



Faster first-order primal-dual methods for linear programming using restarts and sharpness

David Applegate^{1,2} · Oliver Hinder² · Haihao Lu³ · Miles Lubin¹

Received: 26 August 2021 / Accepted: 28 September 2022 / Published online: 19 October 2022
© Springer-Verlag GmbH Germany, part of Springer Nature and Mathematical Optimization Society 2022

Abstract

First-order primal-dual methods are appealing for their low memory overhead, fast iterations, and effective parallelization. However, they are often slow at finding high accuracy solutions, which creates a barrier to their use in traditional linear programming (LP) applications. This paper exploits the sharpness of primal-dual formulations of LP instances to achieve linear convergence using restarts in a general setting that applies to alternating direction method of multipliers (ADMM), primal-dual hybrid gradient method (PDHG) and extragradient method (EGM). In the special case of PDHG, without restarts we show an iteration count lower bound of $\Omega(\kappa^2 \log(1/\epsilon))$, while with restarts we show an iteration count upper bound of $O(\kappa \log(1/\epsilon))$, where κ is a condition number and ϵ is the desired accuracy. Moreover, the upper bound is optimal for a wide class of primal-dual methods, and applies to the strictly more general class of sharp primal-dual problems. We develop an adaptive restart scheme and verify that restarts significantly improve the ability of PDHG, EGM, and ADMM to find high accuracy solutions to LP problems.

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

✉ Oliver Hinder
ohinder@pitt.edu

David Applegate
dapplegate@google.com

Haihao Lu
haihao.lu@chicagobooth.edu

¹ Google Research, New York, USA

² University of Pittsburgh, Pittsburgh, USA

³ University of Chicago Booth School of Business, Chicago, USA

1 Introduction

Linear programming (LP) is a fundamental tool in operations research, with applications in transportation, scheduling, inventory control, revenue management [4, 10, 16, 35, 50, 53]. An LP problem with n variables and m constraints in standard form is

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & c^\top x \\ \text{subject to:} \quad & Ax = b \\ & x \geq 0, \end{aligned}$$

where \mathbb{R} denotes the set of real numbers.

Currently, most practitioners solve LP problems with either simplex [18] or interior-point methods [41]. There are excellent reasons for this, namely that both methods typically find highly accurate solutions in a reasonable amount of time, and reliable implementations are widely available. While the two methods are quite different, both exactly solve linear system subproblems using factorizations (i.e., matrix inversion). These factorizations can be extremely fast, even for large problems. However, there are some drawbacks of relying on methods that use factorizations. Factorization speed heavily depends on the sparsity pattern of the linear systems, and in some examples, can be very slow. Moreover, for large enough problems, factorizations will run out of memory, even when the original problem data can fit into memory. Finally, factorizations are challenging to parallelize or distribute across multiple machines.

For this reason, there is a recent interest in developing methods for LP that forgo factorizations and instead use matrix–vector multiplies [5, 7, 8, 32, 45, 49]. Methods built on matrix–vector multiplication have many computationally appealing aspects. For example, they are readily parallelized or distributed across multiple machines. Moreover, their memory footprint is almost the same as the original problem data, and each matrix–vector multiply is reliably fast. Because of these fundamental differences, Nesterov [58] categorizes optimization methods that use factorizations as capable of handling *medium-scale* problems, and those that instead use matrix–vector multiplies as capable of handling *large-scale* problems.

A major reason why methods based on matrix–vector multiplies remain unsuitable replacements for factorization-based methods for LP is that they often struggle to obtain moderate to high accuracy solutions in a reasonable time-frame [32]. This paper attempts to tackle this issue in an important setting: first-order primal-dual methods for LP. Primal-dual methods solve:

$$\min_{x \in X} \max_{y \in Y} \mathcal{L}(x, y), \quad (1)$$

where $\mathcal{L}(x, y)$ is differentiable, convex in x and concave in y , and X, Y are closed convex sets. For notational convenience, we denote $z = (x, y)$ as the primal-dual variables and $Z = X \times Y$ as the feasible region. Furthermore, denote

$$F(z) = \begin{pmatrix} \nabla_x \mathcal{L}(x, y) \\ -\nabla_y \mathcal{L}(x, y) \end{pmatrix}$$

as the first-order derivatives of $\mathcal{L}(x, y)$. Throughout the paper, we assume the feasible region of (1) is non-empty and that (1) has a bounded optimal value, which are standard assumptions in the convex optimization literature [61].¹ Well-known primal-dual methods include proximal point method (PPM) [74], extragradient method (EGM) [43, 77], primal-dual hybrid gradient (PDHG) [14], and alternating direction method of multipliers (ADMM) [21]. We can solve LP problems with primal-dual methods by writing them in the form:

$$\min_{x \geq 0} \max_{y \in \mathbb{R}^m} \mathcal{L}(x, y) = c^\top x + y^\top b - y^\top Ax. \quad (2)$$

To obtain high accuracy solutions for (2), our key idea is to use restarts and sharpness. A restart feeds the output of an algorithm back in as input; a restart scheme will restart an algorithm whenever certain criteria are met. Sharpness presents a lower bound on the objective function growth in terms of the distance to the optimal solution set. Restart schemes and sharpness conditions are well-studied in the context of unconstrained minimization [64, 73, 78]. In this context they can improve the theoretical and practical convergence of minimization algorithms to high accuracy solutions. We analyze restarts in a different context: primal-dual methods for solving sharp primal-dual problems, such as those derived from LP, obtaining similar improvements.

1.1 Contributions

The contributions of the paper can be summarized as follow:

- We introduce a sharpness condition of primal-dual problems, and show LP problems are sharp.
- We propose restarted algorithms for solving sharp primal-dual problems, and present their linear convergence rates. Their linear convergence rate matches the lower bound for a general class of primal-dual methods.
- We analyze primal-dual hybrid gradient without restarts and show that it is slower than its restarted variant on bilinear problems.
- We present an adaptive restart scheme that requires no hyperparameter searching over restart lengths. Numerical experiments showcase the effectiveness of our scheme.

While not a contribution of this paper, a subsequent computational study [5] evaluates a method named PDLP, which builds on top of the theory in the present paper by implementing a restarted version of PDHG with various additional enhancements. The study shows that adaptive restarts contribute significantly to PDLP's ability to solve a wide variety of benchmark LP problems to high levels of accuracy (see Section C.2 in [5]). Applegate et al. [5] includes an open source Julia prototype of PDLP.²

¹ A robust algorithm for LP would need to detect violations of this assumption, i.e., when the problem is infeasible or unbounded. We refer readers to [6] to understand the behavior of primal-dual methods when applied to infeasible or unbounded LP problems.

² <https://github.com/google-research/FirstOrderLp.jl>.

More recently, a C++ version of PDLP was released within the open source OR-Tools library.³

There are two concepts we introduce to facilitate our implementation and analysis: the normalized duality gap and sharpness.

1.2 Normalized duality gap

The traditional metric for measuring solution quality for primal-dual problems is the primal-dual gap [14, 57] defined as

$$\max_{\hat{z} \in Z} \{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)\} \quad (3)$$

where for brevity we denote $(\hat{x}, \hat{y}) = \hat{z}$. A solution to (1) has zero primal-dual gap. However, when the feasible region is unbounded, such as in LP (2), this quantity can be infinite. For this reason, for unbounded problems, we cannot provide convergence guarantees on this quantity. This makes it a poor progress metric. To overcome this issue we introduce the normalized duality gap:

$$\rho_r(z) := \frac{\max_{\hat{z} \in W_r(z)} \{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)\}}{r}, \quad (4a)$$

where $W_r(z) := \{\hat{z} \in Z \mid \|z - \hat{z}\| \leq r\}$ is the ball centered at z with radius $r \in (0, \infty)$ intersected with the set Z , and $\|\cdot\|$ is a carefully selected semi-norm on Z . If $r = 0$ then for completeness we define⁴:

$$\rho_0(z) := \limsup_{r \rightarrow 0^+} \rho_r(z). \quad (4b)$$

Similar to primal-dual gap, we show the normalized duality gap attains zero if and only if it is evaluated at a solution to (1). Unlike the primal-dual gap, if $\|\cdot\|$ is a norm, then the normalized duality gap $\rho_r(z)$ is always finite at any point $z \in Z$ with $r \in (0, \infty)$.

The normalized duality gap plays two roles in our work. The first is that it is used as a potential function in our restart scheme. In particular, our adaptive restarts are triggered by constant factor reductions in the normalized duality gap. Secondly, it forms the foundations of our analysis for restart schemes and is directly involved in our definition of sharpness for primal-dual problems.

1.3 Sharpness for primal-dual problems

Sharpness, first introduced by Polyak [68], has been an important concept and a useful analytical tool in convex minimization. For a given function f , we say f is α -sharp if

³ <https://developers.google.com/optimization>.

⁴ The RHS of (4b) is well-defined because \limsup always exists [76, Section 5.3]. Later we show (Proposition 5) that $\rho_r(z)$ is monotonically non-increasing in $r \in (0, \infty)$ for fixed z which means that $\rho_0(z) = \limsup_{r \rightarrow 0^+} \rho_r(z) = \lim_{r \rightarrow 0^+} \rho_r(z) < \infty$.

it holds for any $z \in Z$ that

$$f(z) - f^* \geq \alpha \mathbf{dist}(z, Z^*), \quad (5)$$

where Z^* is the set of minima of the function f , f^* is the optimal function value, and $\mathbf{dist}(z, Z) := \inf_{\hat{z} \in Z} \|z - \hat{z}\|$. When f is Lipschitz continuous, convex and sharp, it is well-known that restarted sub-gradient descent obtains linear convergence. In contrast, without sharpness it obtains a slow sublinear rate [78]. In this work, we generalize sharpness from minimization problems to primal-dual problems (1). Informally, we say a primal-dual problem is sharp if and only if the normalized duality gap $\rho_r(z)$ is sharp as per (5) with $f = \rho_r$. This condition is satisfied by LP problems, and the sharpness constant is proportional to the inverse of the Hoffman constant [38] of the matrix that forms the KKT conditions.

1.4 Notation

Let \mathbb{N} denote the set of natural numbers (starting from one). Let $\log(\cdot)$ and $\exp(\cdot)$ refer to the natural log and exponential function respectively. Let $\sigma_{\min}(\cdot)$, $\sigma_{\min}^+(\cdot)$, $\sigma_{\max}(\cdot)$ denote the minimum singular value, minimum nonzero singular value and maximum singular value of a matrix. Let $\|z\|_2 = \sqrt{\sum_j z_j^2}$ be the ℓ_2 norm of z . For a bounded set $S \subseteq Z$, $\mathbf{diam}(S)$ denotes the diameter of S , i.e., $\mathbf{diam}(S) := \sup_{s, s' \in S} \|s - s'\|$. The remainder of this section discusses related literature. In Sect. 2, we present a unified viewpoint of the sublinear rate of classic primal-dual algorithms. In Sect. 3, we discuss the basic properties and examples of sharp primal-dual problems. In Sect. 4, we present two restart schemes: fixed frequency restarts and adaptive restarts, and prove the linear convergence of both algorithms. In Sect. 5, we consider the special case of bilinear problems. In this setting, we show restarts obtain optimal convergence rates and that standard PDHG is provably slower than restarted PDHG for solving bilinear problems. In Sect. 6, we present efficient algorithms to compute the normalized duality gap. In Sect. 7, we present numerical experiments that showcase the effectiveness of our proposed restart schemes.

1.5 Related literature

1.5.1 Convex–concave primal-dual problems

The study of convex-concave primal-dual problems has a long history. Convex-concave primal-dual problems are a special case of monotone variational inequalities. In his seminal work [74], Rockafellar studied PPM for solving monotone variational inequalities. Later on, Tseng [77] shows that both PPM and EGM have a linear convergence rate for solving variational inequalities when certain complicated conditions are satisfied, and these conditions are satisfied for solving strongly-convex-strongly-concave primal-dual problems. In 2004, Nemirovski [57] proposed the Mirror Prox algorithm (a special selection of the prox function recovers EGM), which first showed

that EGM has $O(\frac{1}{\epsilon})$ sub-linear convergence rate for solving convex-concave primal-dual problems over a bounded and compact set.

There are several works that study the special case of (1) when the primal-dual function has bilinear interaction terms, i.e., $\mathcal{L}(x, y) = f(x) + x^\top B y - g(y)$ where $f(\cdot)$ and $g(\cdot)$ are both convex functions. Algorithms for solving these bilinear interaction problems include Douglas–Rachford splitting (a special case is Alternating Direction Method of Multipliers (ADMM)) [21, 22] and Primal-Dual Hybrid Gradient Method (PDHG) [14].

The linear convergence of primal-dual methods is widely studied. Daskalakis et al. [19] studies the Optimistic Gradient Descent Ascent (OGDA), designed for training GANs, and shows that OGDA converges linearly for bilinear problems $\mathcal{L}(x, y) = x^\top A y$ with a full rank matrix A . Mokhtari et al. [56] shows that OGDA and EGM both approximate PPM (indeed, the fact that EGM is an approximation to PPM was first shown in Nemirovski’s earlier work [57]). They further show that these three algorithms have a linear convergence rate when $\mathcal{L}(x, y)$ is strongly-convex strongly-concave or when $\mathcal{L}(x, y) = x^\top A y$ is bilinear with a square and full rank matrix. Lu [51] presents an ODE framework to analyze the dynamics of unconstrained primal-dual algorithms, yielding tight conditions under which different algorithms exhibit linear convergence, including bilinear problems with nonsingular A . Eckstein and Bertsekas [23] use a variant of ADMM to solve LP problems and show the linear convergence of their method. More recently, [44, 46] show that many primal-dual algorithms, including PDHG and ADMM, have eventual linear convergence when applied to LP problems satisfying a non-degeneracy condition. Very recently, for PDHG on LP problems, linear convergence is established without requiring a non-degeneracy assumption [1].

However, not all linear convergence rates are equal. For example, it is well known that both gradient descent and accelerated gradient descent (AGD) obtain linear convergence on smooth, strongly convex functions. However, the gradient descent bound is $O(L/\theta \log(1/\epsilon))$ and accelerated gradient descent’s bound is $O(\sqrt{L/\theta} \log(1/\epsilon))$ where L is the smoothness and θ is strong convexity parameter [61]. Therefore, accelerated gradient descent has a better dependence on the condition number (i.e., the ratio between θ and L) by a square root factor. This significant improvement is seen in practice too. A similar phenomena occurs for PDHG where we show for bilinear problems, the dependency on the condition number improves by a square root with introduction of restarts. We also show that, in practice, restarts improve the performance of PDHG. A more detailed discussion of this issue and related literature appears in Sect. 5.

1.5.2 Sharpness conditions

The concept of sharpness of a minimization problem was first proposed by Polyak [68] in 1970s. The original work by Polyak assumes the optimal solution set is a singleton. This work was then generalized by Burke and Ferris [12] to include the possibility of multiple optima. The early work on sharp minimization problems focused on analyzing the finite termination of certain algorithms, for example, projected gradient descent [69], proximal-point method [26], and Newton’s method [13]. Sharpness of monotone variational inequalities with bounded feasible regions was studied by

Marcotte and Zhu [54], and they focused on the finite convergence of a two-stage algorithm. Notice the primal-dual problem (1) is a special case of a variational inequality (VI) [36]. Indeed, when the feasible region is bounded, the sharpness of primal-dual problems we propose, translating to VI form, is equivalent to the sharpness of variational inequalities up to a constant. Our results generalize the concept for unbounded regions (when translating to VI). More recently, sharpness plays an important role for developing faster convergence rates for first-order methods. In particular, linear convergence for subgradient (restarted) descent on sharp non-smooth functions [78] and sharp non-convex (weakly-convex) minimization [20]. Finally, similar to strong convexity, sharpness provides a lower bound on the optimality gap in distance to optimality. Both sharpness and strong convexity can be viewed as error bound conditions with different parameters [73].

1.5.3 Restart schemes

Restarting is a powerful technique in continuous optimization that can improve the practical and theoretical convergence of a base algorithm [67]. The approach does not require modifications to the base algorithm, and thus is in general easy to incorporate within standard implementations. Recently, there have been extensive work on this technique, in smooth convex optimization [64, 73], non-smooth convex optimization [27, 78], and stochastic convex optimization [40, 47, 75]. Many of the previous works require estimating algorithmic parameters such as the strong convexity constant. This estimation often requires computationally intensive hyperparameter searches. An important topic is to avoid this estimation by developing adaptive methods, which were well-studied for accelerated gradient descent [2, 25, 30, 48, 60, 64]. In this paper, we present two restart schemes for primal-dual methods: fixed frequency restarts and adaptive restarts, where the former needs an estimation of the sharpness constant, while the latter does not.

1.5.4 Hoffman constant

The Hoffman constant [38], upper bounds the ratio between the distance to a non-empty polyhedron and the corresponding constraint violation. More formally, for a given non-empty polyhedron $P = \{z \in \mathbb{R}^{n+m} \mid Kz \geq h\}$, the Hoffman constant $H(K)$ satisfies $\mathbf{dist}(z, P) \leq H(K) \|(h - Kz)^+\|$ for all z . The Hoffman constant can also be viewed as the inverse of the sharpness constant of the function $f(z) = \|(h - Kz)^+\|$, namely,

$$f(z) - f^* \geq \frac{1}{H(K)} \mathbf{dist}(z, Z^*)$$

by noticing the minimal objective is $f^* = 0$ when the polyhedron is non-empty. Furthermore, there is a natural connection between the Hoffman constant and Renegar's condition number [71, 72], which measures the distance to ill-posedness of a system of linear inequalities. The Hoffman constant also characterizes the inherent difficulty

to solve a system of linear inequalities, and plays an important role in establishing the convergence properties of many optimization methods [34, 52, 66, 72].

1.5.5 Concurrent work

After posting the first version of our manuscript, we noticed a concurrent and independent work [24]. This work presents a new regularity condition of primal-dual problems called quadratic error bound of the smoothed gap (QEB). Furthermore, [24] shows that LP problems satisfy QEB, and restarted PDHG has linear convergence under QEB condition if the QEB parameter is known to the user. Our condition is different from the QEB condition as it involves the sharpness of the normalized duality gap. Contributions of our paper that do not appear in [24] include: (i) we present an adaptive scheme that achieves the optimal convergence rate without knowing the sharpness parameter; (ii) we provide a lower bound that proves our restarted algorithms achieve the optimal linear convergence rate under our sharpness condition, whereas it is unclear whether the algorithms presented in [24] are optimal in their setting; (iii) our analysis applies to a variety of primal-dual algorithms including PPM, EGM, ADMM and PDHG, whereas [24] focuses on PDHG. On the other hand, contributions of [24] that do not appear in our paper include that (i) our sharpness condition is tied to LP, while the QEB condition is possibly more general; (ii) we show the linear convergence of the current iterate only in the bilinear setting whereas [24] analyzes the linear convergence of the current iterate under the QEB condition.

2 A unified viewpoint of the sublinear ergodic convergence rate of primal-dual algorithms

While there have been recent works on studying the last iterations of primal-dual algorithms [17], many previous convergence bounds for primal-dual problems (1) are with respect to the average over the algorithm's iterates [15, 37, 57]. This bound is known as the algorithm's ergodic convergence rate. In this section, we discuss a unified viewpoint of the sublinear ergodic convergence rate for classic primal-dual algorithms, which we will later use to develop the analysis for our restart schemes.

We consider a generic class of primal-dual algorithms for primal-dual problems starting from $t = 0$ with iterate update

$$z^{t+1}, \hat{z}^{t+1} = \text{PrimalDualStep}(z^t, \eta), \quad (6)$$

where z^t is the *iterate solution* at time t , z^{t+1} is the *iterate solution* at time $t + 1$, \hat{z}^{t+1} is the *target solution* at time $t + 1$, and $\eta \in (0, \infty)$ is the step size of the algorithm. As we will see later, the convergence guarantee of some primal-dual algorithms may not be on the iterate solution z^t ; hence, we introduce the target solution \hat{z}^t . The output of the algorithm is a function of target sequence $\hat{z}^1, \dots, \hat{z}^t$, and for the standard ergodic rate it is the average of target solutions, namely, $\bar{z}^t = \frac{1}{t} \sum_{j=1}^t \hat{z}^j$. Here we discuss four classic primal-dual algorithms that can be represented as (6):

- (Proximal point method) Proximal point method (PPM) [74] is a classic algorithm for solving (1). The update can be written as:

$$\begin{aligned} \hat{z}^{t+1} = z^{t+1} &= (x^{t+1}, y^{t+1}) \in \arg \min_{x \in X} \max_{y \in Y} \mathcal{L}(x, y) \\ &+ \frac{1}{2\eta} \|x - x^t\|_2^2 - \frac{1}{2\eta} \|y - y^t\|_2^2. \end{aligned} \tag{7}$$

- (Extragradient method) Extragradient method [77] (a special case of mirror prox [57]) is another classic method for solving (1). It is an approximation of PPM [57]. The update can be written as:

$$\begin{aligned} \hat{z}^{t+1} &\in \operatorname{argmin}_{z \in Z} \left\{ F(z^t)^\top z + \frac{1}{2\eta} \|z - z^t\|_2^2 \right\} \\ z^{t+1} &\in \operatorname{argmin}_{z \in Z} \left\{ F(\hat{z}^{t+1})^\top z + \frac{1}{2\eta} \|z - z^t\|_2^2 \right\}. \end{aligned} \tag{8}$$

For this method, we assume $F(z)$ is L -Lipschitz continuous.

- (Primal-dual hybrid gradient) Primal-dual hybrid gradient (also called Chambolle and Pock method) targets primal-dual problems with bilinear interaction term, namely, $\mathcal{L}(x, y) = f(x) + y^\top Ax - g(y)$, and it is widely used in image processing [14, 15]. The update can be written as⁵:

$$\begin{aligned} x^{t+1} &\in \operatorname{argmin}_{x \in X} f(x) + (y^t)^\top Ax + \frac{1}{2\eta} \|x - x^t\|_2^2 \\ y^{t+1} &\in \operatorname{argmin}_{y \in Y} g(y) - y^\top A(2x^{t+1} - x^t) + \frac{1}{2\eta} \|y - y^t\|_2^2 \\ \hat{z}^{t+1} &= z^{t+1} = (x^{t+1}, y^{t+1}). \end{aligned} \tag{9}$$

- (Alternating direction method of multipliers) Alternating direction method of multipliers (ADMM) [11, 21, 22, 37] targets linearly constrained convex optimization

⁵ PDHG is often presented in a form with different primal and dual step sizes [14, 15]. Here, we choose to use the same primal and dual step size for consistent notation with other primal-dual algorithms. Our results can easily extend to the case of different step sizes by rescaling. In particular, by setting $\eta = \sqrt{\sigma\tau}$ and defining a rescaled space: $(\hat{x}, \hat{y}) = (x\sqrt{\eta/\tau}, y\sqrt{\eta/\sigma})$ where $\tau \in (0, \infty)$ is the desired primal step and $\sigma \in (0, \infty)$ is the desired dual step size. Applying (9) to this rescaled space, i.e., replacing $f(x)$ with $\hat{f}(\hat{x}) = f(\hat{x}/\sqrt{\eta/\tau})$, $g(y)$ with $\hat{g}(\hat{y}) = g(\hat{y}/\sqrt{\eta/\sigma})$, X with $\hat{X} = \{x\sqrt{\eta/\tau} : x \in X\}$, Y with $\hat{Y} = \{y\sqrt{\eta/\sigma} : y \in Y\}$, $\|x - x^t\|$ with $\|\hat{x}^t - x^t\|$, and $\|y - y^t\|$ with $\|\hat{y}^t - y^t\|$, then substituting back $(\hat{x}, \hat{y}) = (x\sqrt{\eta/\tau}, y\sqrt{\eta/\sigma})$ and $(\hat{x}^t, \hat{y}^t) = (x^t\sqrt{\eta/\tau}, y^t\sqrt{\eta/\sigma})$ yields the classic PDHG update:

$$\begin{aligned} x^{t+1} &\in \operatorname{argmin}_{x \in X} f(x) + (y^t)^\top Ax + \frac{1}{2\tau} \|x - x^t\|_2^2 \\ y^{t+1} &\in \operatorname{argmin}_{y \in Y} g(y) - y^\top A(2x^{t+1} - x^t) + \frac{1}{2\sigma} \|y - y^t\|_2^2. \end{aligned}$$

problems with separable structure

$$\begin{aligned} \min_{x_U \in X_U, x_V \in X_V} \quad & \theta_1(x_U) + \theta_2(x_V) \\ \text{s.t.} \quad & Ux_U + Vx_V = q, \end{aligned}$$

for which the Lagrangian becomes

$$\min_{x \in X} \max_{y \in Y} \mathcal{L}(x, y) := \theta(x) - y^\top (U \ V) \begin{pmatrix} x_U \\ x_V \end{pmatrix} + q^\top y, \tag{10}$$

where $x = (x_U, x_V)$, $\theta(x) = \theta_U(x_U) + \theta_V(x_V)$, $X = X_U \times X_V$, and $Y = \mathbb{R}^m$. The update of ADMM is

$$\begin{aligned} x_U^{t+1} &\in \operatorname{argmin}_{x_U \in X_U} \left\{ \theta_1(x_U) + \frac{\eta}{2} \left\| (Ux_U + Vx_V^t - q) - \frac{1}{\eta} y^t \right\|_2^2 \right\} \\ x_V^{t+1} &\in \operatorname{argmin}_{x_V \in X_V} \left\{ \theta_2(x_V) + \frac{\eta}{2} \left\| (Ux_U^{t+1} + Vx_V - q) - \frac{1}{\eta} y^t \right\|_2^2 \right\} \\ y^{t+1} &= y^t - \eta (Ux_U^{t+1} + Vx_V^{t+1} - q) \\ \hat{z}^{t+1} &= (x_U^{t+1}, x_V^{t+1}, y^t - \eta (Ux_U^{t+1} + Vx_V^t - q)) \\ z^{t+1} &= (x_U^{t+1}, x_V^{t+1}, y^{t+1}) \end{aligned} \tag{11}$$

Although the above four classic primal-dual algorithms have different update rules and may even target different primal-dual problems, it turns out they all satisfy the following assumption, which provides a unified viewpoint of primal-dual algorithms as well as a unified analysis on their ergodic rate.

Property 1 Consider a primal-dual algorithm with iterate update (6). There exists $C > 0$ and a semi-norm $\| \cdot \|$ such that it holds for any $t \in \mathbb{N}$ and $z \in Z$ that

$$\mathcal{L}(\hat{x}^{t+1}, y) - \mathcal{L}(x, \hat{y}^{t+1}) \leq \frac{C}{2} \|z - z^t\|^2 - \frac{C}{2} \|z - z^{t+1}\|^2. \tag{12}$$

Property 1 presents an upper bound on the primal-dual gap at the target solution \hat{z}^{t+1} . Notice the RHS of (12) involves the iterate solution z^t instead of the target solution \hat{z}^t , which is intentional. The next proposition shows how classic primal-dual algorithms satisfy Property 1 as well as their corresponding constants C and semi-norms $\| \cdot \|$. These are known results for individual algorithms; here we provide a unified viewpoint.

Proposition 1 PPM, EGM, PDHG and ADMM satisfy Property 1 with

- (PPM) $C = \frac{1}{\eta}$ and Euclidean norm for any $\eta > 0$;
- (EGM) $C = \frac{1}{\eta}$ and Euclidean norm for $0 < \eta \leq \frac{1}{L}$;

- (PDHG) $C = \frac{1}{\eta}$ and $\|z\|^2 = z^\top Mz$ for $0 < \eta \leq \frac{1}{\sigma_{\max}(A)}$, where $M = \begin{pmatrix} I & -\eta A^\top \\ -\eta A & I \end{pmatrix}$;
- (ADMM) $C = 1$ and $\|z\|^2 = z^\top Mz$ for $\eta > 0$, where $M = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \eta V^\top V & 0 \\ 0 & 0 & \frac{1}{\eta} I \end{pmatrix}$.

Proof (PPM) It follows from the optimality condition of (7) that for all $z \in Z$,

$$\left\langle F(z^{t+1}) + \frac{1}{\eta}(z^{t+1} - z^t), z^{t+1} - z \right\rangle \leq 0.$$

After rearranging the above inequality, we obtain

$$\begin{aligned} \left\langle F(z^{t+1}), z^{t+1} - z \right\rangle &\leq \frac{1}{2\eta} \|z - z^t\|^2 \\ -\frac{1}{2\eta} \|z - z^{t+1}\|^2 - \frac{1}{2\eta} \|z^t - z^{t+1}\|^2 &\leq \frac{1}{2\eta} \|z - z^t\|^2 - \frac{1}{2\eta} \|z - z^{t+1}\|^2. \end{aligned}$$

We finish the proof for PPM by noticing $\mathcal{L}(x^{t+1}, y) - \mathcal{L}(x, y^{t+1}) \leq \langle F(z^{t+1}), z^{t+1} - z \rangle$.

(EGM) See, for example, the proof of Theorem 3.2 in [57] (a simplified version of the proof can be found in Theorem 18.2 in [63]).

(PDHG) See, for example, Lemma 1 and the proof of Theorem 1 in [15].

(ADMM) See, for example, Lemma 3.1 in [37]. □

The next proposition presents how to obtain the ergodic sublinear convergence rate and non-expansiveness of primal-dual algorithms as a direct consequence of Property 1, which will be later used in developing the convergence results for our restart scheme. Again, these are known results for each individual algorithm.

Proposition 2 Consider the sequences $\{\hat{z}^t, z^t\}$ generated from a generic primal-dual algorithm with iterate update (6) and initial solution z^0 . Denote $\bar{z}^t = \frac{1}{t} \sum_{i=1}^t \hat{z}^i$ as the average target solution. Suppose the primal-dual algorithm satisfies Property 1, then it holds for any $z^* \in Z^*$ and $z \in Z$ that

- (Non-expansiveness) $\|z^{t+1} - z^*\| \leq \|z^t - z^*\|$.
- (Sublinear rate on the primal-dual gap) $\mathcal{L}(\bar{x}^t, y) - \mathcal{L}(x, \bar{y}^t) \leq \frac{C\|z - z^0\|^2}{2t}$.

Proof First we prove i. By setting $z = z^*$ in (12) and rearranging the inequality, it holds for any t that

$$\begin{aligned} \|z^* - z^{t+1}\|^2 &\leq \|z^* - z^t\|^2 + \frac{2}{C} \mathcal{L}(x^*, \hat{y}^{t+1}) - \frac{2}{C} \mathcal{L}(\hat{x}^{t+1}, y^*) \\ &= \|z^* - z^t\|^2 + \frac{2}{C} \mathcal{L}(x^*, \hat{y}^{t+1}) - \frac{2}{C} \mathcal{L}(x^*, y^*) \\ &\quad + \frac{2}{C} \mathcal{L}(x^*, y^*) - \frac{2}{C} \mathcal{L}(\hat{x}^{t+1}, y^*) \\ &\leq \|z^* - z^t\|^2, \end{aligned}$$

where the second inequality is because $\mathcal{L}(x^*, y^*) = \min_x \mathcal{L}(x, y^*) \leq \mathcal{L}(x^{t+1}, y^*)$ and $\mathcal{L}(x^*, y^*) = \max_y \mathcal{L}(x^*, y) \geq \mathcal{L}(x^*, y^{t+1})$.

We now prove ii:

$$\begin{aligned} \mathcal{L}(\bar{x}^t, y) - \mathcal{L}(x, \bar{y}^t) &\leq \frac{1}{t} \left(\sum_{i=1}^t \mathcal{L}(\hat{x}^i, y) - \mathcal{L}(x, \hat{y}^i) \right) \leq \frac{C}{2t} \left(\|z - z^0\|^2 - \|z - \bar{z}^t\|^2 \right) \\ &\leq \frac{C\|z - z^0\|^2}{2t}, \end{aligned}$$

where the first inequality is from convexity-concavity of $\mathcal{L}(x, y)$ and Jensen’s inequality, and the second inequality holds by summing up (12) over t and telescoping. \square

The non-expansiveness in Proposition 2 implies that z^t stays in a bounded region. The next assumption assumes the target solution \hat{z}^t is not far away from either z^t or z^{t-1} , thus \hat{z}^t also stays in a bounded region.

Property 2 *There exists $q > 0$ such that for all $t \in \mathbb{N}$ either $\|\hat{z}^t - z^t\| \leq q \mathbf{dist}(z^t, Z^*)$, or $\|\hat{z}^t - z^{t-1}\| \leq q \mathbf{dist}(z^t, Z^*)$.*

The four classic primal-dual algorithms we discussed above satisfy Property 2:

Proposition 3 *Under the same conditions as Proposition 1, PPM, EGM, PDHG and ADMM satisfy Assumption 2 since*

- (PPM) $\hat{z}^t = z^t$ for all $t \in \mathbb{N}$, and therefore $q = 0$;
- (EGM) $\|\hat{z}^t - z^{t-1}\| \leq 3 \mathbf{dist}(z^t, Z^*)$ for all $t \in \mathbb{N}$, and therefore $q = 3$;
- (PDHG) $\hat{z}^t = z^t$ for all $t \in \mathbb{N}$, and therefore $q = 0$;
- (ADMM) $\|\hat{z}^t - z^t\| \leq 2 \mathbf{dist}(z^t, Z^*)$ for all $t \in \mathbb{N}$, and therefore $q = 2$.

Proof Note that the results for PPM and PDHG hold immediately.

Proof for EGM. The proof for EGM is somewhat technical so we defer it to “Appendix B”.

Proof for ADMM. Let $z^* = \min_{z \in Z^*} \|z^t - z\|$. By (11) we have $\hat{x}_U^{t+1} = x_U^{t+1}$, $\hat{x}_V^{t+1} = x_V^{t+1}$, and $\hat{y}^{t+1} = y^{t+1} + \eta V(x_V^{t+1} - x_V^t)$. Therefore,

$$\begin{aligned} \|\hat{z}^{t+1} - z^{t+1}\| &= \sqrt{(x_V^{t+1} - x_V^t)^\top \eta V^\top V (x_V^{t+1} - x_V^t)} \\ &= \|x_V^{t+1} - x_V^t\|_{\eta V^\top V} \\ &\leq \|z^{t+1} - z^t\| \\ &\leq \|z^{t+1} - z^*\| + \|z^t - z^*\| \\ &\leq 2\|z^t - z^*\|, \end{aligned}$$

where the last inequality uses Proposition i. \square

Property 3 captures the minimal algorithmic properties required for restart schemes to achieve linear convergence, as shown in Sect. 4. This technical assumption essentially states that the primal-dual algorithm imposes an order $1/t$ sublinear convergence

on the decay of the normalized duality gap (Property i.) and the iterates do not move too far away from the starting point (Property ii.). Proposition 4 demonstrates that if Property 1 and 2 hold, then Property 3 also holds (with different parameters q and C), so Property 3 holds for the classic primal-dual algorithms (PPM, EGM, PDHG and ADMM).

Property 3 *There exist scalar constants $q, C > 0$ such that for all initial points $z^0 \in Z$, the output sequence $\{\bar{z}^t\}_{t=1}^\infty$ generated by the algorithm satisfies $\bar{z}^t \in Z$ and*

- i. $\rho_{\|\bar{z}^t - z^0\|}(\bar{z}^t) \leq \frac{2C\|\bar{z}^t - z^0\|}{t}$,
- ii. $\|\bar{z}^t - z^0\| \leq (q + 2)\mathbf{dist}(z^0, Z^*)$.

Proposition 4 *Consider the sequences $\{\hat{z}^t, z^t\}$ generated from a generic primal-dual algorithm with iterate update (6) and initial solution z^0 . Suppose the primal-dual algorithm satisfies Property 1 and 2, then Property 3 holds with $\bar{z}_t = \frac{1}{t} \sum_{i=1}^t \hat{z}_t$.*

Proof Suppose $r = \|\bar{z}^t - z^0\| > 0$, then

$$\begin{aligned} \rho_r(\bar{z}^t) &= \frac{1}{r} \max_{\hat{z} \in W_r(\bar{z}^t)} \{ \mathcal{L}(\bar{x}^t, \hat{y}) - \mathcal{L}(\hat{x}, \bar{y}^t) \} \leq \frac{1}{r} \max_{\hat{z} \in W_{2r}(z^0)} \{ \mathcal{L}(\bar{x}^t, \hat{y}) - \mathcal{L}(\hat{x}, \bar{y}^t) \} \\ &\leq \frac{C(2r)^2}{2tr} = \frac{2Cr}{t} \end{aligned}$$

where the first inequality is from $W_r(\bar{z}^t) \subseteq W_{2r}(z^0)$ by the definition of r , and the second inequality uses Proposition ii.. Alternatively, if $\|\bar{z}^t - z^0\| = 0$ then it suffices to observe that $\rho_r(\bar{z}^t) \leq 2Cr/t$ and $\rho_r(z) \geq 0$ implies $\limsup_{r \rightarrow 0^+} \rho_r(\bar{z}^t) = 0$. This establishes Property i..

We turn to proving Property ii.. Let $z^* \in \operatorname{argmin}_{z \in Z^*} \|z - z^0\|$, then for either $i = t$ or $i = t - 1$,

$$\begin{aligned} \|\hat{z}^t - z^0\| &\leq \|\hat{z}^t - z^i\| + \|z^i - z^*\| + \|z^* - z^0\| \leq \|\hat{z}^t - z^i\| + 2\|z^*\| \\ -z^0\| &\leq (q + 2)\|z^0 - z^*\| \end{aligned}$$

where the first inequality uses the triangle inequality, the second inequality uses Proposition i., and the third inequality uses $\|\hat{z}^t - z^i\| \leq q\mathbf{dist}(z^i, Z^*)$ as per Property 2. From $\|\hat{z}^t - z^0\| \leq (q + 2)\|z^0 - z^*\|$ it follows that

$$\begin{aligned} \|\bar{z}^t - z^0\| &= \left\| \frac{1}{t} \sum_{i=1}^t \hat{z}^i - z^0 \right\| \leq \frac{1}{t} \sum_{i=1}^t \|\hat{z}^i - z^0\| \\ &\leq \frac{1}{t} \sum_{i=1}^t (q + 2)\|z^0 - z^*\| = (q + 2)\mathbf{dist}(z^0, Z^*) \end{aligned}$$

which establishes Property ii.. □

Table 1 The corresponding constants and norms in Property 3 and for different algorithms. Recall L is the Lipschitz constant of F . The values in this table can be found by inspecting Propositions 1, 3 and 4

	C	Valid η values	q	Semi-norm squared
PPM	$\frac{1}{\eta}$	$(0, \infty)$	0	$z^\top z$
EGM	$\frac{1}{\eta}$	$(0, 1/L]$	3	$z^\top z$
PDHG	$\frac{1}{\eta}$	$(0, 1/\sigma_{\max}(A)]$	0	$z^\top \begin{pmatrix} \mathbf{I} & -\eta A^\top \\ -\eta A & \mathbf{I} \end{pmatrix} z$
ADMM	1	$(0, \infty)$	2	$z^\top \begin{pmatrix} 0 & 0 & 0 \\ 0 & \eta V^\top V & 0 \\ 0 & 0 & \frac{1}{\eta} \mathbf{I} \end{pmatrix} z$

3 Sharpness for primal-dual problems

This section discusses the basic properties of the normalized duality gap and sharp primal-dual problems, and then presents the concept of sharp primal-dual problems along with examples.

Critical to the paper is the concept of a sharp primal-dual problem. Indeed, when this condition holds, we are able to establish linear convergence using restarts (Sect. 4).

Definition 1 We say a primal-dual problem (1) is α -sharp on the set $S \subseteq Z$ if ρ_r is α -sharp (see (5)) on S for all $r \in (0, \text{diam}(S)]$, i.e., it holds for all $z \in S$ that

$$\alpha \text{dist}(z, Z^*) \leq \rho_r(z) . \tag{13}$$

As we can see, the sharpness of a primal-dual problem essentially says the localized duality gap $\rho_r(z)$ is sharp in the traditional sense. A key concept is that the sharpness is defined on a set $S \subseteq Z$. Even if the domain Z is unbounded, the sharpness can be defined on a bounded set S .

From the definition of the normalized duality gap given in (4) and the sharpness condition, we can immediately establish the monotonicity of sharpness in r and S (Facts 1 and 2 and Proposition 5).

Fact 1 It holds for any z that $r \rho_r(z)$ is monotonically non-decreasing for $r \in [0, \infty)$.

Fact 2 The sharpness condition is monotone in set S , namely, suppose $S_1 \subseteq S_2 \subseteq Z$ and $\rho_r(z)$ is α -sharp in S_2 , then $\rho_r(z)$ is α -sharp in S_1 .

Proposition 5 It holds for any fixed z that $\rho_r(z)$ is monotonically non-increasing for $r \in [0, \infty)$.

Proof We prove the result for $r \in (0, \infty)$ which immediately implies the result also holds for $r = 0$ by definition of $\rho_0(z)$. Consider any $z \in Z$ and $0 < r_1 < r_2$. Let $z_2 \in \operatorname{argmax}_{\hat{z} \in W_{r_2}(z)} \{ \mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y) \}$. Then we have that $z_1 := z + \frac{r_1}{r_2}(z_2 - z) \in W_{r_1}(z)$.

Moreover, it follows from convexity of $\mathcal{L}(x, y)$ over x that $\mathcal{L}(x_1, y) \leq \frac{r_1}{r_2} \mathcal{L}(x_2, y) +$

$(1 - \frac{r_1}{r_2})\mathcal{L}(x, y)$, which after rearrangement becomes

$$\mathcal{L}(x, y) - \mathcal{L}(x_2, y) \leq \frac{r_2}{r_1} (\mathcal{L}(x, y) - \mathcal{L}(x_1, y)).$$

Similarly, it follows from concavity of $\mathcal{L}(x, y)$ over y that

$$\mathcal{L}(x, y_2) - \mathcal{L}(x, y) \leq \frac{r_2}{r_1} (\mathcal{L}(x, y_1) - \mathcal{L}(x, y)).$$

Therefore, it holds that

$$\begin{aligned} \rho_{r_2}(z) &= \frac{\mathcal{L}(x, y_2) - \mathcal{L}(x_2, y)}{r_2} \\ &= \frac{(\mathcal{L}(x, y_2) - \mathcal{L}(x, y)) + (\mathcal{L}(x, y) - \mathcal{L}(x_2, y))}{r_2} \\ &\leq \frac{(\mathcal{L}(x, y_1) - \mathcal{L}(x, y)) + (\mathcal{L}(x, y) - \mathcal{L}(x_1, y))}{r_1} \\ &\leq \rho_{r_1}(z), \end{aligned}$$

which finishes the proof for the proposition. □

Proposition 6 shows that when the normalized duality gap is a reasonable metric to measure the quality of a solution: the normalized duality gap is zero if and only if the primal-dual gap is also zero.

Proposition 6 *For any $r \in [0, \infty)$, the primal-dual gap (3) at z is zero if and only if $\rho_r(z) = 0$.*

Proof Suppose that the primal-dual gap at z is zero, then by definition of ρ_r we get for all $r \in (0, \infty)$ that $0 \leq \rho_r(z) \leq \frac{1}{r} \max_{\hat{z} \in Z} \{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)\} = 0$. Taking the limit as $r \rightarrow 0^+$ shows $\rho_0(z) = 0$.

For the other direction argue by contrapositive. In particular, we assume the primal-dual gap at z is nonzero, and using this assumption show that for all $r \in [0, \infty)$ that $\rho_r(z) > 0$. From the assumption that the primal-dual gap is nonzero, there exists $\hat{z} \in Z$ such that $\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y) > 0$. For any $r \in (0, \infty)$, setting $\alpha = \min \left\{ 1, \frac{r}{\|\hat{z} - z\|} \right\}$ yields $\|\alpha(\hat{z} - z)\| = \alpha\|\hat{z} - z\| \leq r$ and $z + \alpha(\hat{z} - z) = (1 - \alpha)z + \alpha\hat{z} \in Z$. Therefore $z + \alpha(\hat{z} - z) \in W_r(z)$. By definition of \hat{z} , $\alpha > 0$, convex-concavity of \mathcal{L} , and $z + \alpha(\hat{z} - z) \in W_r(z)$ we get

$$\begin{aligned} 0 &< \frac{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)}{\|\hat{z} - z\|} \leq \frac{\mathcal{L}(x, y + \alpha(\hat{y} - y)) - \mathcal{L}(y + \alpha(\hat{y} - y), y)}{\alpha\|\hat{z} - z\|} \\ &\leq \frac{r}{\alpha\|\hat{z} - z\|} \rho_r(z) \end{aligned}$$

thus $\rho_r(z) > 0$ for $r \in (0, \infty)$. Moreover, since $\rho_r(z) > 0$ for $r \in (0, \infty)$ we deduce $\rho_0(z) > 0$ by Proposition 5. This finishes the proof. □

We next present examples of sharp primal-dual problems. Here we utilize Euclidean norm in the sharpness definition to illustrate the examples. For any other norms, the sharpness condition still holds (upto a constant) by noticing that any norms in finite space are equivalent to each other (See, for example, Theorem 5.36 in [39]). As an illustrating example, we show below that the corresponding norm for PDHG in Proposition 1 is approximately Euclidean (up to a constant):

Proposition 7 *Let $\eta \in (0, 1/\sigma_{\max}(A))$. If $\|z\|^2 = \|x\|_2^2 - 2\eta y^\top Ax + \|y\|_2^2$, which is the PDHG norm, then*

$$(1 - \eta\sigma_{\max}(A))\|z\|_2^2 \leq \|z\|^2 \leq (1 + \eta\sigma_{\max}(A))\|z\|_2^2.$$

Proof Lower bounding $\|z\|^2$ gives, $\|z\|^2 \geq \|x\|_2^2 - 2\eta\sigma_{\max}(A)\|y\|_2\|x\|_2 + \|y\|_2^2 = (1 - \eta\sigma_{\max}(A))\|z\|_2^2 + \eta\sigma_{\max}(A)(\|x\|_2 - \|y\|_2)^2 \geq (1 - \eta\sigma_{\max}(A))\|z\|_2^2$. Similarly, we have the other side. \square

3.1 Sharpness of primal-dual problems with bounded feasible regions

Lemma 1 *Consider a primal-dual problem (1) with a bounded feasible set $Z = X \times Y$. Suppose the constraint set Z has diameter R , and Z is equipped with the Euclidean norm. Let $P(x) = \max_{y \in Y} \mathcal{L}(x, y)$ and $D(y) = \min_{x \in X} \mathcal{L}(x, y)$ denote the primal objective and the dual objective, respectively. If $P(x)$ is α_1 -sharp in x and $D(y)$ is α_2 -sharp in y , then the primal-dual problem (1) is $\min\{\alpha_1, \alpha_2\}/R$ -sharp on $S = Z$.*

Proof Let X^* and Y^* denote the minimizers of $P(x)$ and $D(x)$ respectively. Then it holds for any $z \in Z$ and $r \leq R$ that

$$\begin{aligned} \rho_r(z) &\geq \rho_R(z) = \frac{\max_{\hat{z} \in Z} \{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)\}}{R} = \frac{P(x) - D(y)}{R} \\ &\geq \frac{\alpha_1 \mathbf{dist}(x, X^*) + \alpha_2 \mathbf{dist}(y, Y^*)}{R} \geq \frac{\min\{\alpha_1, \alpha_2\}}{R} \mathbf{dist}(z, Z^*), \end{aligned}$$

where the first inequality utilizes Proposition 5, the first equality is from the definition of $P(x)$ and $D(y)$, the second inequality uses sharpness, and the third inequality uses the triangle inequality. \square

3.2 Sharpness of bilinear problems

We will find the following well-known Lemma useful. The proof of Lemma 2 appears in the ‘‘Appendix A’’ for completeness.

Lemma 2 *Suppose $\|\cdot\|$ is the Euclidean norm. Let n and m be positive integers, $H \in \mathbb{R}^{m \times n}$ and $h \in \mathbb{R}^m$. Define $S := \{z \in \mathbb{R}^n : Hz = h\}$. If $S \neq \emptyset$, then $\mathbf{dist}(S, z) \leq \frac{1}{\sigma_{\min}^+(H)} \|Hz - h\|$.*

Lemma 3 Suppose $\mathcal{L}(x, y) = c^\top x - y^\top Ax + b^\top y$, $Z = \mathbb{R}^{m+n}$ and Z is equipped with the Euclidean norm. Suppose there exists a finite solution to (1). Then the primal-dual problem (1) is α -sharp on Z with $\alpha = \sigma_{\min}^+(A)$.

Proof Since $\mathcal{L}(x, y) = c^\top x - y^\top Ax + b^\top y$ it holds for any z and \hat{z} that

$$\begin{aligned} \mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y) &= \mathcal{L}(x, \hat{y}) - \mathcal{L}(x, y) + \mathcal{L}(x, y) - \mathcal{L}(\hat{x}, y) \\ &= (-\nabla_y \mathcal{L}(x, y))^\top (y - \hat{y}) + \nabla_x \mathcal{L}(x, y)^\top (x - \hat{x}) \\ &= F(z)^\top (z - \hat{z}) . \end{aligned} \tag{14}$$

Therefore, it holds that

$$\rho_r(z) = \frac{1}{r} \max_{z' \in W_r(z)} F(z)^\top (\hat{z} - z) = \|F(z)\| . \tag{15}$$

Meanwhile, notice $Z^* = \{z \in Z \mid F(z) = 0\} = \{z \in Z \mid Hz = h\}$, where $H = \begin{pmatrix} A^\top \\ A \end{pmatrix}$ and $h = \begin{pmatrix} c \\ b \end{pmatrix}$. It then follows from Lemma 2 that

$$\|F(z)\|^2 = \|Hz - h\|^2 = \|A^\top y - c\|^2 + \|Ax - b\|^2 \geq \sigma_{\min}^+(A)^2 \mathbf{dist}(z, Z^*)^2 . \tag{16}$$

Combining (15) and (16), we arrive at

$$\rho_r(z) = \|F(z)\| \geq \sigma_{\min}^+(A) \mathbf{dist}(z, Z^*) ,$$

which finishes the proof. □

Remark 1 In the above bilinear example, we see that the primal-dual problem (1) is α -sharp if and only if $\|F(z)\|$ is α -sharp in the standard sense. Indeed, $\|F(z)\|$ being a sharp function is a necessary condition for $\mathcal{L}(x, y)$ to be sharp for any unconstrained convex-concave primal-dual problem. This can be seen by noticing

$$\rho_0(z) = \|F(z)\| ,$$

and by utilizing the monotonicity of $\rho_r(z)$ stated in Proposition 5.

3.3 Sharpness of standard linear programming

Consider a generic LP problem:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & c^\top x \\ \text{s.t.} \quad & Ax = b \\ & x \geq 0 , \end{aligned} \tag{17}$$

its dual:

$$\begin{aligned} & \max_{y \in \mathbb{R}^m} b^\top y \\ & \text{s.t. } A^\top y \geq c, \end{aligned} \tag{18}$$

and its Lagrangian form:

$$\min_{x \geq 0} \max_{y \in \mathbb{R}^m} \mathcal{L}(x, y) = c^\top x + y^\top b - y^\top Ax . \tag{19}$$

Suppose (17) and (18) have feasible solutions. By strong duality we can reduce (17) and (18) to solving the system $Kz \geq h$ where

$$K := \begin{pmatrix} \mathbf{I} & 0 \\ -A & 0 \\ A & 0 \\ 0 & -A^\top \\ -c^\top & b^\top \end{pmatrix}, \quad h := \begin{pmatrix} 0 \\ -b \\ b \\ -c \\ 0 \end{pmatrix}. \tag{20}$$

The term $\|(h - Kz)^+\|$ is a common metric for the termination of algorithms for LP [3]. Lemma 4 shows that this metric can be bounded via the normalized duality gap.

Lemma 4 *Let $\|\cdot\|$ be the Euclidean norm. Consider the Lagrangian (19), formed from the LP setup. Suppose that there is a finite solution to (19). Then, for all $R \in (0, \infty)$, $r \in (0, R]$, and $z \in W_R(0)$ we have*

$$\|(h - Kz)^+\| \leq \rho_r(z) \sqrt{1 + R^2}.$$

Proof Consider $z \in Z$. Let $v = \begin{pmatrix} (-c + A^\top y)^+ \\ b - Ax \end{pmatrix}$ and note that $F(z) = \begin{pmatrix} c - A^\top y \\ Ax - b \end{pmatrix}$. Recall from (14) that $\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y) = F(z)^\top (z - \hat{z})$. Therefore, it follows from the definition of ρ_r that for all $\hat{z} \in W_r(z)$,

$$\rho_r(z) \geq \frac{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)}{r} = \frac{F(z)^\top (z - \hat{z})}{r}. \tag{21}$$

Let $z_1 = z + r \frac{v}{\|v\|}$, then $\|z_1 - z\| = r$. Meanwhile we have $(-c + A^\top y)^+ \geq 0$, thus $z_1 \in Z$ by noticing $z \in Z$. Therefore, $z_1 \in W_r(z)$ and it follows from (21) that

$$\rho_r(z) \geq -\frac{1}{\|v\|} F(z)^\top v \geq \|v\|, \tag{22}$$

where the last inequality utilizes the definition of v .

If $z = 0$, let $z_2 = 0 \in W_r(z)$, then we have from (21) that

$$\rho_r(z) \geq 0 = c^\top x - b^\top y = \frac{1}{r} (c^\top x - b^\top y). \tag{23}$$

Otherwise, let $z_2 = z - \min \left\{ \frac{r}{\|z\|}, 1 \right\} z$, then $\|z_2 - z\| \leq \frac{r}{\|z\|} \|z\| \leq r$. Meanwhile, we have $x_2 \geq x - x = 0$, thus $z_2 \in Z$. Therefore, $z_2 \in W_r(z)$ and it follows from (21) that

$$\rho_r(z) \geq \frac{1}{r} \min \left\{ \frac{r}{\|z\|}, 1 \right\} F(z)^\top z = \min \left\{ \frac{1}{\|z\|}, \frac{1}{r} \right\} (c^\top x - b^\top y). \tag{24}$$

Taking the worst case bound in (23) and (24) by noticing that $0 < r \leq R, 0 \leq \|z\| \leq R$ and $\rho_r(z) \geq 0$ yields

$$\rho_r(z) \geq \frac{1}{R} (c^\top x - b^\top y)^+. \tag{25}$$

Combining (22) and (25), we obtain

$$\begin{aligned} (1 + R^2)\rho_r(z)^2 &\geq ((c^\top x - b^\top y)^+)^2 + \|(-c + A^\top y)^+\|^2 + \|b - Ax\|^2 \\ &= \|(h - Kz)^+\|^2, \end{aligned}$$

where the equality utilizes the fact $x \geq 0$ and the property of ℓ_2 norm. Taking the square root finishes the proof. \square

Let $H(K)$ be the Hoffman constant [38] of the matrix K in Euclidean norm, i.e., it holds for any $z \in \mathbb{R}^{n+m}$ that

$$\mathbf{dist}(z, Z^*) \leq H(K) \|(h - Kz)^+\|. \tag{26}$$

Remark 2 A popular characterization of $H(A)$ for linear inequalities is (see for example [33, 42, 66])

$$H(A) = \max_{\substack{J \subseteq \{1, \dots, 2m+2n+1\} \\ A_J \text{ has full row rank}}} \frac{1}{\min_{v \in \mathbb{R}_+^J, \|v\|=1} \|A_J^\top v\|}, \tag{27}$$

where A_J is the matrix with the corresponding rows of A indexed by J . In terms of the singular values of A , a slightly looser bound would be

$$H(A) \geq \max_{\substack{J \subseteq \{1, \dots, 2m+2n+1\} \\ A_J \text{ has full row rank}}} \frac{1}{\sigma_{\min}^+(A_J)}$$

which ignores the constraints $v \in \mathbb{R}_+^J$ in (27) and take advantage of the Euclidean norm.

Lemma 5 (Sharpness of linear programming) *Let $R \in (0, \infty)$ and assume (19) has a solution. Then, the primal-dual problem (19) is α -sharp on the set $W_R(0)$ where $\alpha = \frac{1}{H(K)\sqrt{1+4R^2}}$.*

Proof For all $r \in (0, \text{diam}(W_R(0))] \subseteq (0, 2R]$, by Lemma 4 and (26),

$$\rho_r(z)\sqrt{1 + 4R^2} \geq \|(h - Kz)^+\| \geq \frac{1}{H(K)} \text{dist}(z, Z^*).$$

□

3.4 Sharpness of ADMM for linear programming

Following Lemma 5, we can show different formulations of LP may also be sharp. In particular, one popular method for LP is using ADMM [65]. Let $X_U = \{x \in \mathbb{R}^n \mid Ax = b\}$ and $X_V = \{x \in \mathbb{R}^n \mid x \geq 0\}$. Consider the following form of LP:

$$\begin{aligned} \min_{x_U \in X_U, x_V \in X_V} \quad & c^\top x_V \\ \text{s.t.} \quad & x_U - x_V = 0, \end{aligned}$$

and its Lagrangian form

$$\min_{x_U \in X_U, x_V \in X_V} \max_{y \in \mathbb{R}^m} \mathcal{L}(x, y) = c^\top x_V - y^\top (\mathbf{I} - \mathbf{I}) \begin{pmatrix} x_U \\ x_V \end{pmatrix}. \tag{28}$$

Lemma 6 (Sharpness of ADMM for linear programming) *Suppose (28) has a solution. Then there exists a matrix K' and vector h' such that the set of solutions to (28) are equal to $\{z : K'z \geq h'\}$. Let $z^* \in Z^*$, then it holds for any set $S(R, z^*) = \{z \in Z : \|z^* - z\| \leq R\}$ that the primal-dual problem (28) is α -sharp where $\alpha = \frac{1}{\max\{\eta^2, 1/\eta^2\}H(K')\sqrt{1+4R^2}}$.*

The proof of Lemma 6 appears in Section C and follows the same structure as the proof of Lemma 5.

4 Restart schemes for primal-dual methods

In this section, we present a generic restart scheme for primal-dual methods and show its linear convergence rate for solving sharp primal-dual problems. We discuss two restart schemes: fixed frequency restart and adaptive restarts, where the former requires knowledge of the sharpness constant α while the latter does not.

Algorithm 1 presents our nested-loop restarted primal-dual algorithm. We initialize with $z^{0,0} \in Z$, a suitable primal-dual algorithm, and a step-size of the algorithm η . At each outer loop, we keep running PrimalDualStep until one of the restart conditions holds (to be discussed later). More specifically, at the t -th inner loop of the n -th outer loop, we call PrimalDualStep to update the solution $z^{n,t}$ and keep track of the target solution $\hat{z}^{n,t}$ as well as the running average $\bar{z}^{n,t}$. At the end of each outer loop, we restart the next outer loop from the output of a primal-dual algorithm, namely, the running average \bar{z}^{n,τ^n} , and store the length of its inner loops as τ^n . Here we propose two restart schemes:

Algorithm 1: Restarted Primal-Dual Algorithms

```

1 Input: An initial solution  $z^{0,0}$ , a primal-dual algorithm PrimalDualStep, a step-size  $\eta$ ;
2 Initialize outer loop counter  $n \leftarrow 0$ ;
3 repeat
4   initialize the inner loop. inner loop counter  $t \leftarrow 0$ ;
5   repeat
6      $z^{n,t+1}, \hat{z}^{n,t+1} \leftarrow \text{PrimalDualStep}(z^{n,t}, \eta)$ ;
7      $\bar{z}^{n,t+1} \leftarrow \frac{1}{t+1} \sum_{i=1}^{t+1} \hat{z}^{n,i} = \frac{t}{t+1} \bar{z}^{n,t} + \frac{1}{t+1} \hat{z}^{n,t+1}$ ;
8      $t \leftarrow t + 1$ ;
9   until one of the two restart conditions (see texts) holds;
10  restart the outer loop.  $\tau^n \leftarrow t, z^{n+1,0} \leftarrow \bar{z}^{n,\tau^n}, n \leftarrow n + 1$ ;
11 until  $z^{n,0}$  convergence;
12 Output:  $z^{n,0}$ .

```

4.1 Fixed frequency restarts

Suppose we know the sharpness constant α of the primal-dual problem and an upper bound on C . In this scheme, we break the inner loop and restart the outer loop with a fixed frequency t^* . Namely, we restart the algorithm if

$$t \geq t^* := \left\lceil \frac{2C(q+2)}{\alpha\beta} \right\rceil, \tag{29}$$

where C and q are the corresponding parameters of a given primal-dual algorithm, as stated in Property 1 and 2. The parameter $\beta \in (0, 1)$ controls when a restart is triggered. The value of β can be tuned to improve both practical performance and the constants in the theoretical bound. Concretely, it follows from Proposition 1, Proposition 3 and Proposition 4 that

- For PPM with $\eta \in (0, \infty)$, $t^* = \left\lceil \frac{4}{\alpha\beta\eta} \right\rceil$;
- For PDHG with $\eta \in (0, 1/\sigma_{\max}(A)]$, $t^* = \left\lceil \frac{4}{\alpha\beta\eta} \right\rceil$;
- For EGM with $\eta \in (0, 1/L]$, $t^* = \left\lceil \frac{10}{\alpha\beta\eta} \right\rceil$;
- For ADMM, $t^* = \left\lceil \frac{8}{\alpha\beta} \right\rceil$.

4.2 Adaptive restarts

In practice, it can be non-trivial to compute the sharpness constant α of a given primal-dual problem. Here we propose an adaptive restart scheme that does not need an estimate of α . In this scheme, we restart when the normalized duality gap has sufficient decay; more specifically, we restart the algorithm if

$$\begin{cases} \rho_{\|\bar{z}^{n,t} - z^{n,0}\|}(\bar{z}^{n,t}) \leq \beta \rho_{\|z^{n,0} - z^{n-1,0}\|}(z^{n,0}) & \text{if } n \geq 1 \\ t \geq \tau^0 & \text{if } n = 0 \end{cases} \tag{30}$$

where the first restart interval length $\tau^0 \in \mathbb{N}$ can be selected by the user as a hyperparameter of the algorithm. In practice, one can pick $\tau^0 = 1$ for simplicity.

Although the two restarting schemes (29) and (30) look disconnected, it turns out the restart interval length in (30) is upper bounded by t^* , as shown later in Theorem 2. Furthermore, although the parameter $\beta \in (0, 1)$ seems to play different roles in fixed frequency restart scheme (29) and adaptive restart scheme (30), it determines the contraction of the distance to optimal solution set for both schemes.

The rest of this section presents the linear convergence rate of the two restart schemes stated above. We start from the fixed frequency restart scheme.

Proposition 8 is useful throughout this section to ignore awkward cases in our proofs when $\|\bar{z}^{n,t} - z^{n,0}\| = 0$. For example, technically sharpness (Definition 1) is only defined for $r > 0$.

Proposition 8 *If Property 3 holds and $\|\bar{z}^{n,t} - z^{n,0}\| = 0$ then $\bar{z}^{n,t} = z^{n,0} \in Z^*$.*

Proof Property 3 implies $\rho_0(\bar{z}^{n,t}) = 0$ and therefore by Proposition 6 we get $\bar{z}^{n,t} \in Z^*$. □

Theorem 1 (Fixed frequency restarts) *Consider the sequence $\{z^{n,0}\}_{i=0}^\infty, \{\tau^n\}_{i=1}^\infty$ generated by Algorithm 1 with the fixed frequency restart scheme, namely, we restart the outer loop if (29) holds. Suppose the (6) satisfies Property 3, and the primal-dual problem (1) is α -sharp on the set $W_R(z^{0,0})$ with $R = \frac{q+2}{1-\beta} \mathbf{dist}(z^{0,0}, Z^*)$. Then it holds for each outer iteration $n \in \mathbb{N} \cup \{0\}$ that:*

$$\mathbf{dist}(z^{n,0}, Z^*) \leq \beta^n \mathbf{dist}(z^{0,0}, Z^*) . \tag{31}$$

Proof We argue by induction. Note this hypothesis trivially holds for $n = 0$. Suppose the hypothesis (31) holds for all $n \leq N$. Then it holds that

$$\begin{aligned} \|z^{N+1,0} - z^{0,0}\| &\leq \sum_{n=0}^N \|z^{n+1,0} - z^{n,0}\| \leq (q + 2) \sum_{n=0}^N \mathbf{dist}(z^{n,0}, Z^*) \\ &\leq (q + 2) \sum_{n=0}^N \beta^n \mathbf{dist}(z^{0,0}, Z^*) \leq \frac{q + 2}{1 - \beta} \mathbf{dist}(z^{0,0}, Z^*) , \end{aligned}$$

where the first inequality uses triangle inequality, the second inequality is from Property ii., the third inequality uses the induction hypothesis, and the last inequality is from bounds on sum of a geometric sequence. This implies $z^{N+1,0} \in W_R(z^{0,0})$. It then follows that

$$\mathbf{dist}(z^{N+1,0}, Z^*) \leq \frac{\rho_{\|z^{N+1,0} - z^{N,0}\|}(z^{N+1,0})}{\alpha} \leq \frac{2C\|z^{N+1,0} - z^{N,0}\|}{\alpha t^*}$$

$$\begin{aligned} &\leq \frac{2C(q+2)}{\alpha t^*} \mathbf{dist}(z^{N,0}, Z^*) \leq \beta \mathbf{dist}(z^{N,0}, Z^*) \\ &\leq \beta^{N+1} \mathbf{dist}(z^{0,0}, Z^*), \end{aligned} \tag{32}$$

where the first inequality uses α -sharpness of the primal-dual problem by choosing $r = \|z^{N+1,0} - z^{N,0}\|$ and $z = z^{N+1,0} \in W_R(z^{0,0})$ in (13), the second inequality utilizes Property i., the third inequality is from Property ii. by noticing $z^{N+1,0} = \bar{z}^{t^*}$, and the fourth inequality comes from the definition of t^* . This finishes the proof by induction. \square

Remark 3 As a direct consequence of Theorem 1, we need $O(\frac{C}{\alpha} \log(\frac{1}{\epsilon}))$ total iterations of PDHG to find an approximate solution z such that the distance $\mathbf{dist}(z, Z^*) \leq \epsilon$.

Remark 4 We comment the proof logic of Theorem 1, and compare it with the standard linear convergence proofs for restart schemes in convex optimization. Consider restarted accelerated gradient descent to minimize a convex, L -smooth, and that grows μ -quadratically⁶ function f . The standard analysis of restarted accelerated gradient descent can be written as (see, for example, [60])

$$\mathbf{dist}(z^{n+1,0}, Z^*)^2 \leq \frac{f(z^{n+1,0}) - f^*}{\mu} \leq \frac{2L \mathbf{dist}(z^{n+1,0}, Z^*)^2}{\mu(\hat{t}^*)^2} \leq \beta^2 \mathbf{dist}(z^{n,0}, Z^*)^2$$

where the second inequality uses the $O(1/k^2)$ rate of accelerated gradient descent [61, 62] when the inner loop iteration is larger than \hat{t}^* , and the last inequality is from $\hat{t}^* = \frac{1}{\beta} \sqrt{\frac{2L}{\mu}}$. This showcases the (accelerated) linear convergence rate of restarted accelerated gradient descent. Our analysis of Theorem 1 follows from a similar argument: sharpness allows us to transform the sublinear bound on the normalized duality gap $\rho_{\|z^{n+1,0} - z^{n,0}\|}(z^{n+1,0})$ (of a classic primal-dual algorithm) to a constant factor contraction in the distance to optimality after t^* inner iterations. A major difference is that we are able to evaluate the normalized duality gap in our setting, while in convex optimization, it can be non-trivial to evaluate the optimality gap $f(z^{n+1,0}) - f^*$. Indeed, this feature is crucial as we develop the adaptive restarting scheme.

The next theorem presents the linear convergence rate of the adaptive restart scheme:

Theorem 2 (Adaptive restarts) *Consider the sequence $\{z^{n,0}\}_{i=0}^\infty, \{\tau^n\}_{i=1}^\infty$ generated by Algorithm 1 with the adaptive restart scheme, namely, we restart the outer loop if (30) holds. Suppose (6) satisfies Property 3. Suppose there exists a set $S \subseteq Z$ such that $z^{n,0} \in S$ for any $n \geq 0$, and the primal-dual problem (1) is α -sharp on the set S . Then it holds for each outer iteration $n \in \mathbb{N}$ that*

- i. *The restart length, τ^n , is upper bounded by t^* :*

$$\tau^n \leq t^* = \left\lceil \frac{2C(q+2)}{\alpha\beta} \right\rceil; \tag{33}$$

⁶ A function f grows μ -quadratically if $f(z) - f^* \geq \mu \mathbf{dist}(z, Z^*)^2$.

ii. The distance to primal-dual solution set decays linearly:

$$\mathbf{dist}(z^{n,0}, Z^*) \leq \beta^n \frac{t^*}{\tau^0} \mathbf{dist}(z^{0,0}, Z^*).$$

Proof We prove (33) by contradiction. If it is not satisfied, there exists $n \geq 1$ such that $\tau^n > t^*$, and it holds for $t = t^*$ that

$$\begin{aligned} \rho_{\|\bar{z}^{n,t^*} - z^{n,0}\|}(\bar{z}^{n,t^*}) &\leq \frac{2C \|\bar{z}^{n,t^*} - z^{n,0}\|}{t^*} \leq \frac{2C \|\bar{z}^{n,t^*} - z^{n,0}\|}{\alpha t^* \mathbf{dist}(z^{n,0}, Z^*)} \rho_{\|z^{n,0} - z^{n-1,0}\|}(z^{n,0}) \\ &\leq \frac{2C(q+2)}{\alpha t^*} \rho_{\|z^{n,0} - z^{n-1,0}\|}(z^{n,0}) \leq \beta \rho_{\|z^{n,0} - z^{n-1,0}\|}(z^{n,0}), \end{aligned}$$

where the first inequality utilizes Property i., the second inequality follows from sharpness of the primal-dual problem (13) by choosing $r = \|z^{n,0} - z^{n-1,0}\|$ and $z = z^{n,0} \in S$, the third inequality utilizes Property ii., and the last inequality is from the definition of t^* . This is a contradiction: the restart condition (30) should have been triggered at $t = t^* < \tau^n$. Therefore (33) holds.

For ii., it follows from sharpness condition that

$$\begin{aligned} \mathbf{dist}(z^{n,0}, Z^*) &\leq \frac{\rho_{\|z^{n,0} - z^{n-1,0}\|}(z^{n,0})}{\alpha} \leq \beta^{n-1} \frac{\rho_{\|z^{1,0} - z^{0,0}\|}(z^{1,0})}{\alpha} \\ &\leq \beta^{n-1} \frac{2C \|z^{1,0} - z^{0,0}\|}{\alpha \tau^0} \leq \beta^n \frac{2C(q+2)}{\alpha \beta \tau^0} \mathbf{dist}(z^{0,0}, Z^*), \end{aligned}$$

where the second inequality recursively utilizes the restarting condition (30), the third inequality is from Property i. and the last inequality utilizes Property ii. □

Remark 5 (Choice of β) From Theorem 2 we can obtain

$$\sum_{i=0}^n \tau^i \leq \tau^0 + \frac{t^*}{\log(1/\beta)} \left(\log \left(\frac{t^* \mathbf{dist}(z^{0,0}, Z^*)}{\tau^0 \mathbf{dist}(z^{n,0}, Z^*)} \right) \right). \tag{34}$$

If we assume $\frac{\mathbf{dist}(z^{0,0}, Z^*)}{\mathbf{dist}(z^{n,0}, Z^*)} \gg t^*/\tau^0$ and use $t^* \approx \frac{2C(q+2)}{\alpha \beta}$ the right hand side of (34) is approximately:

$$\tau^0 + \frac{\beta}{\log(1/\beta)} \frac{2C(q+2)}{\alpha} \log \left(\frac{\mathbf{dist}(z^{0,0}, Z^*)}{\mathbf{dist}(z^{n,0}, Z^*)} \right),$$

optimizing this with respect to β yields $\beta = \exp(-1)$.

Compared with the assumptions for fixed frequency restarts (Theorem 1), Theorem 2 additionally requires that $z^{n,0}$ stays in a set $S \subseteq Z$, where the primal-dual problem is sharp. Suppose the feasible region Z is bounded, we can simply choose $S = Z$. However, in the LP example (i.e., Lemma 5) where Z can be unbounded, and the

primal-dual problem is sharp only on a bounded region. Below we show that the iterates of restarted PDHG, PPM, EGM and ADMM stay in a bounded region:

Proposition 9 *Consider the sequence $\{z^{n,0}\}_{i=0}^\infty, \{\tau^n\}_{i=1}^\infty$ generated by Algorithm 1 with proper step-size η (see Table 1) and $\beta \in (0, 1)$. For restarted PDHG or restarted PPM, there exists a constant $R > 0$ such that $z^{n,0} \in W_R(z^{0,0})$ for all $n \in \mathbb{N}$.*

Proof For PDHG and PPM, the proof uses the non-expansiveness property of $\hat{z}^{t,n} = z^{t,n}$ (see Proposition i.). In particular, let $z^* \in \operatorname{argmin}_{z \in Z^*} \|z - z^{0,0}\|$, then it holds for any

n, t that

$$\begin{aligned} \|z^* - \bar{z}^{n,t}\| &= \left\| \frac{1}{t} \sum_{i=1}^t (z^* - z^{n,i}) \right\| \leq \frac{1}{t} \sum_{i=1}^t \|z^* - z^{n,i}\| \\ &\leq \frac{1}{t} \sum_{i=1}^t \|z^* - z^{n,0}\| = \|z^* - z^{n,0}\| \end{aligned}$$

where the first inequality is from triangle inequality and the second inequality uses Proposition i.. Thus, we have

$$\|z^* - z^{n+1,0}\| = \|z^* - \bar{z}^{n,\tau^n}\| \leq \|z^* - z^{n,0}\|$$

whereby $\|z^* - z^{n,0}\| \leq \|z^* - z^{0,0}\|$ by induction. It then follows from the triangle inequality that

$$\|z^{n,0} - z^{0,0}\| \leq \|z^{0,0} - z^*\| + \|z^* - z^{n,0}\| \leq 2\|z^{0,0} - z^*\| = 2\mathbf{dist}(z^{0,0}, Z^*)$$

which concludes the proof for PDHG and PPM with $R = 2\mathbf{dist}(z^{0,0}, Z^*)$. □

Proposition 10 *Consider the sequence $\{z^{n,0}\}_{i=0}^\infty, \{\tau^n\}_{i=1}^\infty$ generated by Algorithm 1 with the adaptive restart scheme, namely, we restart the outer loop if (30) holds. If $0 < \beta < \frac{1}{q+3}$ then, there exists a constant $R > 0$ such that $z^{n,0} \in W_R(z^{0,0})$ for all $n \in \mathbb{N}$ in the following settings:*

- restarted EGM applied to LP (Lemma 5);
- restarted ADMM applied to LP (Lemma 6).

The proof of Proposition 10 is more technical than Proposition 9 because $\hat{z}^{t,n} = z^{t,n}$ no longer holds, and we defer it to ‘‘Appendix D’’.

5 Tight convergence results for bilinear problems

Bilinear problems are a special case of LP where there are no nonnegativity or inequality constraints. Moreover, in many situations the asymptotic behaviour of primal-dual methods applied to LP is equivalent to applying the method to a bilinear problem. In

particular, if after a certain number of iterations the set of active variables⁷ does not change, then the subset of the matrix, objective, and right hand side corresponding to this active set defines such a bilinear subproblem. Recent results [44, 46] show that for many primal-dual algorithms, including PDHG and ADMM, if the problem is non-degenerate⁸ then after a finite number of iterations the active set freezes. See [44] for a differential geometry argument for this phenomenon. Thus, the convergence of methods for bilinear problems can characterize their asymptotic behavior on nondegenerate LP problems.

Already we have some basic results in this setup. In particular, Lemma 3, and Proposition 7 imply that the primal-dual problem is $\Omega(\sigma_{\min}^+(A))$ -sharp on bilinear problems with respect to the PDHG norm for $\eta \leq 1/(2\sigma_{\max}(A))$. Combining the sharpness result with Table 1, and Theorem 1 (or Theorem 2 with a small loss in the log term) implies that after

$$O\left(\frac{\sigma_{\max}(A)}{\sigma_{\min}^+(A)} \log\left(\frac{1}{\epsilon}\right)\right)$$

total matrix–vector multiplications restarted PDHG finds a point $z^{n,0}$ satisfying

$$\frac{\text{dist}(z^{n,0}, Z^*)}{\text{dist}(z^{0,0}, Z^*)} \leq \epsilon.$$

The remainder of this section proceeds as follows. First we present a lower bound showing that, up to a constant factor, restarted PDHG yields the best convergence rates for solving bilinear problems. Next, we show tight convergence bounds for the average and last iterate of PDHG which we use to demonstrate our restarted PDHG convergence rate represents a strict improvement over the latter two methods. Finally, we compare the performance of other methods from literature for solving bilinear problems.

5.1 Lower bounds for primal-dual algorithms

In this subsection, we present a lower bound for solving a sharp primal-dual problem (1) demonstrating that restarting gives the best possible worst-case convergence bound for a large class of primal-dual methods. First, we review related lower bounds for first-order unconstrained function minimization.

Definition 2 An algorithm is span-respecting for a minimization problem $\min_{x \in \mathbb{R}^m} f(x)$ if

$$x^t \in x^0 + \text{Span}\left\{\nabla f(x^i) : \forall i \in \{1, \dots, t-1\}\right\}$$

⁷ The active variables are variables not at their bounds.

⁸ In the LP case (19), non-degeneracy means that the algorithm converges to a primal-dual solution that satisfies strict complementarity.

for all $t \in \mathbb{N}$.

Span-respecting algorithms cover a large number of unconstrained optimization algorithms including gradient descent, conjugate gradient, and accelerated gradient descent. As is well-known one can form a lower bound using convex quadratics for these types of algorithms as given in Theorem 3.

Theorem 3 (Theorem 2.1.12 of [61]) *For all $\gamma_{\max} > \gamma_{\min} > 0$, $m \in \mathbb{N}$, there exists a positive definite matrix $H \in \mathbb{R}^{m \times m}$ and vector $h \in \mathbb{R}^m$ such that $\sigma_{\max}(H) = \gamma_{\max}$, $\sigma_{\min}(H) = \gamma_{\min}$, and any span-respecting unconstrained minimization algorithm for solving*

$$\min_{x \in \mathbb{R}^m} \frac{1}{2} x^\top H x + h^\top x. \tag{35}$$

satisfies for $t < m$ that

$$\mathbf{dist}(x^t, X^*) \geq \left(1 - \sqrt{\frac{\sigma_{\min}(H)}{\sigma_{\max}(H)}} \right)^t \mathbf{dist}(x^0, X^*).$$

Our goal is to supply a similar result to Theorem 3 but in the primal-dual setting.

Definition 3 An algorithm is span-respecting for a unconstrained primal-dual problem $\max_{y \in \mathbb{R}^n} \min_{x \in \mathbb{R}^m} \mathcal{L}(x, y)$ if:

$$\begin{aligned} x^t &\in x^0 + \text{Span} \left\{ \nabla_x \mathcal{L}(x^i, y^j) : \forall i, j \in \{1, \dots, t-1\} \right\} \\ y^t &\in y^0 + \text{Span} \left\{ \nabla_y \mathcal{L}(x^i, y^j) : \forall i \in \{1, \dots, t\}, j \in \{1, \dots, t-1\} \right\}. \end{aligned}$$

Definition 3 is analogous to Definition 2 but in the primal-dual setting. If \mathcal{L} is bilinear then with appropriate indexing of their iterates, primal-dual algorithms including primal-dual hybrid gradient, extragradient and their restarted variants satisfy Definition 3. Corollary 1 provides a lower bound on span-respecting primal-dual algorithms. The proof is by reduction to Theorem 3.

Corollary 1 *Let $\| \cdot \|$ be the Euclidean norm. For all $\alpha, C \in (0, \infty)$, $m \in \mathbb{N}$, there exists a matrix $A \in \mathbb{R}^{m \times m}$ and vector $b \in \mathbb{R}^m$ such that*

$$\max_{y \in \mathbb{R}^m} \min_{x \in \mathbb{R}^m} \mathcal{L}(x, y) = y^\top A x + b^\top y$$

is α -sharp on $\mathbb{R}^{m \times m}$, C -smooth and any span-respecting unconstrained primal-dual algorithm satisfies

$$\mathbf{dist}(z^t, Z^*) \geq \left(1 - \frac{\alpha}{C} \right)^t \mathbf{dist}(z^0, Z^*)$$

for $t < m$.

Proof Consider initial solution $y^0 = x^0 = \mathbf{0}$, where A and b . Then, for any span-respecting primal-dual method satisfies

$$\begin{aligned} x^t &\in \text{Span} \left\{ A^\top y^0, \dots, A^\top y^t \right\} \\ y^t &\in \text{Span} \left\{ Ax^0 - b, \dots, Ax^{t-1} - b \right\}, \end{aligned}$$

thus it holds that

$$x^t \in \text{Span} \left\{ A^\top (Ax^0 - b), \dots, A^\top (Ax^{t-1} - b) \right\}. \quad (37)$$

Let $\gamma_{\max} = C^2$, $\gamma_{\min} = \alpha^2$, and consider the positive definite matrix H and vector h supplied by Theorem 3. Define $A = H^{1/2}$, $b = H^{-1/2}b$. Note $H^{1/2}$ exists and is positive definite as H is positive definite. Also, note that the unique solution to the primal-dual problem is $x^* = A^{-1}b$ and $y^* = A^{-1}\mathbf{0} = \mathbf{0}$ with x^* matching the optimal solution to (35). Then by (37) we have $x^t \in \text{Span} \{ Hx^0 - h, \dots, Hx^{t-1} - h \}$ which implies x^t is span-respecting for (35). Applying the iteration bound of Theorem 3 and using that $y^* = y^0 = \mathbf{0}$ yields the result. \square

Therefore we have established that the convergence bounds of our restart schemes (Theorems 1 and 2) matches the lower bound (Corollary 1) up to a constant. In other words, a restarted method gives the best worst-case bounds.

5.2 Tight bounds for the average and last iterate

Previous section shows that our restart schemes match the lower bound (up to a constant) for solving sharp primal-dual problems. However, at this point it is unclear if standard approaches, e.g., the last and average iterate do not also obtain this worst-case bound with a careful analysis for sharp primal-dual problems. In this section, we show for bilinear problems that the last iterate and the average iterate of PDHG cannot match convergence bounds of restarted PDHG. Such statement is genuinely the case for other primal-dual first-order methods with a similar arguments.

Simple empirical tests support the hypothesis that the restarted PDHG is genuinely faster than the average or last iterate. In particular, Fig. 1 plots the last iterate (z_t) and the average iterate (\bar{z}_t) of both non-restarted PDHG (i.e., PDHG without restarting) and restarted PDHG (with fixed frequency) for solving a simple two-dimensional bilinear problem $\mathcal{L}(x, y) = xy$ with step-size $\eta = 0.2$. As we can see, the average solution of the restarted algorithm has faster convergence to the primal-dual solution $(0, 0)$ than the average or last iterate of PDHG.

We can formalize this improvement from restarted PDHG. Table 2 compares the complexity of the last iterate of PDHG, the average iterate of PDHG, and restarted PDHG on bilinear problems. From the table we can see that only the last iterate achieves linear convergence and its linear convergence constant is a square root factor worse than restarted PDHG. We expect similar results would also apply to other primal-dual algorithms, such as EGM. The remainder of this subsection is devoted to proving the results in Table 2.

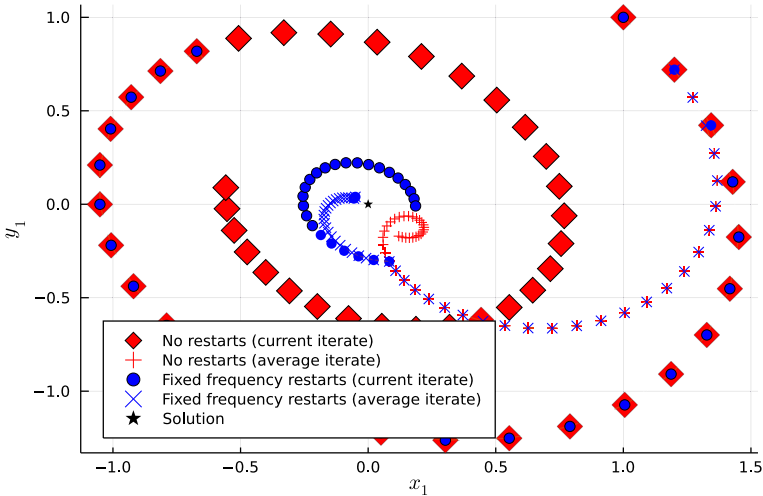


Fig. 1 Plot of the first 50 iterates of non-restarted and restarted PDHG for a simple two-dimensional bilinear problem $\mathcal{L}(x, y) = xy$ with $\eta = 0.2$. The restart length is for the fixed frequency restarts is 25. The unique optimal solution is $(0, 0)$

Table 2 Comparison of the convergence of the last iterate, average iterate, and restarted PDHG on bilinear problems with $\eta = 1/(2\sigma_{\max}(A))$. The algorithms terminate when $\frac{\|z^t - z^*\|_2^2}{\|z^0 - z^*\|_2^2} \leq \epsilon$. We also assume $\epsilon \in (0, 1/2]$. Additionally for the average iterate lower bound we assume $\epsilon \leq \sigma_{\min}^+(A)^2/\sigma_{\max}(A)^2$

Last iterate	Average iterate	Restarted
$\Theta\left(\left(\frac{\sigma_{\max}(A)}{\sigma_{\min}^+(A)}\right)^2 \log\left(\frac{1}{\epsilon}\right)\right)$	$\Theta\left(\frac{\sigma_{\max}(A)}{\sigma_{\min}^+(A)\epsilon}\right)$	$\Theta\left(\frac{\sigma_{\max}(A)}{\sigma_{\min}^+(A)} \log\left(\frac{1}{\epsilon}\right)\right)$

Consider an bilinear problem with $\mathcal{L}(x, y) = y^\top Ax$. Without loss of generality, we can assume A is a diagonal matrix with entries $0 < \sigma_1 \leq \dots \leq \sigma_m$ and $x, y \in \mathbb{R}^m$ (See the beginning of the proof of Theorem 6 in Applegate et. al [6] for an explanation.)

The unique optimal solution to this primal-dual problem is $(0, 0)$. PDHG with step-size $\eta \leq \frac{1}{\sigma_m}$ has iterate update

$$z^{t+1} = \begin{pmatrix} 1 & -\eta A^\top \\ \eta A & 1 - 2\eta^2 AA^\top \end{pmatrix} z^t, \tag{38}$$

where $z^t = \begin{pmatrix} x^t \\ y^t \end{pmatrix}$. We study the convergence of three sequences: the last iterates of PDHG (z^t), the average iterates of PDHG ($\bar{z}^K = \frac{1}{K} \sum_{t=1}^K z^t$) and the solution $\bar{z}^{n,t}$ obtained in our restart scheme with PDHG as a subroutine.

Notice the PDHG dynamic (38) is separable over coordinates, and there are m independently evolving 2-d linear dynamical systems given by:

$$z_i^{t+1} = P_i z_i^t, \tag{39}$$

for $i = 1, \dots, m$ where $z_i = \begin{pmatrix} x_i \\ y_i \end{pmatrix}$ and $P_i = \begin{pmatrix} 1 & -\eta\sigma_i \\ \eta\sigma_i & 1 - 2\eta^2\sigma_i^2 \end{pmatrix}$. Let $P_i = Q_i \Gamma_i Q_i^{-1}$ be the eigen-decomposition of P_i , where

$$\Gamma_i = \begin{pmatrix} 1 - \eta^2\sigma_i^2 - \mathbf{i}\eta\sigma_i\sqrt{1 - \eta^2\sigma_i^2} & \\ & 1 - \eta^2\sigma_i^2 + \mathbf{i}\eta\sigma_i\sqrt{1 - \eta^2\sigma_i^2} \end{pmatrix}$$

$$Q_i = \begin{pmatrix} \eta\sigma_i - \mathbf{i}\sqrt{1 - \eta^2\sigma_i^2} & \eta\sigma_i + \mathbf{i}\sqrt{1 - \eta^2\sigma_i^2} \\ 1 & 1 \end{pmatrix}$$

and \mathbf{i} denotes the imaginary part of a complex number.

Moreover, with $\tilde{z}_i^t = Q_i^{-1} z_i^t$, $B_i = (Q_i^{-1})^\dagger Q_i^{-1}$ and $\|v\|_M = \sqrt{v^\dagger M v}$, where \dagger is the conjugate transpose, we get

$$\|z_i^t - z_i^*\|_{B_i}^2 = \|z_i^t\|_{B_i}^2 = \|Q_i^{-1} z_i^t\|_2^2 = \|\tilde{z}_i^t\|_2^2 = \|\Gamma_i^t \tilde{z}_i^0\|_2^2$$

$$= (1 - \eta^2\sigma_i^2)^t \|\tilde{z}_i^0\|_2^2 = (1 - \eta^2\sigma_i^2)^t \|z_i^0 - z_i^*\|_{B_i}^2, \tag{40}$$

where the second to last equality uses

$$|\gamma_i^\pm|^2 = (1 - \eta^2\sigma_i^2)^2 + \eta^2\sigma_i^2(1 - \eta^2\sigma_i^2) = 1 - \eta^2\sigma_i^2. \tag{41}$$

Moreover, note that

$$Q_i Q_i^\dagger = \begin{pmatrix} \eta\sigma_i - \mathbf{i}\sqrt{1 - \eta^2\sigma_i^2} & \eta\sigma_i + \mathbf{i}\sqrt{1 - \eta^2\sigma_i^2} \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \eta\sigma_i + \mathbf{i}\sqrt{1 - \eta^2\sigma_i^2} & 1 \\ \eta\sigma_i - \mathbf{i}\sqrt{1 - \eta^2\sigma_i^2} & 1 \end{pmatrix}$$

$$= 2 \begin{pmatrix} 1 & \eta\sigma_i \\ \eta\sigma_i & 1 \end{pmatrix}$$

with eigenvalues $2 \pm 2\sigma_i\eta$. Therefore, as $\eta \in [0, 1/(2\sigma_m)]$ we get $1 \leq 2 \pm 2\sigma_i\eta \leq 3$, which implies

$$\|v\|_2^2 \leq v^\dagger Q_i Q_i^\dagger v \leq 3\|v\|_2^2 \Rightarrow \frac{\|v\|_2^2}{3} \leq v^\dagger B_i v \leq \|v\|_2^2 \tag{42}$$

for all $v \in \mathcal{C}^n$ where \mathcal{C} is the set of complex numbers and the implication uses that $B_i Q_i Q_i^\dagger = \mathbf{I}$.

By (42) we get $\frac{1}{3}\|z^t - z^*\|_2^2 \leq \sum_{i=1}^n \|z_i - z_i^*\|_{B_i}^2 \leq \|z^t - z^*\|_2^2$ and therefore (40) implies that we reach a solution z^t such that

$$\frac{\|z^t - z^*\|_2^2}{\|z^0 - z^*\|_2^2} \leq \epsilon$$

within $O\left(\left(\frac{1}{\eta\sigma_1}\right)^2 \log\left(\frac{1}{\epsilon}\right)\right)$ iterations.

From (42) and (40) we deduce that

$$\|z_1^t - z_1^*\|_2^2 \geq \frac{1}{3}(1 - \eta^2\sigma_1^2)^t \|z_1^t - z_1^*\|_2^2.$$

Therefore by $\sigma_1 \leq \sigma_m$ and $\eta = 1/(2\sigma_m)$ we get $\Omega\left(\left(\frac{\sigma_m}{\sigma_1}\right)^2 \log\left(\frac{1}{\epsilon}\right)\right)$ as a lower bound for the last iterate convergence of PDHG.

Define $\|M\|_2 := \sup_{\|v\|_2=1} \|Mv\|_2$ be the ℓ_2 norm of a matrix M . Now we study the solution of the average iterates $\bar{z}^K = \frac{1}{K} \sum_{t=1}^K z^t$, and we claim \bar{z}^K converges to the unique optimal solution $z^* = 0$ with sublinear complexity $\Omega\left(\frac{\sigma_m}{\sigma_1\epsilon}\right)$. Notice that

$$\begin{aligned} & \left\| \bar{z}_i^K - z_i^* \right\|_{B_i} \\ &= \left\| \frac{1}{K} \sum_{t=1}^K z_i^t \right\|_{B_i} \\ &= \left\| \frac{1}{K} \sum_{t=1}^K \tilde{z}_i^t \right\|_2 \\ &= \frac{1}{K} \left\| \sum_{t=1}^K \Gamma_i^t \tilde{z}_i^0 \right\|_2 \\ &= \frac{1}{K} \left\| \Gamma_i (\mathbf{I} - \Gamma_i)^{-1} (\mathbf{I} - \Gamma_i^K) \tilde{z}_i^0 \right\|_2 \\ &\leq \frac{1}{K} \|\Gamma_i\|_2 \left\| (\mathbf{I} - \Gamma_i)^{-1} \right\|_2 \left\| \mathbf{I} - \Gamma_i^K \right\|_2 \left\| \tilde{z}_i^0 \right\|_2 \\ &\leq \frac{1}{K(\eta\sigma_1)} \|\Gamma_i\|_2 \left(\|\mathbf{I}\|_2 + \left\| \Gamma_i^K \right\|_2 \right) \left\| \tilde{z}_i^0 \right\|_2 \\ &\leq \frac{2}{K\eta\sigma_1} \left\| \tilde{z}_i^0 \right\|_2 \\ &= \frac{2}{K\eta\sigma_1} \left\| z_i^0 - z_i^* \right\|_{B_i}, \end{aligned} \tag{43}$$

where the second inequality uses

$$|1 - \gamma_i^\pm| = \sqrt{\eta^4\sigma_i^4 + \eta^2\sigma_i^2(1 - \eta^2\sigma_i^2)} = \eta\sigma_i \tag{44}$$

and the third inequality utilizes $\|\Gamma_i\|_2 \leq 1$. Setting $\eta = \frac{1}{2\sigma_m}$ and using (42) demonstrates \bar{z}^t converges to z^* in $O\left(\frac{\sigma_m}{\sigma_1\epsilon}\right)$ iterations, which is consistent with the ergodic rate of PDHG.

We will now establish a matching $\Omega\left(\frac{\sigma_m}{\sigma_1 \epsilon}\right)$ lower bound for the average iterate. Observe that

$$\begin{aligned} \left|1 - (\gamma_i^\pm)^K\right| &\geq 1 - |\gamma_i^\pm|^K \\ &\geq 1 - (1 - \eta^2 \sigma_i^2)^{K/2} \\ &\geq 1 - \frac{1}{1 + (K/2)\eta^2 \sigma_i^2} \\ &\geq \frac{K \eta^2 \sigma_i^2}{2 + K \eta^2 \sigma_i^2}. \end{aligned} \tag{45}$$

where the third inequality uses the fact that $(1 - a)^b \leq \frac{1}{1+ba}$ for $a \geq 0, b \geq 0$ (to see this note that $f(a, b) = (1 - a)^b(1 + ba)$ satisfies $f(0, b) = 1$ and $\frac{\partial f}{\partial a} \leq 0$). Therefore,

$$\begin{aligned} &\left\| \bar{z}_1^K - z_1^* \right\|_{B_1} \\ &= \frac{1}{K} \left\| \Gamma_1 (\mathbf{I} - \Gamma_1)^{-1} (\mathbf{I} - \Gamma_1^K) \bar{z}_1^0 \right\|_2 \\ &= \frac{1}{K} \sqrt{(\bar{x}_1^0)^2 |\gamma_1^-|^2 |1 - \gamma_1^-|^{-2} |1 - (\gamma_1^-)^K|^2 + (\bar{y}_1^0)^2 |\gamma_1^+|^2 |1 - \gamma_1^+|^{-2} |1 - (\gamma_1^+)^K|^2} \\ &\geq \frac{1}{K} \|\bar{z}_1^0\|_2 |\gamma_1^\pm| |1 - \gamma_1^\pm|^{-1} (1 - |\gamma_1^\pm|^K) \\ &\geq \frac{\|\bar{z}_1^0\|_2 \sqrt{1 - \eta^2 \sigma_1^2}}{K \eta \sigma_1} \times \frac{K \eta^2 \sigma_1^2}{2 + K \eta^2 \sigma_1^2} \\ &\geq \frac{\|\bar{z}_1^0\|_2 \sqrt{3/4}}{K \eta \sigma_1} \times \frac{K \eta^2 \sigma_1^2}{2 + K \eta^2 \sigma_1^2} \\ &= \frac{\sqrt{3}}{2} \times \frac{\eta \sigma_1 \|z_1^0 - z_1^*\|_{B_1}}{2 + K \eta^2 \sigma_1^2}, \end{aligned}$$

where the second equality uses that the absolute value is completely multiplicative, the second inequality uses (41), (44) and (45) and the final inequality uses $\eta = \sigma_m/2$. This establishes the lower bound for the average iterate. We note that this lower bound only matches the upper bound for $K = \Omega\left(\frac{1}{\eta^2 \sigma_1^2}\right)$. We believe matching the upper bound across all K is possible but the analysis appears to be much more technical.

5.3 Comparison with other methods

Mokhtari et al. [56] study EGM and the optimistic gradient descent, in this setting they show an

$$O\left(\frac{\sigma_{\max}(A^\top A)}{\sigma_{\min}(A^\top A)} \log(1/\varepsilon)\right) = O\left(\left(\frac{\sigma_{\max}(A)}{\sigma_{\min}(A)}\right)^2 \log(1/\varepsilon)\right)$$

bound for both methods on bilinear problems (recall $\sigma_{\min}(A^\top A) = \sigma_{\min}(A)$ and $\sigma_{\max}(A^\top A) = \sigma_{\max}(A)$). As expected, this matches the last iterate performance of PDHG.

We can use other methods to solve bilinear problems. One approach is reduce the bilinear game to a pair of nonsmooth optimization problems:

$$\begin{aligned} &\min_{x \in \mathbb{R}^m} \|Ax - b\|_2 \\ &\min_{y \in \mathbb{R}^m} \|A^\top y - c\|_2. \end{aligned}$$

Each of these problems is $\sigma_{\min}^+(A)$ -sharp (Lemma 2) and $\sigma_{\max}(A)$ -Lipschitz. Therefore according to [78, Corollary 9] restarted subgradient method requires

$$O\left(\left(\frac{\sigma_{\max}(A)}{\sigma_{\min}^+(A)}\right)^2 \log(1/\varepsilon)\right)$$

iterations on each problem to find an approximate solution satisfying $\frac{\|Ax-b\|_2}{\|Ax_0-b\|_2} \leq \varepsilon$ and $\frac{\|A^\top y_0-c\|_2}{\|A^\top y_0-c\|_2} \leq \varepsilon$. Another approach is to reformulate solving bilinear game as a pair of smooth unconstrained problems:

$$\begin{aligned} &\min_{x \in \mathbb{R}^m} \|Ax - b\|_2^2 \\ &\min_{y \in \mathbb{R}^m} \|A^\top y - c\|_2^2. \end{aligned}$$

One can apply conjugate gradient or accelerated gradient descent to each of these subproblems. The well-known convergence bounds for these methods [61] imply that they have the following iteration bounds

$$O\left(\sqrt{\frac{\sigma_{\max}(A^\top A)}{\sigma_{\min}^+(A^\top A)}} \log(1/\varepsilon)\right) = O\left(\frac{\sigma_{\max}(A)}{\sigma_{\min}^+(A)} \log(1/\varepsilon)\right)$$

which matches the bounds we obtained by restarted PDHG. We also point the reader to Gilpin et al. [29], who obtain similar results to ours in the LP (with bounded feasible region) and bilinear setting. The main difference is that they modify Nesterov’s smoothing algorithm for nonsmooth convex optimization [59], which is a primal method, whereas we combine restarts with primal-dual methods to achieve our result.

6 Efficient computation of the normalized duality gap

6.1 Separable norm

This subsection shows that the normalized duality gap can be computed efficiently under the mild assumption that the norm is separable:

$$\|(x, y)\|^2 = \|x\|^2 + \|y\|^2.$$

To evaluate (4a) we need to find a solution to

$$\max_{\hat{z} \in W_r(z)} \mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y). \quad (46)$$

Recall that due to Proposition 8 solving (4b) is unnecessary for adaptive restarts.

Solving (46) can be done using standard methods which we briefly outline here. By the assumption that \mathcal{L} is convex-concave and by duality, there exists a $\lambda \in [0, \infty)$ such that (46) is equivalent to

$$\max_{(\hat{x}, \hat{y}) \in Z} \mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y) - \frac{\lambda}{2} (\|\hat{x} - x\|^2 + \|\hat{y} - y\|^2). \quad (47)$$

For any $\lambda \in (0, \infty)$, define

$$h(\lambda) := \|\bar{z}(\lambda) - z\| - r$$

where $\bar{z}(\tau)$ is the unique solution to (47) (the solution is unique because \mathcal{L} is convex-concave which means (47) is strongly concave). Observe the function $h(\lambda)$ is nonincreasing. If $h(\lambda) < 0$ for all $\lambda \in (0, \infty)$ then the solution to (47) with $\lambda = 0$ that is closest to z solves (46) (and can be approximated by picking an arbitrarily tiny λ). If $h(\lambda) > 0$ for all $\lambda \in (0, \infty)$ then $r = 0$, and the optimal solution to (46) is $\hat{z} = z$. The final possibility is that there exist $\lambda, \lambda' \in (0, \infty)$ such that $h(\lambda) < 0 < h(\lambda')$ in which case there exists a root of h which we can find using a standard root finding methods such as the bisection method. This root corresponds to the optimal solution of (46).

Problem (47) is separable into

$$\operatorname{argmin}_{\hat{x} \in X} \mathcal{L}(\hat{x}, y) + \frac{\lambda}{2} \|\hat{x} - x\|^2 \quad (48a)$$

$$\operatorname{argmin}_{\hat{y} \in Y} \mathcal{L}(x, \hat{y}) + \frac{\lambda}{2} \|\hat{y} - y\|^2, \quad (48b)$$

which is equivalent to the subproblem iterations for PPM, PDHG, and ADMM (Sect. 2). Even if this root finding only requires evaluating (48) a small number of times it could significantly increase the cost of running the algorithm. Therefore, it makes sense to only test condition (30) at fixed intervals with large separation, e.g., every 200 iterations.

6.2 Primal-dual hybrid gradient

Section 6.1 considered the case where $\|\cdot\|$ is separable. However, the (natural) norm that we previously used to analyze PDHG is $\|z\| = \sqrt{\|x\|_2^2 - 2\eta y^\top Ax + \|y\|_2^2}$ which is not separable. Fortunately, the following Corollary shows that Property 3 also holds for PDHG with the Euclidean norm. This allows us the possibility of computing ρ_r with $\|\cdot\|$ being the Euclidean norm and retaining our theoretical guarantees. Indeed this is how we implement the adaptive restart scheme for PDHG in Sect. 7.

Corollary 2 *For PDHG with $\eta \in (0, 1/\sigma_{\max}(A))$, Property 3 holds when $\|\cdot\|$ is the Euclidean norm, $C = \frac{2}{\eta(1-\eta\sigma_{\max}(A))}$ and $q = 4\frac{1+\eta\sigma_{\max}(A)}{1-\eta\sigma_{\max}(A)}$.*

The proof uses that Property 3 holds for the PDHG norm, and Proposition 7 which establishes that $\|z\|$ is approximately Euclidean. The proof appears in ‘‘Appendix E’’.

By Proposition 9 we also know that the iterates will remain bounded even after this change of norm. Therefore, using Lemma 5 we have established the premise of Theorem 2 for PDHG with the Euclidean norm on LP problems.

6.3 Linear time implementation for linear programming

This subsection shows that we can more efficiently compute the normalized duality gap in the LP setting.

6.3.1 Linear time algorithm for linear trust region problems

The key ingredient to the linear time computation of the normalized duality gap is a linear time algorithm for linear trust region problems. In particular, consider a trust region problem with a linear objective g , center z , Euclidean norm, and lower bounds l (possibly infinite) on the variables:

$$\operatorname{argmin}_{\hat{z} \in \mathbb{R}^{n+m} : l \leq \hat{z}, \|\hat{z} - z\| \leq r} g^\top \hat{z} . \tag{49}$$

If $g_i < 0$ then setting $g_i = -g_i$ and $l_i = -\infty$ will create an equivalent problem satisfying $g_i > 0$. Therefore without loss of generality we can assume $g_i > 0$ for all i .

The optimal \hat{z} will be of the form

$$\hat{z}(\lambda) := \max(z - \lambda g, l) , \tag{49}$$

for some $\lambda \in [0, \infty)$, so (49) reduces to the univariate problem

$$\max \lambda : \|\hat{z}(\lambda) - z\| \leq r . \tag{51}$$

Letting $\hat{\lambda}_i$ be the value of λ at which $\hat{z}(\lambda)_i$ encounters the bound, that is, $\hat{\lambda}_i := (z_i - l_i)/g_i$, we have

$$\|\hat{z}(\lambda) - z\|^2 = \sum_{i : \hat{\lambda}_i \leq \lambda} (l_i - z_i)^2 + \lambda^2 \sum_{i : \hat{\lambda}_i > \lambda} g_i^2. \tag{52}$$

This forms the basis for a linear time algorithm: by computing (52) for the median value λ_{med} of $\hat{\lambda}_i$ (in linear time [9]), and comparing it with r^2 , we determine whether or not the optimal λ for (50) is greater or less than λ_{med} . This lets us eliminate half of the $\hat{\lambda}_i$, collapsing their contributions into the appropriate summand. Applying this recursively gives a running time recurrence $f(n + m) = O(n + m) + f((n + m)/2)$, i.e., $f(n + m) = O(n + m)$. Details are in ‘‘Appendix F’’.

6.3.2 Extragradient and primal-dual hybrid gradient

Recall that for extragradient the norm for defining ρ_r is Euclidean. For LP with extragradient (Sect. 3.3), we use the Lagrangian:

$$\min_{x \geq 0} \max_{y \in \mathbb{R}^m} \mathcal{L}(x, y) = c^\top x + y^\top b - y^\top Ax \tag{53}$$

for which the trust region problem (46) is equivalent to

$$\underset{\hat{x} \geq 0, \hat{y} \in \mathbb{R}^m : \|(\hat{x}, \hat{y}) - (x, y)\| \leq r}{\text{argmin}} (c^\top - y^\top A)\hat{x} - (b^\top - (Ax)^\top)\hat{y}. \tag{54}$$

which is a trust region problem with a linear objective, Euclidean norm, and lower bounds on variables (the dual lower bounds are $-\infty$). Given $y^\top A$ and Ax , this can be solved in time $O(n + m)$ using the algorithm in Sect. 6.3.1. The same result hold for PDHG with Euclidean norm as described in Corollary 2.

6.3.3 ADMM

Recall for ADMM for LP (Lemma 6) that

$$\min_{x_U \in X_U, x_V \in X_V} \max_y \mathcal{L}(x, y) = c^\top x_V - y^\top (\mathbf{I} - \mathbf{I}) \begin{pmatrix} x_U \\ x_V \end{pmatrix}$$

and $\|\cdot\| = \|\cdot\|_M$ with

$$M = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \eta \mathbf{I} & 0 \\ 0 & 0 & \frac{1}{\eta} \mathbf{I} \end{pmatrix}.$$

Therefore, for $Q_r(x_V, y) = \{(\hat{x}_V, \hat{y}) : \hat{x}_V \geq 0, \eta \|x_V - \hat{x}_V\|_2^2 + \frac{1}{\eta} \|y - \hat{y}\|_2^2 \leq r^2\}$ we have

$$\begin{aligned} \rho_r(\bar{z}^t) &= \frac{1}{r} \sup_{\hat{z} \in W_r(\bar{z}^t)} c \cdot (\bar{x}_V^t - \hat{x}_V) + \bar{y}^t \cdot (\hat{x}_U - \hat{x}_V) - \hat{y} \cdot (\bar{x}_U^t - \bar{x}_V^t) \\ &= \frac{1}{r} \left(\sup_{\hat{x}_U \in X_U} \bar{y}^t \cdot \hat{x}_U + \sup_{(\hat{x}_V, \hat{y}) \in Q_r(\bar{x}_V^t, \bar{y}^t)} c \cdot (\bar{x}_V^t - \hat{x}_V) - \bar{y}^t \cdot \hat{x}_V - \hat{y} \cdot (\bar{x}_U^t - \bar{x}_V^t) \right) \\ &= \frac{1}{r} \left(\bar{y}^t \cdot \bar{x}_U^t + \sup_{(\hat{x}_V, \hat{y}) \in Q_r(\bar{x}_V^t, \bar{y}^t)} c \cdot (\bar{x}_V^t - \hat{x}_V) - \bar{y}^t \cdot \hat{x}_V - \hat{y} \cdot (\bar{x}_U^t - \bar{x}_V^t) \right) \\ &= \frac{1}{r} \sup_{(\hat{x}_V, \hat{y}) \in Q_r(\bar{x}_V^t, \bar{y}^t)} -(\bar{y}^t + c) \cdot (\hat{x}_V - \bar{x}_V^t) + (\bar{x}_V^t - \bar{x}_U^t) \cdot (\hat{y} - \bar{y}^t), \quad (55) \end{aligned}$$

where the first equality follows from definition of the normalized duality gap, the second equality uses the definition of M and $W_r(z)$, and the third equality uses that $\bar{y}^t \cdot \hat{x}_U = \bar{y}^t \cdot \bar{x}_U$. To see this claim that $\bar{y}^t \cdot \hat{x}_U = \bar{y}^t \cdot \bar{x}_U$ note that $\bar{x}_U^t \in X_U = \{x : Ax = b\}$ and that if $\bar{y}^t \cdot \hat{x}_U \neq \bar{y}^t \cdot \bar{x}_U^t$ for $\hat{x}_U \in X_U$ then $\bar{x}_U^t + \alpha(\hat{x}_U - \bar{x}_U^t) \in X_U$ for all $\alpha \in \mathbb{R}$ and either $\lim_{\alpha \rightarrow \infty} \bar{y}^t \cdot (\bar{x}_U^t + \alpha(\hat{x}_U - \bar{x}_U^t)) = \infty$ or $\lim_{\alpha \rightarrow -\infty} \bar{y}^t \cdot (\bar{x}_U^t + \alpha(\hat{x}_U - \bar{x}_U^t)) = \infty$, this contradicts that $\rho_r(\bar{z}^t) < \infty$ (Proposition 4).

Equation (55) is again a trust region problem with a linear objective, Euclidean norm, and bounds on the variables, so can be solved in linear time using the algorithm from Sect. 6.3.1.

7 Numerical experiments

As validation of the linear convergence rates predicted by Theorems 1 and 2, we implemented and tested the proposed restart schemes for PDHG, EGM and ADMM applied to LP. Indeed, for our test problems we are able to achieve very high accuracy solutions, commensurate with the accuracy expected from second-order methods on LP problems. The experiments also show that our adaptive scheme is competitive with the best fixed frequency restart scheme without any need for hyperparameter search. This section primarily covers the results for PDHG. Additional details and results for EGM and ADMM are presented in ‘‘Appendix G’’. The code for the numerical experiments is publicly available at https://github.com/google-research/google-research/tree/master/restarting_FOM_for_LP.

We select four LP instances, `qap10`, `qap15`, `nug08-3rd`, and `nug20`, from the Mittelmann collection set [55], a standard benchmark set for LP. These four problems are relaxations of quadratic assignment problems [70], a classical NP-hard combinatorial optimization problem. These instances are known to be challenging for traditional solvers and amenable to first-order methods [28]. The instance sizes are given in Table 3.

For the step size of PDHG we use $\eta = 0.9\sigma_{\max}(A)$ which ensures the theoretical bounds for PDHG hold, where $\sigma_{\max}(A)$ is estimated using power iteration.

Table 3 Test problem sizes

	Rows	Cols	Nonzeros
qap10	1820	4150	18,200
qap15	6330	22,275	94,950
nug8-3rd	19,728	20,448	139,008
nug20	15,240	72,600	304,800

It is known that a major factor in the practical performance of PDHG is the relative scaling of the primal and dual [14, 31]. We call this the primal weight (ω). It can be viewed as a scaling factor on the objective value, i.e., setting $c_{\text{new}} = c/\omega^2$, or as the relative step size for the primal and dual, i.e., the primal step size is η/ω and the dual step size is $\eta\omega$. Before running PDHG with restarts on each of these problems we estimate the optimal primal weight by running the non-restarted PDHG algorithm for 5000 iterations across the range $\omega \in \{4^{-5}, 4^{-4}, 4^{-3}, 4^{-2}, 4^{-1}, 1, 4^1, 4^2, 4^3, 4^4, 4^5\}$. We then choose the ω value with the smallest value of the KKT error $\|(Kz - h)^+\|_2$ where K and h are defined as per (20). We use this ω value for our experiment across all restart schemes.

On these problems we run PDHG with: no restarts, fixed frequency restarts with different restart lengths, and adaptive restarts (where the norm used for computing the normalized duality gap is Euclidean as described in Sect. 6.3.2). Fixed restart lengths of $4^1, 4^2, \dots, 4^9$ were tested and the best three restart lengths⁹ were plotted for each problem. In Fig. 2, we plot the value of the normalized duality gap against the number of iterations for each of the method where the input radius for the normalized duality gap is equal to the distance traveled since the last restart (for no restarts this is the starting point). For the adaptive restarts we use $\beta = \exp(-1)$ as per Remark 5 and only evaluate (30) every 30 iterations to ensure the cost of evaluating the normalized duality gap is negligible. There are a few consistent patterns. One is that often a few iterations after a restart, the normalized duality gap spikes and then falls quickly below its previous value. This is particularly apparent for nug20. Another observation is that after the active set is identified, the linear convergence is often apparent in the plot. However, for qap10, qap15, and nug20, restarted PDHG outperforms non-restarted PDHG, much earlier than when the active set freezes.

The normalized duality gap is not a typical metric for determining solution quality. Typically, the KKT error is a more standard [65] and algorithm-independent metric. Indeed as Lemma 4 shows, the KKT error can be bounded above by the normalized duality gap times a constant.

For completeness, Table 4 reports the number of iterations required by each approach to drive the KKT error below 10^{-6} (either for the average or last iterate). This table exhibits very similar patterns to Fig. 2. In particular, non-restarted PDHG fails to drive the KKT error below 10^{-6} except for the nug08-3rd problem. In Table 4

⁹ The restart lengths were ordered descending by iterations to find a normalized duality gap below 10^{-7} , and then by the normalized duality gap at the maximum number of iterations. The top three restart lengths were then selected for display.

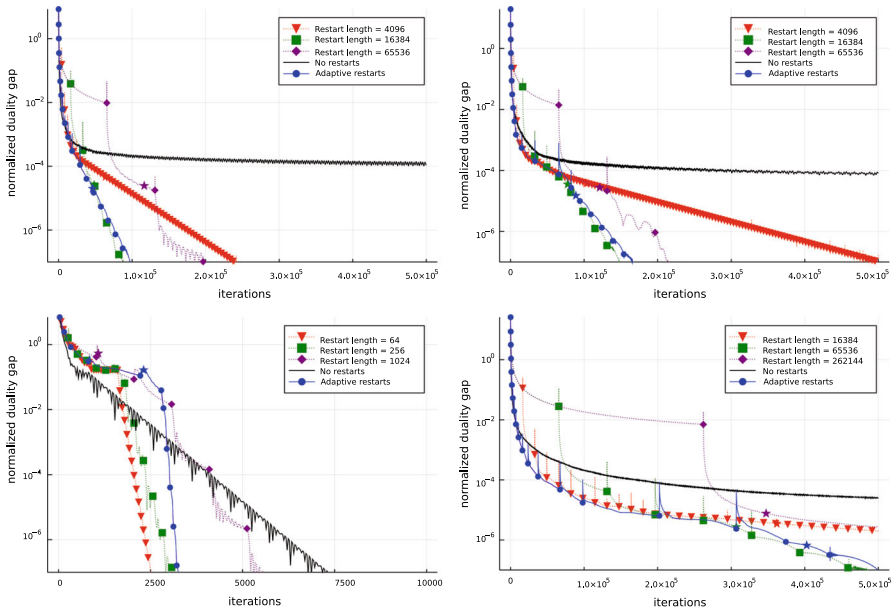


Fig. 2 Plots show normalized duality gap (in log scale) versus number of iterations for restarted PDHG. Results from left to right, top to bottom for qap10, qap15, nug08-3rd, and nug20. Each plot compares no restarts, adaptive restarts, and the three best fixed restart lengths. For the restart schemes we evaluate the normalized duality gap at the average, for no restarts we evaluate it at the last iterate as the average performs much worse. A star indicates the last iteration that the active set changed before the iteration limit was reached. Series markers indicate when restarts occur

we compare the adaptive restarts with the best fixed frequency restart scheme for PDHG, EGM and ADMM. Notably, the number of iterations is fairly similar for both approaches. However, adaptive restarts do not require any hyperparameter search over the restart lengths and therefore require much less computational effort.

Remark 6 (Flexible restarts) A major benefit of using the normalized duality gap as a potential function is that provides substantial flexibility in algorithm design while preserving the convergence properties given in Theorem 2. For example, we can develop an algorithm that chooses either to restart to the average or last iterate. To do this we add an additional line below Line 7 of Algorithm 1 as follows

$$\bar{z}^{n,t} \leftarrow \begin{cases} z^{n,t} & \rho_{\|z^{n,t}-z^{n,0}\|}(z^{n,t}) < \rho_{\|\bar{z}^{n,t}-z^{n,0}\|}(\bar{z}^{n,t}) \\ \bar{z}^{n,t} & \text{otherwise.} \end{cases} \tag{56}$$

In other words, we consider the last iterate as a candidate for restarting, if it has a better normalized duality gap. Inspection of Theorem 2 shows that it still holds with modification to (56). This is because Property ii. also holds for $z^{n,t}$ (Proposition 1 and i. establishes $\|z^{t+1} - z^*\| \leq \|z^t - z^*\|$, then use Proposition 7 to change the norm from PDHG to Euclidean). Indeed one can imagine many other ways of producing restart candidates that satisfy Property ii. (e.g., convex combinations of iterates) and

Table 4 Number of iterations required for the KKT error (the maximal absolute primal residual, dual residual and primal dual gap) to fall below 10^{-6} . The stopping condition is checked every 30 iterations. The – symbol indicates the iteration limit (500,000 for PDHG and EGM; 200,000 for ADMM as its iterations are slower) was reached. The best fixed frequency restart scheme is the fixed frequency restart scheme from Fig. 2 that requires the least iterations to achieve 10^{-6} KKT error. Flexible restarts are described in Remark 6

Problem name	Method	No restarts	Best fixed frequency	Adaptive restarts	Flexible restarts
qap10	PDHG	–	76,230	86,820	84,240
	EGM	–	76,170	86,820	84,240
	ADMM	–	19,590	26,340	21,180
qap15	PDHG	–	144,060	153,780	154,560
	EGM	–	144,060	153,780	154,590
	ADMM	–	46,620	44,880	36,660
nug08-3rd	PDHG	6600	2280	3300	1590
	EGM	6300	2520	3180	1590
	ADMM	930	2520	2700	2520
nug20	PDHG	–	399,840	447,300	425,610
	EGM	–	399,840	445,590	426,630
	ADMM	–	148,110	124,380	132,360

therefore maintain the guarantees of Theorem 2. We implement (56) and report the results Table 4 under the name ‘flexible restarts’. One can see that the flexible restart scheme is in general slightly faster than the original adaptive restart scheme; for nug08-3rd it is much faster.

Acknowledgements We thank Brendan O’Donoghue, Vasilis Charisopoulos, and Warren Schudy for useful discussions and feedback on this work.

A Proof of Lemma 2

Proof Let $U\Sigma V^T$ be a singular value decomposition of H . Recall U and V are orthogonal matrices ($U^T U = U U^T = \mathbf{I}$ and $\|Uz\| = \|z\|$), and Σ is diagonal. First, by assumption there exists some \bar{z} such that $H\bar{z} = h$. This implies that for $\bar{p} = V^T \bar{z}$ we have $U^T h = U^T (U \Sigma V^T) \bar{z} = \Sigma V^T \bar{z} = \Sigma \bar{p}$. Consider some $z \in \mathbb{R}^n$. Define $p = V^T z$,

$$p_i^* = \begin{cases} \Sigma_{ii}^{-1} (U^T h)_i & \Sigma_{ii} \neq 0 \\ z_i & \text{otherwise} \end{cases}$$

and $z^* = Vp^*$. If $\Sigma_{ii} \neq 0$ then by definition of p^* , $(\Sigma p^*)_i = (U^\top h)_i$. If $\Sigma_{ii} = 0$ then $(U^\top h)_i = (\Sigma \bar{p})_i = \Sigma_{ii} \bar{p}_i = 0 = \Sigma_{ii} p_i^* = (\Sigma p^*)_i$. Hence $\Sigma p^* = U^\top h$. Therefore,

$$\begin{aligned} \|z^* - z\| &= \|V^\top z^* - V^\top z\| \\ &= \|p^* - p\| \\ &\leq \frac{1}{\sigma_{\min}^+(H)} \|\Sigma p^* - \Sigma p\| \\ &= \frac{1}{\sigma_{\min}^+(H)} \|(U^\top h) - \Sigma p\| \\ &= \frac{1}{\sigma_{\min}^+(H)} \|h - U\Sigma V^\top z\| \\ &= \frac{1}{\sigma_{\min}^+(H)} \|h - Hz\|. \end{aligned}$$

□

B Proof of Proposition 3 for EGM

Proof Let $z^* = \min_{z \in Z^*} \|z^t - z\|$. By definition of \hat{z}^t and that F is L -Lipschitz, we have

$$\begin{aligned} &F(z^{t-1})^\top (\hat{z}^t - z^{t-1}) + \frac{1}{2\eta} \|\hat{z}^t - z^{t-1}\|^2 \\ &\leq F(z^{t-1})^\top (z^* - z^{t-1}) + \frac{1}{2\eta} \|z^* - z^{t-1}\|^2 \\ &= F(z^*)^\top (z^* - z^{t-1}) + (F(z^{t-1}) - F(z^*))^\top (z^* - z^{t-1}) + \frac{1}{2\eta} \|z^* - z^{t-1}\|^2. \end{aligned} \tag{57}$$

Note that

$$\begin{aligned} F(z^{t-1})^\top (\hat{z}^t - z^{t-1}) &= F(z^*)^\top (\hat{z}^t - z^{t-1}) + (F(z^{t-1}) - F(z^*))^\top (\hat{z}^t - z^*) + \\ &\quad (F(z^{t-1}) - F(z^*))^\top (z^* - z^{t-1}). \end{aligned} \tag{58}$$

Substituting (58) into the LHS of (57) and cancelling terms yields

$$\begin{aligned} &F(z^*)^\top (\hat{z}^t - z^{t-1}) + (F(z^{t-1}) - F(z^*))^\top (\hat{z}^t - z^*) + \frac{1}{2\eta} \|\hat{z}^t - z^{t-1}\|^2 \\ &\leq F(z^*)^\top (z^* - z^{t-1}) + \frac{1}{2\eta} \|z^* - z^{t-1}\|^2. \end{aligned} \tag{59}$$

By $F(z^*)^\top(z^* - z^{t-1}) \leq F(z^*)^\top(\hat{z}^t - z^{t-1})$, (59), F is L -Lipschitz, and the triangle inequality one gets

$$\begin{aligned}
 0 &\geq F(z^*)^\top(z^* - z^{t-1}) - F(z^*)^\top(\hat{z}^t - z^{t-1}) \\
 &\geq \frac{1}{2\eta} \|\hat{z}^t - z^{t-1}\|^2 + (F(z^{t-1}) - F(z^*))^\top(\hat{z}^t - z^*) - \frac{1}{2\eta} \|z^* - z^{t-1}\|^2 \\
 &\geq \frac{1}{2\eta} \|\hat{z}^t - z^{t-1}\|^2 - L\|z^{t-1} - z^*\| \|\hat{z}^t - z^*\| - \frac{1}{2\eta} \|z^* - z^{t-1}\|^2 \\
 &= \frac{1}{2\eta} \|\hat{z}^t - z^{t-1}\|^2 - L\|z^{t-1} - z^*\| \|z^{t-1} - z^* + \hat{z}^t - z^{t-1}\| - \frac{1}{2\eta} \|z^* - z^{t-1}\|^2 \\
 &\geq \frac{1}{2\eta} \|\hat{z}^t - z^{t-1}\|^2 - L\|z^{t-1} - z^*\| \|\hat{z}^t - z^{t-1}\| - \left(L + \frac{1}{2\eta}\right) \|z^* - z^{t-1}\|^2 \\
 &= \left(\frac{1}{2\eta} \|\hat{z}^t - z^{t-1}\| - \left(\frac{1}{2\eta} + L\right) \|z^* - z^{t-1}\|\right) \left(\|\hat{z}^t - z^{t-1}\| + \|z^* - z^{t-1}\|\right)
 \end{aligned} \tag{60}$$

where the first inequality uses $F(z^*)^\top(z^* - z^{t-1}) \leq F(z^*)^\top(\hat{z}^t - z^{t-1})$, the second inequality uses (59), the third inequality uses that F is L -Lipschitz, and the final inequality uses the triangle inequality.

Notice that $\|\hat{z}^t - z^{t-1}\| + \|z^* - z^{t-1}\| > 0$ otherwise the result trivially holds, thus by (60) we get $\frac{1}{2\eta} \|\hat{z}^t - z^{t-1}\| - \left(\frac{1}{2\eta} + L\right) \|z^* - z^{t-1}\| \leq 0$, rearranging yields

$$\|\hat{z}^t - z^{t-1}\| \leq (1 + 2\eta L) \|z^* - z^{t-1}\| \leq 3 \|z^* - z^{t-1}\|$$

where the last inequality uses that $\eta \leq 1/L$ from the premise of Proposition 1. □

C Proof of Lemma 6

Proof Recall for ADMM the norm is $\|z\| = \sqrt{z^\top M z}$ with

$$M = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \eta V^\top V & 0 \\ 0 & 0 & \frac{1}{\eta} \mathbf{I} \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \eta \mathbf{I} & 0 \\ 0 & 0 & \frac{1}{\eta} \mathbf{I} \end{pmatrix}$$

where we utilize $V = I$ in (28). Suppose that $\eta = 1$. Consider $z \in Z, r \in (0, 2R]$. Let

$$v = \begin{pmatrix} 0 \\ (c + y)^- \\ x_V - x_U \end{pmatrix} \text{ and note that } F(z) = \begin{pmatrix} -y \\ c + y \\ x_U - x_V \end{pmatrix}.$$

Let $z_1 = z + r \frac{v}{\|v\|}$, then $z_1 \in Z$ by the definition of Z . It follows from (21) that

$$\rho_r(z) \geq \frac{F(z)^\top(z - z_1)}{r} = -\frac{1}{\|v\|} F(z)^\top v = \|v\|_2 . \tag{61}$$

Consider any optimal solution $z^* = (x_V^*, x_V^*, y^*)$, define $\hat{z} = (x_V^*, x_V^*, y) \in Z^*$. Let $z_2 = z - \mu(z - \hat{z})$ for $\mu = \min \left\{ \frac{r}{\|z - \hat{z}\|}, 1 \right\}$, then $\|z_2 - z\| \leq \frac{r}{\|z - \hat{z}\|} \|z - \hat{z}\| \leq r$. Meanwhile, we have $z_2 = (1 - \mu)z + \mu\hat{z}$, thus $z_2 \in Z$ by convexity. We conclude that $z_2 \in W_r(z)$. Therefore, it follows from (21) that

$$\begin{aligned} \rho_r(z) &\geq \frac{1}{r} \min \left\{ \frac{r}{\|z - \hat{z}\|}, 1 \right\} F(z)^\top(z - \hat{z}) \\ &= \min \left\{ \frac{1}{\|z - \hat{z}\|}, \frac{1}{r} \right\} \begin{pmatrix} -y \\ c + y \\ x_U - x_V \end{pmatrix} \cdot \begin{pmatrix} x_U - x_V^* \\ x_V - x_V^* \\ y \end{pmatrix} \\ &= \min \left\{ \frac{1}{\|z - \hat{z}\|}, \frac{1}{r} \right\} c^\top(x_V - x_V^*) . \end{aligned} \tag{62}$$

Substituting $r \in (0, 2R]$ and $\|z - \hat{z}\| \leq \|z - z^*\| \in [0, 2R]$ into (62) and noticing $\rho_r(z) \geq 0$ yields

$$\rho_r(z) \geq \frac{1}{2R} (c^\top x_V - c^\top x_V^*)^+ . \tag{63}$$

Combining (61) and (63), we deduce there exists K' and h' such that

$$\begin{aligned} (4R^2 + 1)\rho_r(z)^2 &\geq ((c^\top x_V - c^\top x_V^*)^+)^2 + \|(c + y)^-\|_2^2 + \|x_V - x_U\|_2^2 \\ &= ((c^\top x_V - c^\top x_V^*)^+)^2 + \|(c + y)^-\|_2^2 \\ &\quad + \|x_V - x_U\|_2^2 + \|x_U^-\|_2^2 + \|Ax_V - b\|_2^2 \\ &= \|(h' - K'z)^+\|_2^2 \geq \frac{1}{H(K')^2} \mathbf{dist}(z, Z^*)^2 , \end{aligned}$$

where the last inequality is from duality, Hoffman condition (Equation (26)), and the fact that $\|z\|_2 \geq \|z\|$. Taking the square root obtains the result for $\eta = 1$.

Next, we consider the case $\eta \neq 1$. Let us denote the corresponding normalized duality gap as ρ_r^η and distance as \mathbf{dist}^η then with $\theta = \max\{\eta, 1/\eta\}$ we get

$$\begin{aligned} \rho_r^\eta(z) &\geq \frac{\rho_{\theta r}(z)}{\theta} \geq \frac{1}{\theta H(K')\sqrt{4R^2 + 1}} \mathbf{dist}^1(z, Z^*) \\ &\geq \frac{1}{\theta^2 H(K')\sqrt{4R^2 + 1}} \mathbf{dist}^\eta(z, Z^*) . \end{aligned}$$

□

D Proof of Proposition 10

The next lemma is used in the proof of Proposition 10.

Lemma 7 Consider the sequence $\{z^{n,0}\}_{i=0}^\infty, \{\tau^n\}_{i=1}^\infty$ generated by Algorithm 1 with the adaptive restart scheme and $\beta < \frac{1}{q+3}$. Suppose (6) satisfies Property 3, and there exists $a \in (0, \infty)$ s.t. $\rho_r(z) \geq \frac{a}{1+\|z\|} \mathbf{dist}(z, Z^*)$. Then $z^{n,0}$ stays in a bounded region, in particular, $z^{n,0} \in W_R(z^{0,0})$ for all $n \in \mathbb{N}$ with $R = \frac{2C(q+3)^3}{\beta a(1-\beta(q+3))} \mathbf{dist}(z^{0,0}, Z^*)(1 + \|z^{0,0}\| + \mathbf{dist}(z^{0,0}, Z^*))$.

Proof First, notice that it holds for any $n > 0$ that $\mathbf{dist}(z^{n,0}, Z^*) \leq \mathbf{dist}(z^{n-1,0}, Z^*) + \|z^{n,0} - z^{n-1,0}\| \leq (q + 3)\mathbf{dist}(z^{n-1,0}, Z^*)$ where the first inequality uses the triangle inequality, and the second inequality is from Property ii.. Recursively applying this inequality yields

$$\mathbf{dist}(z^{n,0}, Z^*) \leq (q + 3)^n \mathbf{dist}(z^{0,0}, Z^*) . \tag{64}$$

Thus, we have

$$\begin{aligned} \|z^{N,0} - z^{0,0}\| &\leq \sum_{n=0}^{N-1} \|z^{n+1,0} - z^{n,0}\| \leq (q + 3) \sum_{n=0}^{N-1} \mathbf{dist}(z^{n,0}, Z^*) \\ &\leq \mathbf{dist}(z^{0,0}, Z^*) \sum_{n=0}^{N-1} (q + 3)^{n+1} \\ &= \mathbf{dist}(z^{0,0}, Z^*) \frac{(q + 3)^{N+1} - (q + 3)}{q + 2} \\ &\leq \frac{(q + 3)^{N+1}}{2} \mathbf{dist}(z^{0,0}, Z^*) , \end{aligned} \tag{65}$$

where the first inequality uses the triangle inequality, the second inequality is from Property ii., the third inequality utilizes (64) for $n = 0, \dots, N - 1$, the equality uses the formula for the sum of a geometric series, and the last inequality uses $q \geq 0$.

Furthermore, we have

$$\begin{aligned} \rho_{\|z^{n,0} - z^{n-1,0}\|}(z^{n,0}) &\leq \beta^{n-1} \rho_{\|z^{1,0} - z^{0,0}\|}(z^{1,0}) \leq \beta^{n-1} 2C \|z^{1,0} - z^{0,0}\| \\ &\leq \beta^{n-1} 2C(q + 2) \mathbf{dist}(z^{0,0}, Z^*) , \end{aligned} \tag{66}$$

where the first inequality recursively uses (30), the second inequality is from Property i. and $\tau^0 \geq 1$, and the last inequality uses Property ii..

Therefore, it holds that

$$\begin{aligned}
 \|z^{n+1,0} - z^{n,0}\| &\leq (q + 2)\mathbf{dist}(z^{n,0}, Z^*) \\
 &\leq (q + 2)\frac{\rho_{\|z^{n,0} - z^{n-1,0}\|}(z^{n,0})}{a} \left(1 + \|z^{n,0}\|\right) \\
 &\leq \beta^{n-1}\frac{2C(q + 2)^2\mathbf{dist}(z^{0,0}, Z^*)}{a} \left(1 + \|z^{n,0}\|\right) \\
 &\leq \beta^{n-1}\frac{2C(q + 2)^2\mathbf{dist}(z^{0,0}, Z^*)}{a} \left(1 + \|z^{0,0}\| + \|z^{n,0} - z^{0,0}\|\right) \\
 &\leq \beta^{n-1}\frac{2C(q + 2)^2\mathbf{dist}(z^{0,0}, Z^*)}{a} \\
 &\quad \left(1 + \|z^{0,0}\| + \frac{(q + 3)^{n+1}}{2}\mathbf{dist}(z^{0,0}, Z^*)\right) \\
 &\leq (\beta(q + 3))^n\frac{C(q + 3)^3}{\beta a}\mathbf{dist}(z^{0,0}, Z^*) \\
 &\quad \left(1 + \|z^{0,0}\| + \mathbf{dist}(z^{0,0}, Z^*)\right),
 \end{aligned}
 \tag{67}$$

where the first inequality uses Property **ii.**, the second inequality is from that $\rho_r(z) \geq \frac{a}{1+\|z\|}\mathbf{dist}(z, Z^*)$, the third inequality uses (66), the fourth inequality uses the triangle inequality, the fifth inequality uses (65), and the last inequality uses $q + 3 \geq 2$.

Therefore, we have

$$\begin{aligned}
 \|z^{N,0} - z^{0,0}\| &\leq \sum_{n=0}^{N-1} \|z^{n+1,0} - z^{n,0}\| \\
 &\leq \sum_{n=0}^{N-1} (\beta(q + 3))^n\frac{C(q + 3)^3}{\beta a}\mathbf{dist}(z^{0,0}, Z^*) \\
 &\quad \left(1 + \|z^{0,0}\| + \mathbf{dist}(z^{0,0}, Z^*)\right) \\
 &\leq \frac{C(q + 3)^3}{\beta a(1 - \beta(q + 3))}\mathbf{dist}(z^{0,0}, Z^*) \left(1 + \|z^{0,0}\| + \mathbf{dist}(z^{0,0}, Z^*)\right),
 \end{aligned}$$

where the first inequality is the triangle inequality, the second inequality uses (67), and the last inequality is from the bound on the sum of a geometric series by noticing $\beta(q + 3) < 1$. This finishes the proof. \square

Proof of Proposition 10 for EGM and ADMM applied to LP. Recall that the Lagrangian form of an LP instance is α -sharp on $S(R) = \{z \in Z : \|z\| \leq R\}$ with $\alpha = \frac{1}{H(K)\sqrt{1+4R^2}}$ (see Lemma 5) and the ADMM form of a LP instance is α -sharp on $S(R) = \{z \in Z : \|z\| \leq R\}$ with $\alpha = \frac{1}{\max\{\eta^2, 1/\eta^2\}H(K')\sqrt{1+4R^2}}$ (see Lemma 6). Thus the sharpness constant α satisfies the condition in Lemma 7 with $a = \frac{1}{2H(K)}$ for

standard form LP (Lemma 5) and $a = \frac{1}{2 \max\{\eta^2, 1/\eta^2\} H(K')}$ for ADMM form LP. We finish the proof by setting $R = \frac{2C(q+3)^3}{\beta\alpha(1-\beta(q+3))} \mathbf{dist}(z^{0,0}, Z^*)(1 + \|z^{0,0}\| + \mathbf{dist}(z^{0,0}, Z^*))$ as stated in Lemma 7. \square

E Proof of Corollary 2

Proof Let $\|z\|_M = \sqrt{\|x\|^2 - 2\eta y^\top Ax + \|y\|^2}$. Recall that Property 3 holds for $\|\cdot\|_M$ with $q = 0$ and $C = 2/\eta$ (refer to Table 1). Define $\hat{B}_r(z) = \{z \in Z : \|z\|_M \leq r\}$ then with $\hat{r} = r\sqrt{\frac{1}{1-\eta\sigma_{\max}(A)}}$ for any $z \in Z$ we get

$$\begin{aligned} \rho_r(z) &= \frac{\max_{\hat{z} \in W_r(z)} \{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)\}}{r} \\ &\leq \frac{\max_{\hat{z} \in \hat{B}_r(z)} \{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)\}}{r} \\ &= \sqrt{\frac{1}{1-\eta\sigma_{\max}(A)}} \frac{\max_{\hat{z} \in \hat{B}_r(z)} \{\mathcal{L}(x, \hat{y}) - \mathcal{L}(\hat{x}, y)\}}{\hat{r}} \\ &\leq \frac{1}{\eta} \sqrt{\frac{1}{1-\eta\sigma_{\max}(A)}} \frac{2\hat{r}}{t} = \frac{1}{\eta(1-\eta\sigma_{\max}(A))} \frac{2r}{t} \end{aligned}$$

where the first inequality uses that $W_r(z) \subseteq \hat{B}_r(z)$ by Proposition 7, the second inequality uses Property i.. This establishes $C = \frac{2}{\eta(1-\eta\sigma_{\max}(A))}$. Moreover, with $z^* \in \operatorname{argmin}_{z \in Z^*} \|z^0 - z\|_M$ we have

$$\|\bar{z}^t - z^0\|^2 \leq \frac{\|\bar{z}^t - z^0\|_M^2}{1-\eta\sigma_{\max}(A)} \leq \frac{4\|z^* - z^0\|_M^2}{1-\eta\sigma_{\max}(A)} \leq \frac{4(1+\eta\sigma_{\max}(A))}{1-\eta\sigma_{\max}(A)} \|z^* - z^0\|^2$$

where the first inequality uses by Proposition 7, the second inequality uses Property ii., and the third inequality uses Proposition 7. This establishes $q = 4\frac{1+\eta\sigma_{\max}(A)}{1-\eta\sigma_{\max}(A)} - 2$. \square

F Linear time algorithm for linear trust region problems

Theorem 4 Algorithm 2 exactly solves (49), the trust region problem with linear objective, Euclidean norm, and bounds.

Proof If the algorithm returns from line 1, (51) is unbounded, and the algorithm returns $\hat{z}(\infty)$. Otherwise, in each iteration of the while loop (line 5), the algorithm maintains the invariants

- $\mathcal{I} = \{i : \lambda_{lo} < \hat{\lambda}_i < \lambda_{hi}\}$
- For $\lambda_{lo} < \lambda < \lambda_{hi}$, $\|\hat{z}(\lambda) - z\|^2 = f_{lo} + \lambda^2 f_{hi} + \sum_{i \in \mathcal{I}} (\hat{z}(\lambda)_i - z_i)^2$

Algorithm 2: Linear objective trust region algorithm

```

Input: Center  $z \in \mathbb{R}^{n+m}$ , lower bound  $l \in (\mathbb{R} \cup \{-\infty\})^{n+m}$ , objective  $g \in \mathbb{R}^{n+m}$ , radius  $r \geq 0$ 
Assumptions:  $l \leq z, g > 0, r \geq 0$ 
Result: Trust region solution  $\hat{z}$  solving (49)
1 if  $\|l - z\| \leq r$  then return  $\hat{z}(\infty) = l, \lambda_{lo} \leftarrow 0, \lambda_{hi} \leftarrow \infty$ 
2  $f_{lo} \leftarrow \sum_{i: \hat{\lambda}_i \leq \lambda_{lo}} (l_i - z_i)^2, f_{hi} \leftarrow \sum_{i: \hat{\lambda}_i \geq \lambda_{hi}} g_i^2$ 
3  $\mathcal{I} \leftarrow \{i : \lambda_{lo} < \hat{\lambda}_i < \lambda_{hi}\}$ 
4 while  $\mathcal{I} \neq \emptyset$  do
    Invariants:  $\mathcal{I} = \{i : \lambda_{lo} < \hat{\lambda}_i < \lambda_{hi}\}$ 
    For  $\lambda_{lo} < \lambda < \lambda_{hi}, \|\hat{z}(\lambda) - z\|^2 = f_{lo} + \lambda^2 f_{hi} + \sum_{i \in \mathcal{I}} (\hat{z}(\lambda)_i - z_i)^2$ 
     $\|\hat{z}(\lambda_{lo}) - z\|^2 \leq r^2 \leq \|\hat{z}(\lambda_{hi}) - z\|^2$ 
5  $\lambda_{mid} \leftarrow \text{median}(\hat{\lambda}_i : i \in \mathcal{I})$ 
6  $f_{mid} \leftarrow f_{lo} + \sum_{i \in \mathcal{I}} (\hat{z}(\lambda_{mid})_i - z_i)^2 + f_{hi} \lambda_{mid}^2$  //  $f_{mid} = \|\hat{z}(\lambda_{mid}) - z\|^2$ 
7 if  $f_{mid} < r^2$  then
8    $\lambda_{lo} \leftarrow \lambda_{mid}$ 
9    $f_{lo} \leftarrow f_{lo} + \sum_{i \in \mathcal{I}: \hat{\lambda}_i \leq \lambda_{mid}} (l_i - z_i)^2$ 
10   $\mathcal{I} \leftarrow \{i \in \mathcal{I} : \hat{\lambda}_i > \lambda_{mid}\}$ 
11 else
12   $\lambda_{hi} \leftarrow \lambda_{mid}$ 
13   $f_{hi} \leftarrow f_{hi} + \sum_{i \in \mathcal{I}: \hat{\lambda}_i \geq \lambda_{mid}} g_i^2$ 
14   $\mathcal{I} \leftarrow \{i \in \mathcal{I} : \hat{\lambda}_i < \lambda_{mid}\}$ 
15 end
16 end
17  $\lambda_{mid} \leftarrow \sqrt{(r^2 - f_{lo}) / f_{hi}}$  // solving  $r^2 = \|\hat{z}(\lambda_{mid}) - z\|^2 = f_{lo} + \lambda^2 f_{hi}$ 
18 return  $\hat{z}(\lambda_{mid})$ 

```

- $\|\hat{z}(\lambda_{lo}) - z\|^2 \leq r^2 \leq \|\hat{z}(\lambda_{hi}) - z\|^2$

The while loop (line 5) is finite, since on each iteration $|\mathcal{I}|$ is reduced by at least a factor of 2. When the while loop exits (with $\mathcal{I} = \emptyset$), these invariants mean that the final λ_{mid} computed on line 18 optimizes (51) and the returned $\hat{z}(\lambda_{mid})$ solves (49). □

Theorem 5 Algorithm 2 runs in $O(m + n)$ time.

Proof The work outside the while loop (line 5) is clearly $O(m + n)$. In each pass through the while loop, the median (line 6) can be found in $O(|\mathcal{I}|)$ time ([9]), and the rest of the loop also takes $O(|\mathcal{I}|)$ time. Since at least half of the elements of \mathcal{I} are removed in each iteration, and initially $|\mathcal{I}| \leq m + n$, the total time in the while loop is at most $O(\sum_{i=0}^{\infty} \frac{m+n}{2^i}) = O(m + n)$ by the formula for the sum of an infinite geometric series. □

G Numerical results for restarted EGM and ADMM

In this section, we repeat the numerical experiments of Sect. 7 with EGM and ADMM instead of PDHG. We tune the primal weight for EGM and the step size η for ADMM using the same procedure as with PDHG (see Sect. 7).

Figure 3 (EGM) is almost identical to Fig. 2 (PDHG), which verifies again the improved performance of restarted algorithms. Furthermore, note that one iteration of

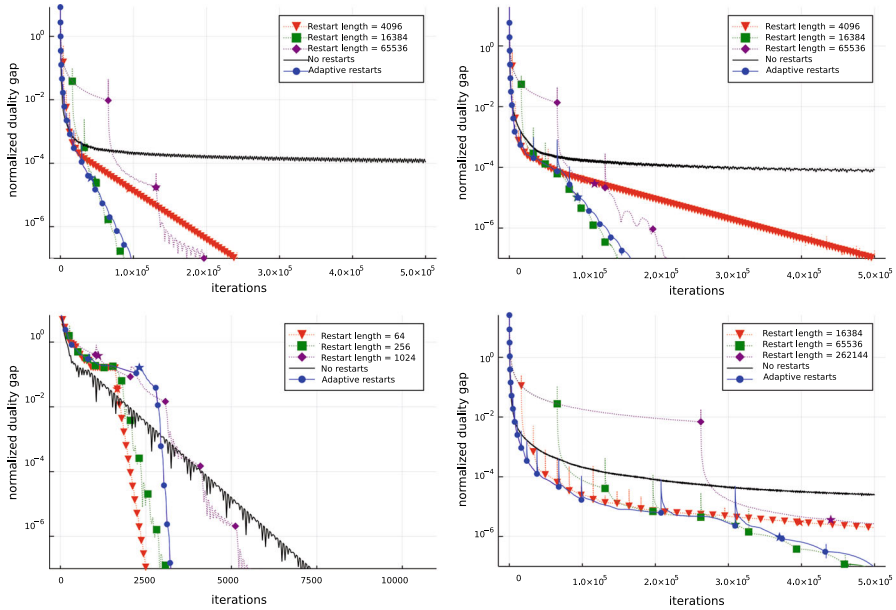


Fig. 3 Plots show normalized duality gap (in log scale) versus number of iterations for (restarted) EGM. Results from left to right, top to bottom for qap10, qap15, nug08-3rd, and nug20. Each plot compares no restarts, adaptive restarts, and the three best fixed restart lengths. For the restart schemes we evaluate the normalized duality gap at the average, for no restarts we evaluate it at the last iterate as the average performs much worse. A star indicates the last iteration that the active set changed before the iteration limit was reached. Series markers indicate when restarts occur

EGM is about twice as expensive as a PDHG iteration, because EGM requires four matrix–vector multiplications per iteration, while PDHG requires two matrix–vector multiplications.

Indeed, to see the difference between EGM and PDHG requires a close examination of Table 4. Recall that termination criteria is only checked every 30 iterations, which causes the iteration counts to be identical in several instances.

Figure 4 shows the performance of restarted ADMM for the same problems, which verifies again the improved performance of restarted algorithms. We observe that restarts do appear to make performance slower for nug08-3rd but this appears to be a very easy problem—the number of iterations for all methods is very low.

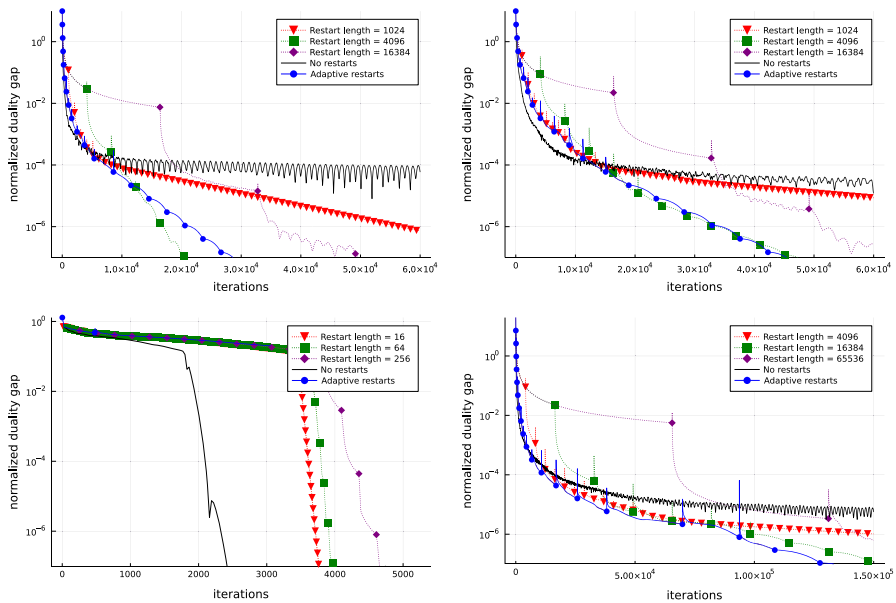


Fig. 4 Plots show normalized duality gap (in log scale) versus number of iterations for (restarted) ADMM. Results from left to right, top to bottom for qap10, qap15, nug08-3rd, and nug20. Each plot compares no restarts, adaptive restarts, and the three best fixed restart lengths. For the restart schemes we evaluate the normalized duality gap at the average, for no restarts we evaluate it at the last iterate as the average performs much worse. A star indicates the last iteration that the active set changed before the iteration limit was reached. Series markers indicate when restarts occur

References

- Alacaoglu, A., Fercoq, O., Cevher, V.: On the convergence of stochastic primal-dual hybrid gradient, arXiv preprint [arXiv:1911.00799](https://arxiv.org/abs/1911.00799) (2019)
- Alamo, T., Limon, D., Krupa, P.: Restart FISTA with global linear convergence. In: 18th European Control Conference (ECC). IEEE, vol. 2019, pp. 1969–1974 (2019)
- Andersen, E.D., Andersen, K.D.: The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In: High Performance Optimization, pp. 197–232. Springer (2000)
- Anderson, R.I., Fok, R., Scott, J.: Hotel industry efficiency: an advanced linear programming examination. *Am. Bus. Rev.* **18**(1), 40 (2000)
- Applegate, D., Díaz, M., Hinder, O., Lu, H., Lubin, M., O’Donoghue, B., Schudy, W.: Practical large-scale linear programming using primal-dual hybrid gradient. In: Advances in Neural Information Processing Systems, vol. 34 (2021)
- Applegate, D., Díaz, M., Lu, H., Lubin, M.: Infeasibility detection with primal-dual hybrid gradient for large-scale linear programming, arXiv preprint [arXiv:2102.04592](https://arxiv.org/abs/2102.04592) (2021)
- Basu, K., Ghoting, A., Mazumder, R., Pan, Y.: ECLIPSE: an extreme-scale linear program solver for web-applications. In: Daumé III, H., Singh, A. (eds.) Proceedings of the 37th International Conference on Machine Learning (Virtual), Proceedings of Machine Learning Research, PMLR, vol. 119, pp. 704–714 (2020)
- Basu, K., Ghoting, A., Mazumder, R., Pan, Y.: Eclipse: an extreme-scale linear program solver for web-applications. In: International Conference on Machine Learning, PMLR, pp. 704–714 (2020)
- Blum, M., Floyd, R.W., Pratt, V., Rivest, R.L., Tarjan, R.E.: Time bounds for selection. *J. Comput. Syst. Sci.* **7**(4), 448–461 (1973)

10. Bowman, E.H.: Production scheduling by the transportation method of linear programming. *Oper. Res.* **4**(1), 100–103 (1956)
11. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., et al.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends@ Mach. Learn.* **3**(1), 1–122 (2011)
12. Burke, J.V., Ferris, M.C.: Weak sharp minima in mathematical programming. *SIAM J. Control Optim.* **31**(5), 1340–1359 (1993)
13. Burke, V.J., Ferris, M.C.: A Gauss–Newton method for convex composite optimization. *Math. Program.* **71**(2), 179–194 (1995)
14. Chambolle, A., Pock, T.: A first-order primal–dual algorithm for convex problems with applications to imaging. *J. Math. Imaging Vis.* **40**(1), 120–145 (2011)
15. Chambolle, A., Pock, T.: On the ergodic convergence rates of a first-order primal–dual algorithm. *Math. Program.* **159**(1–2), 253–287 (2016)
16. Charnes, A., Cooper, W.W.: The stepping stone method of explaining linear programming calculations in transportation problems. *Manag. Sci.* **1**(1), 49–69 (1954)
17. Condat, L., Malinovsky, G., Richtárik, P.: Distributed proximal splitting algorithms with rates and acceleration. *Front. Signal Process.* **12** (2022)
18. Dantzig, G.B.: *Linear Programming and Extensions*, vol. 48. Princeton University Press (1998)
19. Daskalakis, C., Andrew, I., Syrgkanis, V., Zeng, H.: Training GANs with optimism. In: *International Conference on Learning Representations* (2018)
20. Davis, D., Drusvyatskiy, D., MacPhee, K.J., Paquette, C.: Subgradient methods for sharp weakly convex functions. *J. Optim. Theory Appl.* **179**(3), 962–982 (2018)
21. Douglas, J., Rachford, H.H.: On the numerical solution of heat conduction problems in two and three space variables. *Trans. Am. Math. Soc.* **82**(2), 421–439 (1956)
22. Eckstein, J., Bertsekas, D.P.: On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Math. Program.* **55**(1–3), 293–318 (1992)
23. Eckstein, J., Bertsekas, D.P.: et al., An alternating direction method for linear programming (1990)
24. Fercoq, O.: Quadratic error bound of the smoothed gap and the restarted averaged primal–dual hybrid gradient (2021)
25. Fercoq, O., Zheng, Q.: Adaptive restart of accelerated gradient methods under local quadratic growth condition. *IMA J. Numer. Anal.* **39**(4), 2069–2095 (2019)
26. Ferris, M.C.: Finite termination of the proximal point algorithm. *Math. Program.* **50**(1–3), 359–366 (1991)
27. Freund, R.M., Haihao, L.: New computational guarantees for solving convex optimization problems with first order methods, via a function growth condition measure. *Math. Program.* **170**(2), 445–477 (2018)
28. Galabova, I.L., Hall, J.A.J.: The ‘idiot’ crash quadratic penalty algorithm for linear programming and its application to linearizations of quadratic assignment problems. *Optim. Methods Softw.* **35**(3), 488–501 (2020)
29. Gilpin, A., Pena, J., Sandholm, T.: First-order algorithm with $\mathcal{O}(\ln(1/\epsilon))$ -convergence for ϵ -equilibrium in two-person zero-sum games. *Math. Program.* **133**(1), 279–298 (2012)
30. Giselsson, P., Boyd, S.: Monotonicity and restart in fast gradient methods. In: *53rd IEEE Conference on Decision and Control*, pp. 5058–5063. IEEE (2014)
31. Goldstein, T., Li, M., Yuan, X.: Adaptive primal–dual splitting methods for statistical learning and image processing. In: *Advances in Neural Information Processing Systems*, pp. 2089–2097 (2015)
32. Gondzio, J.: Interior point methods 25 years later. *Eur. J. Oper. Res.* **218**(3), 587–601 (2012)
33. Güler, O., Hoffman, A.J., Rothblum, U.G.: Approximations to solutions to systems of linear inequalities. *SIAM J. Matrix Anal. Appl.* **16**(2), 688–696 (1995)
34. Gutman, D.H., Peña, J.F.: The condition number of a function relative to a set. *Math. Program.* (2020), to appear
35. Hanssmann, F., Hess, S.W.: A linear programming approach to production and employment scheduling. *Manag. Sci.* **1**, 46–51 (1960)
36. Harker, P.T., Pang, J.-S.: Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Math. Program.* **48**(1–3), 161–220 (1990)
37. He, B., Yuan, X.: On the $O(1/n)$ convergence rate of the Douglas–Rachford alternating direction method. *SIAM J. Numer. Anal.* **50**(2), 700–709 (2012)

38. Hoffman, A.J.: On approximate solutions of systems of linear inequalities. *J. Res. Natl. Bur. Stand.* **49**, 263–265 (1952)
39. Hunter, J.K., Natchergale, B.: *Applied Analysis*. World Scientific Publishing Company (2001)
40. Johnson, R., Zhang, T.: Accelerating stochastic gradient descent using predictive variance reduction. *Adv. Neural. Inf. Process. Syst.* **26**, 315–323 (2013)
41. Karmarkar, N.: A new polynomial-time algorithm for linear programming. In: *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, pp. 302–311 (1984)
42. Klatte, D., Thiere, G.: Error bounds for solutions of linear equations and inequalities. *Z. Oper. Res.* **41**(2), 191–214 (1995)
43. Korpelevich, G.M.: The extragradient method for finding saddle points and other problems. *Matecon* **12**, 747–756 (1976)
44. Lewis, A.S., Liang, J.: Partial smoothness and constant rank, arXiv preprint [arXiv:1807.03134](https://arxiv.org/abs/1807.03134) (2018)
45. Li, X., Sun, D., Toh, K.-C.: An asymptotically superlinearly convergent semismooth newton augmented Lagrangian method for linear programming. *SIAM J. Optim.* **30**(3), 2410–2440 (2020)
46. Liang, J., Fadili, J., Peyré, G.: Local linear convergence analysis of primal-dual splitting methods. *Optimization* **67**(6), 821–853 (2018)
47. Lin, H., Mairal, J., Harchaoui, Z.: A universal catalyst for first-order optimization. In: *Advances in Neural Information Processing Systems*, pp. 3384–3392 (2015)
48. Lin, Q., Xiao, L.: An adaptive accelerated proximal gradient method and its homotopy continuation for sparse optimization. In: *International Conference on Machine Learning*, pp. 73–81 (2014)
49. Lin, T., Ma, S., Ye, Y., Zhang, S.: An ADMM-based interior-point method for large-scale linear programming. *Optim. Methods Softw.* **36**(2–3), 389–424 (2021)
50. Liu, Q., Van Ryzin, G.: On the choice-based linear programming model for network revenue management. *Manuf. Serv. Oper. Manag.* **10**(2), 288–310 (2008)
51. Lu, H.: An $O(s^r)$ -resolution ODE framework for discrete-time optimization algorithms and applications to convex-concave saddle-point problems, arXiv preprint [arXiv:2001.08826](https://arxiv.org/abs/2001.08826) (2020)
52. Luo, Z.-Q., Tseng, P.: Error bounds and convergence analysis of feasible descent methods: a general approach. *Ann. Oper. Res.* **46**(1), 157–178 (1993)
53. Manne, A.S.: Linear programming and sequential decisions. *Manag. Sci.* **6**(3), 259–267 (1960)
54. Marcotte, P., Zhu, D.: Weak sharp solutions of variational inequalities. *SIAM J. Optim.* **9**(1), 179–189 (1998)
55. Mittelman, H.D.: Benchmark of simplex LP solvers (2020). <http://plato.asu.edu/ftp/lpsimp.html>
56. Mokhtari, A., Ozdaglar, A., Pattathil, S.: A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: proximal point approach. In: *International Conference on Artificial Intelligence and Statistics* (2020)
57. Nemirovski, A.: 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 J. Optim.* **15**(1), 229–251 (2004)
58. Nesterov, Yu.: Subgradient methods for huge-scale optimization problems. *Math. Program.* **146**(1), 275–297 (2014)
59. Nesterov, Y.: Smooth minimization of non-smooth functions. *Math. Program.* **103**(1), 127–152 (2005)
60. Nesterov, Y.: Gradient methods for minimizing composite functions. *Math. Program.* **140**(1), 125–161 (2013)
61. Nesterov, Y.: *Introductory Lectures on Convex Optimization: a basic course*, vol. 87. Springer (2013)
62. Nesterov, Y.E.: A method for solving the convex programming problem with convergence rate $O(1/k^2)$. *Soviet Math. Doklady* **27**, 372–376 (1983)
63. Niao, H.: Mirror-prox algorithm, Fall (2016), http://niaohe.ise.illinois.edu/IE598_2016/pdf/IE598-lecture18-mirror%20prox%20algorithm%20for%20saddle%20point%20problems.pdf
64. O’Donoghue, B., Candes, E.: Adaptive restart for accelerated gradient schemes. *Found. Comput. Math.* **15**(3), 715–732 (2015)
65. O’Donoghue, B., Chu, E., Parikh, N., Boyd, S.: Conic optimization via operator splitting and homogeneous self-dual embedding. *J. Optim. Theory Appl.* **169**(3), 1042–1068 (2016)
66. Peña, J., Vera, J.C., Zuluaga, L.F.: New characterizations of Hoffman constants for systems of linear constraints. *Math. Program.* (2020), to appear
67. Pokutta, S.: Restarting algorithms: sometimes there is free lunch. In: *International Conference on Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pp. 22–38. Springer (2020)

68. Polyak, B.: Sharp minima. In: Proceedings of the IIASA Workshop on Generalized Lagrangians and Their Applications, Laxenburg, Austria. Institute of Control Sciences Lecture Notes, Moscow (1979)
69. Polyak, B.: Introduction to Optimization. Optimization Software Inc, New York (1987)
70. Ramakrishnan, K.G., Resende, M.G.C., Ramachandran, B., Pekny, J.F.: Tight QAP Bounds via Linear Programming, pp. 297–303. World Scientific Publishing Co. (2002)
71. Renegar, J.: Incorporating condition measures into the complexity theory of linear programming. *SIAM J. Optim.* **5**(3), 506–524 (1995)
72. Renegar, J.: Linear programming, complexity theory and elementary functional analysis. *Math. Program.* **70**(1–3), 279–351 (1995)
73. Roulet, V., d’Aspremont, A.: Sharpness, restart, and acceleration. *SIAM J. Optim.* **30**(1), 262–289 (2020)
74. Tyrrell Rockafellar, R.: Monotone operators and the proximal point algorithm. *SIAM J. Control. Optim.* **14**(5), 877–898 (1976)
75. Tang, J., Golbabaee, M., Bach, F. et al.: Rest-katyusha: exploiting the solution’s structure via scheduled restart schemes. In: Advances in Neural Information Processing Systems, pp. 429–440 (2018)
76. Thomson, B.S., Bruckner, J.B., Bruckner, A.M.: Elementary real analysis, vol. 1, ClassicalRealAnalysis.com (2008)
77. Tseng, P.: On linear convergence of iterative methods for the variational inequality problem. *J. Comput. Appl. Math.* **60**(1–2), 237–252 (1995)
78. Yang, T., Lin, Q.: RSG: beating subgradient method without smoothness and strong convexity. *J. Mach. Learn. Res.* **19**(1), 236–268 (2018)

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

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.