a vector z , we define z_i to be the vector with the two components that multiply the block B_i in the matrix-vector product Qz . Using the same expansion as above, we get

$$\begin{aligned} \|z^{K+k+1} - z^{K+k} - v\|_2^2 &= \|Q^k(((Q - I)z^K - p) - Q^\infty((Q - I)z^K - p))\|_2^2 \\ &= \sum_{i \in [n]} \|(B_i^k - B_i^\infty)((Q - I)z^K - p)_i\|_2^2 \\ &= \sum_{i \in \Delta} \|B_i^k((Q - I)z^K - p)_i\|_2^2 \\ &\geq \sum_{i \in \Delta} \sigma_{\min}(B_i^k)^2 \|((Q - I)z^K - p)_i\|_2^2 \\ &\geq \min_{i \in \Delta} \sigma_{\min}(B_i^k)^2 \sum_{i \in \Delta} \|((Q - I)z^K - p)_i\|_2^2 \\ &= \min_{i \in \Delta} \sigma_{\min}(B_i^k)^2 \|(I - Q^\infty)((Q - I)z^K - p)\|_2^2 \\ &= \left(\min_{i \in \Delta} \sigma_{\min}(B_i) \right)^{2k} \|(Q - I)z^K + (Q^\infty - I)p\|_2^2, \end{aligned}$$

where the second equality and the penultimate one use the block-diagonal structure of the matrix Q to decompose the norm of the matrix-vector product into orthogonal components, and the first inequality follows since B_i^k is a rank two matrix for any $i \in \Delta$.

Normalized iterates. We now turn our attention to the normalized iterates. The upper bound follows almost immediately from Theorem (3) if we consider the PDHG algorithm started at z^K . To show the bound with the theorem, it suffices to note that 1) in this case $v = -Q^\infty p$ and so we might pick $z^* = (Q - I)^\dagger (I - Q^\infty)p$, where $(I - Q)^\dagger$ is the pseudo inverse of $(I - Q)$; and 2) all the norms in finite dimensional spaces are equivalent so we can upper bound $\|\cdot\|_M \leq C \|\cdot\|_2$ for some constant $C > 0$. To see the first point note that

$$\begin{aligned} Qz^* - p - z^* &= -Q^\infty p \iff (Q - I)z^* = (I - Q^\infty)p \\ &\iff z^* = (Q - I)^\dagger (I - Q^\infty)p. \end{aligned}$$

Now we turn our attention to the lower bound. Just as before we analyze the dynamics of Q^k by studying the individual blocks B_i^k . We will use the following identity for blocks satisfying $\rho(B) < 1$:

$$\sum_{j=0}^k B^j = (I - B)^{-1}(I - B^{k+1}). \quad (46)$$

Recall that $\Delta = \{i \mid \sigma_i > 0\}$, which corresponds with blocks satisfying $\rho(B_i) < 1$. Additionally, recall that p_i is the vector with the two components of p that multiply the block B_i in the matrix-vector product Qp . Expanding we get

$$\begin{aligned}
\left\| v - \frac{1}{k} z^{K+k+1} \right\|^2 &= \left\| Q^\infty p + \frac{1}{k} \sum_{j=0}^k Q^j p - Q^{k+1} z^K \right\|^2 \\
&= \sum_{i \in [n]} \left\| B_i^\infty p_i + \frac{1}{k} \sum_{j=0}^k B_i^j p_i - \frac{1}{k} B_i^{k+1} z_i^K \right\|^2 \\
&= \frac{1}{k^2} \sum_{i \in \Delta} \left\| (I - B_i)^{-1} (I - B_i^{k+1}) p_i - B_i^{k+1} z_i^K \right\|^2 + \frac{1}{k^2} \sum_{i \notin \Delta} \|B_i^{k+1} z_i^K\|^2 \\
&\geq \frac{1}{k^2} \sum_{i \in \Delta} (\sigma_{\max}(I - B_i))^{-2} \left\| (I - B_i^{k+1}) p_i - (I - B_i) B_i^{k+1} z_i^K \right\|^2,
\end{aligned}$$

where for the last two equalities we used the fact that Q is block diagonal, and the last inequality follows since $I - B_i$ is invertible for $i \in \Delta$. Then, taking the minimum coefficient we get

$$\begin{aligned}
&\frac{1}{k^2} \sum_{i \in \Delta} \sigma_{\max}(I - B_i)^{-2} \left\| (I - B_i^{k+1}) p_i - (I - B_i) B_i^{k+1} z_i^K \right\|^2 \\
&\geq \frac{1}{k^2} \min_{j \in \Delta} \{ \sigma_{\max}(I - B_j)^{-2} \} \sum_{i \in \Delta} \left\| (I - B_i^{k+1}) p_i - (I - B_i) B_i^{k+1} z_i^K \right\|^2 \\
&= \frac{1}{k^2} \min_{j \in \Delta} \{ \sigma_{\max}(I - B_j)^{-2} \} \left\| (I - Q^{k+1}) p - (I - Q) Q^{k+1} z^K \right\|^2 \\
&\geq \frac{1}{2k^2} \min_{j \in \Delta} \{ \sigma_{\max}(I - B_j)^{-2} \} \left\| (I - Q^\infty) p \right\|^2,
\end{aligned}$$

where the last equality uses the fact that the matrices we handle are block diagonal, and the last line follows since $\|(I - Q^{k+1})p - (I - Q)Q^{k+1}z^K\| \rightarrow \|(I - Q^\infty)p\|$, and so the inequality holds for sufficiently large $k \geq 0$.

Normalized average. The upper bound follows from the exact same argument as the normalized iterates by using Theorem 3.

The lower bound for this case is slightly more intricate. We expand and apply (46)

$$\begin{aligned}
&\left\| v - \frac{2}{k(k+1)} \sum_{j=1}^k z^{K+j} \right\|^2 \\
&= \left\| Q^\infty p - \frac{2}{k(k+1)} \sum_{j=1}^k \sum_{l=0}^{j-1} Q^l p + \frac{2}{k(k+1)} \sum_{j=1}^k Q^j z^K \right\|^2 \\
&= \sum_{i \in [n]} \left\| B_i^\infty p_i - \frac{2}{k(k+1)} \sum_{j=1}^k \sum_{l=0}^{j-1} B_i^l p_i + \frac{2}{k(k+1)} \sum_{j=1}^k B_i^j z_i^K \right\|^2 \\
&= \sum_{i \in \Delta} \left\| \frac{2}{k(k+1)} \sum_{j=1}^k \sum_{l=0}^{j-1} B_i^l p_i + \frac{2}{k(k+1)} \sum_{j=1}^k B_i^j z_i^K \right\|^2 + \sum_{i \notin \Delta} \left\| \frac{2}{k(k+1)} \sum_{j=1}^k B_i^j z_i^K \right\|^2 \\
&= \frac{4}{k^2(k+1)^2} \sum_{i \in \Delta} \left\| (I - B_i)^{-1} \left(\sum_{j=1}^k (I - B_i^j) p_i + (I - B_i^{k+1}) z_i^K \right) - z_i^K \right\|^2 + \sum_{i \notin \Delta} \left\| \frac{2}{k(k+1)} \sum_{j=1}^k B_i^j z_i^K \right\|^2.
\end{aligned}$$

The second equality follows from the fact that Q and Q^∞ are block diagonal. Then, dropping the second sum and using the fact that $\|(I - B_i)^{-1}z\|_2 \geq \sigma_{\max}(I - B_i)\|z\|_2$ for all z and $i \in \Delta$, we can lower bound

$$\begin{aligned}
& \frac{4}{k^2(k+1)^2} \sum_{i \in \Delta} \left\| (I - B_i)^{-1} \left(\sum_{j=1}^k (I - B_i^j)p_i + B_i(I - B_i^k)z_i^K \right) \right\|^2 + \sum_{i \notin \Delta} \left\| \frac{2}{k(k+1)} \sum_{j=1}^k B_i^j z_i^K \right\|^2 \\
& \geq \frac{4}{k^2(k+1)^2} \sum_{i \in \Delta} \sigma_{\max}(I - B_i)^{-2} \left\| \sum_{j=1}^k (I - B_i^j)p_i + B_i(I - B_i^k)z_i^K \right\|^2 \\
& \geq \frac{4}{k^2(k+1)^2} \min_{i \in \Delta} \{ \sigma_{\max}(I - B_i)^{-2} \} \sum_{i \in \Delta} \left\| kp_i - (I - B_i)^{-1}B_i(I - B_i^k)p_i + B_i(I - B_i^k)z_i^K \right\|^2 \\
& \geq \frac{4}{k^2(k+1)^2} \min_{i \in \Delta} \{ \sigma_{\max}(I - B_i)^{-4} \} \sum_{i \in \Delta} \left\| k(I - B_i)p_i - B_i(I - B_i^k)p_i + (I - B_i)B_i(I - B_i^k)z_i^K \right\|^2 \\
& = \frac{4}{k^2(k+1)^2} \min_{i \in \Delta} \{ \sigma_{\max}(I - B_i)^{-4} \} \left\| k(I - Q)p - Q(I - Q^k)p + (I - Q)Q(I - Q^k)z^K \right\|^2 \\
& = \frac{4}{(k+1)^2} \min_{i \in \Delta} \{ \sigma_{\max}(I - B_i)^{-4} \} \left\| (I - Q)p - \frac{1}{k}Q(I - Q^k)p + \frac{1}{k}(I - Q)Q(I - Q^k)z^K \right\|^2 \\
& \geq \frac{2}{(k+1)^2} \min_{i \in \Delta} \{ \sigma_{\max}(I - B_i)^{-4} \} \left\| (I - Q)p \right\|^2,
\end{aligned}$$

where the last inequality follows for large enough k since $\left\| -\frac{1}{k}Q(I - Q^k)p + \frac{1}{k}(I - Q)Q(I - Q^k)z^K \right\|^2 \rightarrow 0$. This completes the proof. \square

6 Numerical experiments

In this section, we test numerically the proposed approach to check infeasibility using PDHG. For the experiments, we implemented PDHG for LP problems with the following primal and dual form

$$\begin{aligned}
& \text{minimize} & c^\top x & & \text{maximize} & b^\top y + l^\top r_+ - u^\top r_- \\
& \text{subject to} & Ax \geq b & & \text{subject to} & c - A^\top y = r \\
& & l \leq x \leq u & & & y \geq 0
\end{aligned} \tag{47}$$

where $b \in \mathbf{R}^m, l \in (\mathbf{R} \cup \{-\infty\})^n, u \in (\mathbf{R} \cup \{\infty\})^n, A \in \mathbf{R}^{m \times n}$ are given and $r_+ = \mathbf{proj}_{\mathbf{R}_+}(r)$ and $r_- = -\mathbf{proj}_{\mathbf{R}_+}(-r)$ are the projections of r onto the positive and negative orthant, respectively. We chose this form over the standard form **(P)**-**(D)** since it is algorithmically easier to reduce arbitrary LP problems to it. Given that the PDHG algorithm for this formulation generates iterates (x^k, y^k) , for our computations we generate r^k by finding the closest point to $c - A^\top y^k$ that makes the dual objective finite; i.e., $r^k = \mathbf{proj}_{C_v}(c - A^\top y^k)$ with C_v defined below. All the results proved in this work also apply to this form under suitable modifications of the statements.

For our experiments we use the *Netlib dataset of infeasible LP instances* [28]. We use this dataset to illustrate the different dynamics that PDHG exhibits. For all our experiments we measure statistics that quantify how close are the candidate iterates (2) to being approximate certificates of infeasibility.

Before we describe these statistics, let us define what we mean by approximate infeasibility certificates. The set of (exact) primal infeasibility certificates for (47) is given by all the vectors $(y, r) \in \mathbf{R}_+^m \times \mathbf{R}^n$ satisfying

$$b^\top y + l^\top r_+ - u^\top r_- > 0, \quad \text{and} \quad r = -A^\top y, \tag{48}$$

while the set of (exact) dual infeasibility certificates is given by all the vectors $x \in \mathbf{R}^n$ satisfying

$$c^\top x < 0, \quad x \in C_v, \quad \text{and} \quad Ax \geq 0, \tag{49}$$

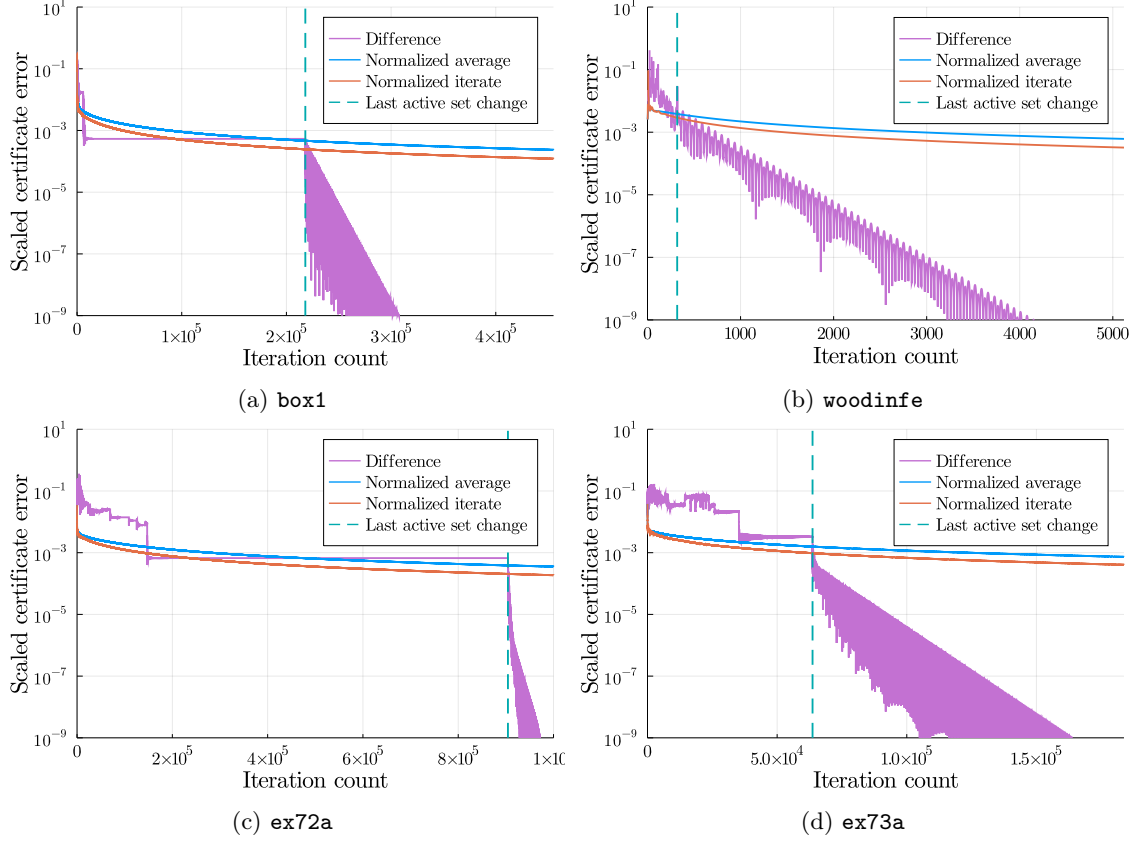


Figure 2: Scaled certificate error (52) for the three sequences defined in (2) for four instances of the Netlib infeasible dataset [28]. Vertical dotted lines denote the last observed active set change.

where the set C_v is given by

$$C_v = \left\{ x : \text{for } i \in [n], x_i \begin{cases} = 0 & \text{if } l_i, u_i \in \mathbf{R}, \\ \geq 0 & \text{if } l_i \in \mathbf{R}, u_i = \infty, \\ \leq 0 & \text{if } l_i = -\infty, u_i \in \mathbf{R}, \end{cases} \right\}.$$

We define an ε -approximate primal infeasibility certificate to be any point $(y, r) \in \mathbf{R}_+^m \times \mathbf{R}^n$ satisfying

$$b^\top y + l^\top r_+ - u^\top r_- > 0, \quad \text{and} \quad (b^\top y + l^\top r_+ - u^\top r_-)^{-1} \|r + A^\top y\|_\infty \leq \varepsilon. \quad (50)$$

Similarly we say that a point $x \in \mathbf{R}^m$ is an ε -approximate dual infeasibility certificate if it satisfies

$$c^\top x < 0, \quad \frac{1}{-c^\top x} \cdot \|x - \mathbf{proj}_{C_v}(x)\|_\infty \leq \varepsilon, \quad \text{and} \quad \frac{1}{-c^\top x} \cdot \|Ax - \mathbf{proj}_{\mathbf{R}_+^m}(Ax)\|_\infty \leq \varepsilon. \quad (51)$$

These definitions parallel the criteria to detect infeasibility used by SCS [60], a popular open-source solver.

Since all the instances in the Netlib infeasible data set are primal infeasible, we will only plot information about the dual components of the candidate certificates (2). To illustrate how close is each candidate to being a certificate we will plot

$$\frac{\|r^k + A^\top y^k\|_\infty}{b^\top y^k + l^\top r_+^k - u^\top r_-^k}. \quad (52)$$

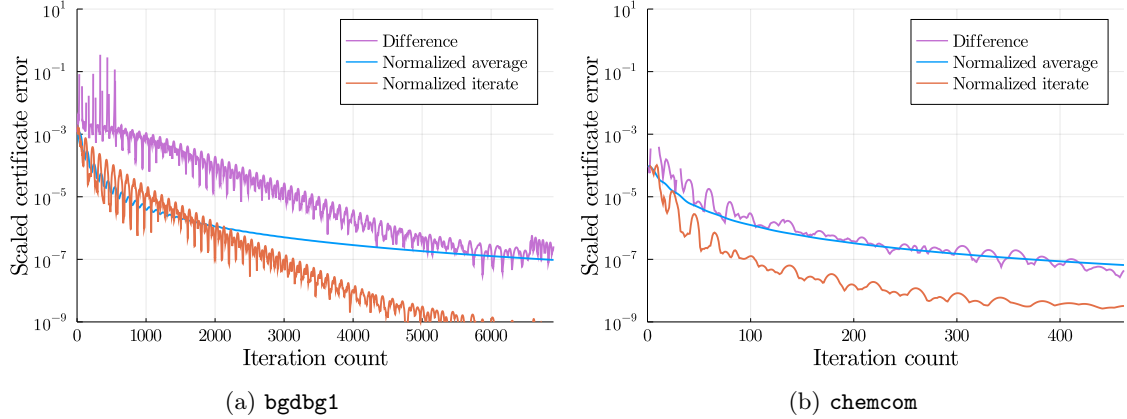


Figure 3: Scaled certificate error (52) for the three sequences defined in (2) for two instances of the Netlib infeasible dataset [28].

We call this quantity the *scaled certificate error*. If the objective term, i.e., the denominator, is negative at any iteration then we do not plot it for that iteration. For almost all the problems we consider (52) remains positive for almost all iterations.

Nondegeneracy in practice. Our first batch of experiments showcases the faster convergence of the difference of iterates in practice. We found empirically that for a subset of instances in the dataset, the difference of iterates (2a) detects infeasibility faster than the other two sequences. Based on the theory, we expect that for these instances, the difference exhibits eventual faster convergence. To test this claim, we run an experiment on four of these instances: `box1`, `woodinfe`, `ex73a` and `ext72a`.

Figure 2 displays the scaled certificate error (52) against the number of iterations for the four instances. For all of them, we can see a clear phase transition between a first stage of slow convergence and a second stage that displays linear convergence. This transition is unequivocally marked by the last change of the active set of the solution (also depicted in the figure). Notice, however, that the iteration number at which the active set is fixed might be large; the point at which this happens ranges among multiple orders of magnitude in our experiments.

Normalized iterates can be faster. Even if eventual identifiability holds, this might take a significant number of iterations. In these cases it might be beneficial to check infeasibility using the normalized iterated (2b) and the normalized average (2c). In this batch of experiments we run PDHG on `bgdbg1` and `chemcom`, the results are displayed in Figure 3. Just as before we plot the scaled certificate error against the number of iterations.

The normalized average is consistently slower at converging than the normalized iterates. This is most likely due to the fact that it retains a tail of initial iterates, which are far away from the infimal displacement vector. For both these problems, the difference takes at least twice the number of iterations than the normalized iterates to obtain a highly accurate certificate, i.e., $\varepsilon = 10^{-8}$. This suggests that solvers may benefit from checking infeasibility with both the normalized iterates (2b) and difference of iterates (2a).

7 Conclusions and future directions

In this work, we showed how to detect infeasibility of LP programs using the iterates of PDHG. The proposed approach is simple, doesn't require changing the logic of the algorithm, and has an almost negligible computational cost. Thus, it is suitable for large-scale LP problems. The ideas developed in this work pave the road for a number of research directions. Below we list some of these directions and superficially address their main difficulties.

It would be interesting to understand how to integrate our theoretical results with enhanced versions of

PDHG. For example, for practical applications it is often useful to have varying stepsizes η_k and τ_k that compensate for the different scales that the primal and dual problems might have [41]. In this case, there is no unique PDHG operator T applied at each iteration, rather a sequence of operators T_k and even the definition of v is not well-posed.

In a similar vein, we wonder if it is possible to speed up the convergence of some of the sequences in (2) using a restart scheme. This is known to be the true in the feasible case. Tackling this question would require designing a more involved method that restarts the PDHG iterates and the candidate certificates periodically. It is unclear if it is possible to have a scheme that works seamlessly for both the feasible and infeasible cases.

The proofs of most of our results rely on the form of the PDHG updates. For instance, one of our key arguments established a connection between the iterates of PDHG running on an infeasible problem and the iterates of PDHG running on an auxiliary feasible problem. It is natural to wonder if this is a phenomenon that only holds for PDHG or if it generalizes to other algorithms and/or other problem families such as quadratic and semidefinite programming.

Another potential research direction is the application of PDHG for solving the homogeneous self-dual embedding problem akin to [53]. A naive implementation of this idea would require doubling the dimension of the primal problem. The authors of [53] showed that for ADMM, this increase in the dimension doesn't affect the computational complexity of each iteration. At its core, their argument for this fact used the symmetry between primal and dual updates in ADMM; it is unclear how to extend this argument to PDHG given its asymmetric updates.

Acknowledgements

We would like to thank Brendan O'Donoghue, Oliver Hinder, and Warren Schudy for insightful conversations during the first stages of this work.

References

- [1] Tobias Achterberg. Conflict analysis in mixed integer programming. *Discrete Optimization*, 4(1):4 – 20, 2007. Mixed Integer Programming.
- [2] F. Al-Khayyal and J. Kyriasis. Finite convergence of algorithms for nonlinear programs and variational inequalities. *J. Optim. Theory Appl.*, 70(2):319–332, 1991.
- [3] JB Bailion, Ronald E Bruck, and Simeon Reich. On the asymptotic behavior of nonexpansive mappings and semigroups in banach spaces. *Houston Journal of Mathematics*, 4(1):1–9, 1978.
- [4] Goran Banjac, Paul Goulart, Bartolomeo Stellato, and Stephen Boyd. Infeasibility detection in the alternating direction method of multipliers for convex optimization. *Journal of Optimization Theory and Applications*, 183(2):490–519, 2019.
- [5] Kinjal Basu, Amol Ghoting, Rahul Mazumder, and Yao Pan. ECLIPSE: An extreme-scale linear program solver for web-applications. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 704–714, Virtual, 13–18 Jul 2020. PMLR.
- [6] H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. CMS Books in Mathematics/Ouvrages de Mathématiques de la SMC. Springer, Cham, second edition, 2017. With a foreword by Hedy Attouch.
- [7] Heinz H Bauschke, Patrick L Combettes, and D Russell Luke. Finding best approximation pairs relative to two closed convex sets in hilbert spaces. *Journal of Approximation theory*, 127(2):178–192, 2004.

- [8] Heinz H Bauschke and Walaa M Moursi. The Douglas–Rachford algorithm for two (not necessarily intersecting) affine subspaces. *SIAM Journal on Optimization*, 26(2):968–985, 2016.
- [9] Heinz H Bauschke, Xianfu Wang, and Liangjin Yao. General resolvents for monotone operators: characterization and extension. *arXiv preprint arXiv:0810.3905*, 2008.
- [10] Amir Beck. *First-Order Methods in Optimization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017.
- [11] J. M. Borwein and A. S. Lewis. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. Springer, 2006.
- [12] J.V. Burke. On the identification of active constraints. II. The nonconvex case. *SIAM J. Numer. Anal.*, 27(4):1081–1103, 1990.
- [13] J.V. Burke and J.J. Moré. On the identification of active constraints. *SIAM J. Numer. Anal.*, 25(5):1197–1211, 1988.
- [14] P.H. Calamai and J.J. Moré. Projected gradient methods for linearly constrained problems. *Math. Prog.*, 39(1):93–116, 1987.
- [15] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.
- [16] Antonin Chambolle and Thomas Pock. An introduction to continuous optimization for imaging. *Acta Numerica*, 25:161–319, 2016.
- [17] John W. Chinneck. Computer codes for the analysis of infeasible linear programs. *The Journal of the Operational Research Society*, 47(1):61–72, 1996.
- [18] G.B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, Princeton, 1963.
- [19] Damek Davis and Wotao Yin. Convergence rate analysis of several splitting schemes. In *Splitting methods in communication, imaging, science, and engineering*, pages 115–163. Springer, 2016.
- [20] Jim Douglas and Henry H Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American mathematical Society*, 82(2):421–439, 1956.
- [21] Dmitriy Drusvyatskiy and Adrian S Lewis. Optimality, identifiability, and sensitivity. *Mathematical Programming*, 147(1-2):467–498, 2014.
- [22] J.C. Dunn. On the convergence of projected gradient processes to singular critical points. *J. Optim. Theory Appl.*, 55(2):203–216, 1987.
- [23] Jonathan Eckstein and Dimitri P Bertsekas. On the Douglas—Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3):293–318, 1992.
- [24] Jonathan Eckstein and Gyorgy Matyasfalvi. Efficient distributed-memory parallel matrix-vector multiplication with wide or tall unstructured sparse matrices. *CoRR*, abs/1812.00904, 2018.
- [25] M.C. Ferris. Finite termination of the proximal point algorithm. *Math. Program.*, 50(3, (Ser. A)):359–366, 1991.
- [26] S.D. Flåm. On finite convergence and constraint identification of subgradient projection methods. *Math. Program.*, 57:427–437, 1992.
- [27] Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & mathematics with applications*, 2(1):17–40, 1976.

- [28] David M. Gay. *Netlib: infeasible linear programming test problems*, 2013 (accessed January 5, 2021). <http://netlib.org/lp/infeas/index.html>.
- [29] Andrew Gilpin, Javier Peña, and Tuomas Sandholm. First-order algorithm with $\mathcal{O}(\ln(1/\epsilon))$ convergence for ϵ -equilibrium in two-person zero-sum games. *Mathematical Programming*, 133(1):279–298, Jun 2012.
- [30] Roland Glowinski and A Marroco. Sur l’approximation, par éléments finis d’ordre un, et la résolution, par pénalisation-dualité d’une classe de problèmes de dirichlet non linéaires. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique*, 9(R2):41–76, 1975.
- [31] Bingsheng He, Yanfei You, and Xiaoming Yuan. On the convergence of primal-dual hybrid gradient algorithm. *SIAM Journal on Imaging Sciences*, 7(4):2526–2537, 2014.
- [32] Jourdain Lamperski, Robert M. Freund, and Michael J. Todd. An oblivious ellipsoid algorithm for solving a system of (in)feasible linear inequalities, 2020.
- [33] Guanghui Lan, Zhaosong Lu, and Renato D. C. Monteiro. Primal-dual first-order methods with $\mathcal{O}(1/\epsilon)$ iteration-complexity for cone programming. *Mathematical Programming*, 126(1):1–29, Jan 2011.
- [34] A. S. Lewis. Active sets, nonsmoothness, and sensitivity. *SIAM J. Optim.*, 13(3):702–725 (electronic) (2003), 2002.
- [35] Adrian S Lewis and Jingwei Liang. Partial smoothness and constant rank. *arXiv preprint arXiv:1807.03134*, 2018.
- [36] Jingwei Liang, Jalal Fadili, and Gabriel Peyré. Local convergence properties of Douglas–Rachford and alternating direction method of multipliers. *Journal of Optimization Theory and Applications*, 172(3):874–913, 2017.
- [37] Jingwei Liang, Jalal Fadili, and Gabriel Peyré. Local linear convergence analysis of primal–dual splitting methods. *Optimization*, 67(6):821–853, 2018.
- [38] Pierre-Louis Lions and Bertrand Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.
- [39] Yanli Liu, Ernest K Ryu, and Wotao Yin. A new use of Douglas–Rachford splitting for identifying infeasible, unbounded, and pathological conic programs. *Mathematical Programming*, 177(1-2):225–253, 2019.
- [40] Yura Malitsky. The primal-dual hybrid gradient method reduces to a primal method for linearly constrained optimization problems, 2019.
- [41] Yura Malitsky and Thomas Pock. A first-order primal-dual algorithm with linesearch. *SIAM Journal on Optimization*, 28(1):411–432, 2018.
- [42] István Maros. *Computational Techniques of the Simplex Method*. Kluwer Academic Publishers, 2003.
- [43] Richard Kipp Martin. *Large Scale Linear and Integer Optimization: A Unified Approach*. Springer US, Boston, MA, 1999.
- [44] R. Mifflin and C. Sagastizábal. Primal-dual gradient structured functions: second-order results; links to epi-derivatives and partly smooth functions. *SIAM J. Optim.*, 13(4):1174–1194 (electronic), 2003.
- [45] R. Mifflin and C. Sagastizábal. \mathcal{VU} -smoothness and proximal point results for some nonconvex functions. *Optim. Methods Softw.*, 19(5):463–478, 2004.
- [46] R. Mifflin and C. Sagastizábal. A \mathcal{VU} -algorithm for convex minimization. *Math. Program.*, 104(2-3, Ser. B):583–608, 2005.

- [47] Cesare Molinari, Jingwei Liang, and Jalal Fadili. Convergence rates of forward–douglas–rachford splitting method. *Journal of Optimization Theory and Applications*, 182(2):606–639, 2019.
- [48] Walaa M Moursi. *The Douglas–Rachford operator in the possibly inconsistent case: static properties and dynamic behaviour*. PhD thesis, University of British Columbia, 2016.
- [49] Arkadi Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- [50] Yu. Nesterov. Subgradient methods for huge-scale optimization problems. *Mathematical Programming*, 146(1):275–297, Aug 2014.
- [51] Daniel O’Connor and Lieven Vandenbergh. On the equivalence of the primal-dual hybrid gradient method and Douglas–Rachford splitting. *Mathematical Programming*, 179(1):85–108, Jan 2020.
- [52] Brendan O’Donoghue. Operator splitting for a homogeneous embedding of the monotone linear complementarity problem. *arXiv preprint arXiv:2004.02177*, 2020.
- [53] Brendan O’Donoghue, Eric Chu, Neal Parikh, and Stephen Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 169(3):1042–1068, 2016.
- [54] A Pazy. Asymptotic behavior of contractions in hilbert space. *Israel Journal of Mathematics*, 9(2):235–240, 1971.
- [55] T. Pock and A. Chambolle. Diagonal preconditioning for first order primal-dual algorithms in convex optimization. In *2011 International Conference on Computer Vision*, pages 1762–1769, 2011.
- [56] Michael JD Powell. Algorithms for nonlinear constraints that use lagrangian functions. *Mathematical programming*, 14(1):224–248, 1978.
- [57] Arvind U Raghunathan and Stefano Di Cairano. Infeasibility detection in alternating direction method of multipliers for convex quadratic programs. In *53rd IEEE Conference on Decision and Control*, pages 5819–5824. IEEE, 2014.
- [58] James Renegar. Accelerated first-order methods for hyperbolic programming. *Mathematical Programming*, 173(1):1–35, Jan 2019.
- [59] R. T. Rockafellar. *Convex analysis*. Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J., 1970.
- [60] Shu-Chung Shi. Semi-continuités génériques de multi-applications. *C.R. Acad. Sci. Paris (Série A)*, 293:27–29, 1981.
- [61] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd. OSQP: an operator splitting solver for quadratic programs. *Mathematical Programming Computation*, 12(4):637–672, 2020.
- [62] M.J. Todd. Detecting infeasibility. In Felipe Cucker, Ron DeVore, Peter Olver, and Endre Süli, editors, *Foundations of Computational Mathematics, Minneapolis 2002*, London Mathematical Society Lecture Note Series, page 157–192. Cambridge University Press, 2004.
- [63] Yuhsiang M. Tsai, Terry Cojean, and Hartwig Anzt. Sparse linear algebra on AMD and NVIDIA GPUs – the race is on. In Ponnuswamy Sadayappan, Bradford L. Chamberlain, Guido Juckeland, and Hatem Ltaief, editors, *High Performance Computing*, pages 309–327, Cham, 2020. Springer International Publishing.
- [64] S.J. Wright. Identifiable surfaces in constrained optimization. *SIAM J. Control Optim.*, 31:1063–1079, July 1993.

A Additional proofs

A.1 Proof of Proposition 1

Proof. Assume that $(z^{k+1} - z^k) \rightarrow v$. Fix $\varepsilon > 0$. Our goal is to show that for all k large enough $\|z^k/k - v\| \leq \varepsilon$. Due to convergence, there exist $K_1 \in \mathbf{N}$ such that for all $k \geq K_2$ we have $\|z^{k+1} - z^k - v\| \leq \varepsilon/3$. Define

$$B := \max_{k \leq K_1} \|z^{k+1} - z^k - v\|,$$

let $K_2 \in \mathbf{N}$ be such that for all $k \geq K_2$ we get $K_1 B/k \leq \varepsilon/3$, and let $K_3 \in \mathbf{N}$ be such that if $k \geq K_3$ then $\|z^0\|/k \leq \varepsilon/3$. Then, for any $k \geq \max\{K_1, K_2, K_3\}$ we have

$$\begin{aligned} \left\| \frac{z^k}{k} - v \right\| &\leq \left\| \frac{1}{k}(z^k - z^0) - v \right\| + \frac{1}{k}\|z^0\| \\ &= \left\| \frac{1}{k} \sum_{j=1}^k (z^j - z^{j-1}) - v \right\| + \frac{1}{k}\|z^0\| \\ &\leq \frac{1}{k} \sum_{j=1}^k \|(z^j - z^{j-1}) - v\| + \frac{1}{k}\|z^0\| \\ &\leq \frac{K_1}{k} B + \frac{1}{k} \sum_{j=K_1}^k \|(z^j - z^{j-1}) - v\| + \frac{1}{k}\|z^0\| \\ &\leq \frac{2}{3}\varepsilon + \frac{1}{3k} \sum_{j=K_1}^k \varepsilon \leq \varepsilon. \end{aligned}$$

This proves the first statement.

Now, assume that $\frac{z^k}{k} \rightarrow v$, and fix $\varepsilon > 0$. Just as before define $K_1 \in \mathbf{N}$ to be such that for all $\|z^k/k - v\| \leq \varepsilon/2$, define the constant $B = \max_{k \leq K_1} \|z^k/k - v\|$, and let K_2 be such that $B(K_1 + 1)K_1/((K_2 + 1)K_2) \leq \varepsilon/2$. Then, we have that for any $k \geq \max\{K_1, K_2\}$,

$$\begin{aligned} \left\| \frac{2}{(k+1)k} \sum_{j=1}^k z^j - v \right\| &= \frac{2}{(k+1)k} \left\| \sum_{j=1}^k z^j - \frac{k(k+1)}{2} v \right\| \\ &= \frac{2}{(k+1)k} \left\| \sum_{j=1}^k (z^j - jv) \right\| \\ &\leq \frac{2}{(k+1)k} \sum_{j=1}^k j \left\| \frac{z^j}{j} - v \right\| \\ &\leq \frac{(K_1 + 1)K_1}{(k+1)k} B + \frac{2}{(k+1)k} \sum_{j=K_1}^k j \left\| \frac{z^j}{j} - v \right\| \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} \left(\frac{2}{(k+1)k} \sum_{j=K_1}^k j \right) \leq \varepsilon. \end{aligned}$$

□

B Counterexamples

Example 2 (Differences don't converge, but normalized iterates do). Consider the sequence $(z^k) \subseteq \mathbf{R}^2$ generated by applying a 90° counterclockwise rotation repeatedly starting from $z^0 = e_1 = (1, 0)$. Thus, the sequence cycles as

$$z^1 = (0, 1), z^2 = (-1, 0), z^3 = (0, -1), z^4 = (1, 0), \dots$$

For this example, the differences of iterates $z^{k+1} - z^k$ also cycle among four possibilities and do not converge. Nonetheless, since the iterates are bounded $\frac{1}{k}z^k \rightarrow 0$.

Example 3 (Normalized iterates diverge, but normalized averages converge). Consider the sequence $(z^k) \subseteq \mathbf{R}^2$ given by $z^k = (-1)^k k^{\frac{3}{2}}$ with $k \in \mathbf{N}$. Then, it is clear that $|z^k|/k > \sqrt{k}$, and so the normalized iterates diverge. On the other hand, notice that

$$\frac{2}{(k+1)}\bar{z}^k = \sum_{j=1}^k (-1)^j \frac{2k^{\frac{1}{2}}}{k+1},$$

and it is easy to show that this series converges using the Leibniz Test.

Example 4 (Nonexpansive operator with divergent z_ϵ). Let $T : \mathbf{R} \rightarrow \mathbf{R}$ given by

$$T(z) = z + f(z) \quad \text{where} \quad f(z) = \begin{cases} \exp(-z^2) + 1 & \text{if } z > 0, \text{ or} \\ 2 & \text{otherwise.} \end{cases} \quad (53)$$

Since the derivative of T is bounded by 1, we get that T is a nonexpansive operator. Furthermore, $\text{range}(T-I) = \text{range}(f) = (1, 2]$, and so $v = 1$. If we define z_ϵ to be a point such that $|v - (T-I)(z_\epsilon)| \leq \epsilon$, we see that $z_\epsilon > \Omega\left(\log\left(\frac{1}{\epsilon}\right)^{\frac{1}{2}}\right)$, and thus it diverges as $\epsilon \rightarrow 0$.