
Parameter-free Locally Accelerated Conditional Gradients

Alejandro Carderera^{*1} Jelena Diakonikolas^{*2} Cheuk Yin (Eric) Lin^{*2} Sebastian Pokutta^{*3}

Abstract

Projection-free conditional gradient (CG) methods are the algorithms of choice for constrained optimization setups in which projections are often computationally prohibitive but linear optimization over the constraint set remains computationally feasible. Unlike in projection-based methods, globally accelerated convergence rates are in general unattainable for CG. However, a very recent work on *Locally accelerated CG* (LaCG) has demonstrated that local acceleration for CG is possible for many settings of interest. The main downside of LaCG is that it requires knowledge of the smoothness and strong convexity parameters of the objective function. We remove this limitation by introducing a novel, Parameter-Free Locally accelerated CG (PF-LaCG) algorithm, for which we provide rigorous convergence guarantees. Our theoretical results are complemented by numerical experiments, which demonstrate local acceleration and showcase the practical improvements of PF-LaCG over non-accelerated algorithms, both in terms of iteration count and wall-clock time.

1. Introduction

Conditional gradient (CG) (or Frank-Wolfe (FW)) methods (Frank & Wolfe, 1956; Levitin & Polyak, 1966) are a fundamental class of projection-free optimization methods, most frequently used to minimize smooth convex objective functions over constrained sets onto which projections are computationally prohibitive; see Combettes & Pokutta (2021) for an overview. These methods have received significant recent attention in the machine learning and optimization communities, due to the fact that they eschew projections and produce solutions with sparse representa-

tions (Jaggi, 2013; Garber, 2016; Hazan & Luo, 2016; Braun et al., 2017; 2019; Lei et al., 2019; Tsiligkaridis & Roberts, 2020; Combettes et al., 2020).

While CG methods have been applied to many different problem settings (see, e.g., Hazan & Luo (2016); Zhou et al. (2018); Pedregosa et al. (2020); Lei et al. (2019); Tsiligkaridis & Roberts (2020); Dvurechensky et al. (2020); Zhang et al. (2020); Négier et al. (2020); Carderera & Pokutta (2020); Kerdreux et al. (2021)), in this paper, we are interested in using CG-type methods to solve problems of the form:

$$\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}), \quad (\text{P})$$

where f is an L -smooth (gradient-Lipschitz) and m -strongly convex function and $\mathcal{X} \subseteq \mathbb{R}^n$ is a polytope.

We assume that we are given access to the objective function f and to the feasible set \mathcal{X} via the following two oracles:

Oracle 1.1 (First Order Oracle (FOO)). Given $\mathbf{x} \in \mathcal{X}$, the FOO returns $f(\mathbf{x})$ and $\nabla f(\mathbf{x})$.

Oracle 1.2 (Linear Minimization Oracle (LMO)). Given $\mathbf{c} \in \mathbb{R}^n$, the LMO returns $\arg\min_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{c}, \mathbf{x} \rangle$.

While being of extreme practical importance, these LMO-based methods in general do not achieve the *globally* optimal rates for smooth (and possibly strongly convex) minimization that are attained by projection-based methods. In particular, LMO-based algorithms cannot converge globally faster than $O(1/k)$ for the class of smooth strongly convex functions where k is the number of iterations (up to the dimension threshold n) (Lan, 2013; Jaggi, 2013). Moreover, the dependence of the convergence rate on the dimension is unavoidable in general.

At the same time, it was shown by Diakonikolas et al. (2020) that optimal rates can be obtained asymptotically, referred to as *locally optimal rates*. That is, after a finite number of iterations (independent of the target accuracy ϵ) with potentially sub-optimal convergence, optimal convergence rates can be achieved. While Diakonikolas et al. (2020) resolve the question of acceleration for CG-type methods, it unfortunately depends on the knowledge of m and L . Although the latter can be removed with the common line search-based arguments, the knowledge of a good estimate of the former is crucial in achieving acceleration. This makes the *Locally-*

^{*}Equal contribution ¹Georgia Institute of Technology, Atlanta, GA, USA ²Department of Computer Sciences, University of Wisconsin-Madison, Madison, WI, USA ³Zuse Institute Berlin and Technische Universität Berlin, Berlin, Germany. Correspondence to: Jelena Diakonikolas <jelena@cs.wisc.edu>.

accelerated Conditional Gradient (LaCG) algorithm from Diakonikolas et al. (2020), despite being of theoretical interest, of limited use in practice. Not only is it usually hard to come by a good estimate of the parameter m , but working with an estimated lower bound does not take advantage of the potentially better local strong convexity behavior in the vicinity of an iterate.

To remedy these shortcomings, we devise a new *Parameter-Free Locally-accelerated Conditional Gradient algorithm (PF-LaCG)* that is based on a similar coupling between a variant of the *Away-step Frank-Wolfe (AFW)* method (Guélat & Marcotte, 1986; Lacoste-Julien & Jaggi, 2015) and an accelerated method as used in LaCG. However, beyond this basic inspiration, not many things can be reused from Diakonikolas et al. (2020), as in order to achieve parameter-freeness, we need to devise a completely new algorithm employing a gradient-mapping technique that out-of-the-box is incompatible with the approach used in LaCG.

1.1. Contributions and Further Related Work

Our main contributions can be summarized as follows (see Section 3.1 for a detailed overview of the main technical ideas).

Near-optimal and parameter-free acceleration with inexact projections. To devise PF-LaCG, we introduce a parameter-free accelerated algorithm for smooth strongly convex optimization that utilizes *inexact* projections onto low-dimensional simplices. While near-optimal (i.e., optimal up to poly-log factors) parameter-free projection-based algorithms were known in the literature prior to our work (Nesterov, 2013; Ito & Fukuda, 2019), their reliance on exact projections which are computationally infeasible makes them unsuitable for our setting.

Parameter-free Locally-accelerated Conditional Gradient (PF-LaCG) algorithm. We propose a novel, parameter-free and locally accelerated CG-type method. Up to poly-logarithmic factors, our algorithm PF-LaCG attains an optimal accelerated local rate of convergence. PF-LaCG leverages efficiently computable projections onto low dimensional simplices, but is otherwise projection-free (i.e., it does not assume access to a projection oracle for \mathcal{X}). Local acceleration is achieved by coupling the parameter-free accelerated method with inexact projections described in the previous paragraph and AFW with a fractional exit condition (Kerdreux et al., 2019). This coupling idea is inspired by the coupling between μ AGD+ (Cohen et al., 2018) (where AGD stands for Accelerated Gradient Descent) and AFW (Guélat & Marcotte, 1986) used in Diakonikolas et al. (2020); however, most of the similarities between our work and Diakonikolas et al. (2020) stop there, as there are major technical challenges that have to be overcome to attain the

results in the parameter-free setting.

Computational experiments. We demonstrate the efficacy of PF-LaCG using numerical experiments, comparing the performance of the proposed algorithms to several relevant CG-type algorithms. The use of PF-LaCG brings demonstrably faster local convergence in primal gap with respect to both iteration count and wall-clock time.

1.2. Outline

In Section 2 we introduce the notation and preliminaries that are required for stating our main results. We then present our approach to parameter-free local acceleration in Section 3 and derive our main results. Finally, we demonstrate the practicality of our approach with computational experiments in Section 4, and conclude with a discussion in Section 5.

2. Notation and Preliminaries

We denote the unique minimizer of Problem (P) by \mathbf{x}^* . Let $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$ denote the Euclidean norm and the standard inner product, respectively. We denote the *diameter* of the polytope \mathcal{X} by $D = \max_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{x} - \mathbf{y}\|$, and its *vertices* by $\text{vert}(\mathcal{X}) \subseteq \mathcal{X}$. Given a non-empty set $\mathcal{S} \subset \mathbb{R}^n$, we denote its *convex hull* by $\text{co}(\mathcal{S})$. For any $\mathbf{x} \in \mathcal{X}$ we denote by $\mathcal{F}(\mathbf{x})$ the *minimal face* of \mathcal{X} that contains \mathbf{x} . We call a subset of vertices $\mathcal{S} \subseteq \text{vert}(\mathcal{X})$ a *support* of $\mathbf{x} \in \mathcal{X}$ if \mathbf{x} can be expressed as a convex combination of the elements of \mathcal{S} . A support \mathcal{S} of \mathbf{x} is a *proper support* of \mathbf{x} if the weights associated with the convex decomposition are positive. Let $\mathcal{B}(\mathbf{x}, r)$ denote the ball around \mathbf{x} with radius r with respect to $\|\cdot\|$. We say that \mathbf{x} is *r -deep* in a convex set $\mathcal{C} \subseteq \mathbb{R}^n$ if $\mathcal{B}(\mathbf{x}, r) \cap \text{aff}(\mathcal{C}) \subseteq \mathcal{C}$, where $\text{aff}(\mathcal{C})$ denotes the smallest affine space that contains \mathcal{C} . The point \mathbf{x} is contained in the *relative interior* of \mathcal{C} , $\mathbf{x} \in \text{rel.int}(\mathcal{C})$, if there exists an $r > 0$ such that \mathbf{x} is r -deep in \mathcal{C} .

Measures of optimality. The two key measures of optimality that we will use in this paper are the *Strong Wolfe Gap* and the *Gradient Mapping*. We define the former as:

Definition 2.1 (Strong Wolfe Gap). Given $\mathbf{x} \in \mathcal{X}$, the *strong Wolfe gap* $w(\mathbf{x})$ of f over \mathcal{X} is defined as

$$w(\mathbf{x}) := \min_{\mathcal{S} \in \mathcal{S}_{\mathbf{x}}} w(\mathbf{x}, \mathcal{S}),$$

where $\mathcal{S}_{\mathbf{x}}$ denotes the set of all proper supports of \mathbf{x} and $w(\mathbf{x}, \mathcal{S}) := \max_{\mathbf{y} \in \mathcal{S}, \mathbf{z} \in \mathcal{X}} \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{z} \rangle$.

Note that any polytope \mathcal{X} satisfies what is known as a $\delta(\mathcal{X})$ -scaling inequality, where $\delta(\mathcal{X})$ is the pyramidal width of the polytope (see, e.g., (Lacoste-Julien & Jaggi, 2015; Beck & Shtern, 2017; Peña & Rodríguez, 2019; Gutman & Peña, 2018)) and where the inequality is defined as:

Definition 2.2 (δ -scaling inequality). There exists $\delta > 0$ such that for all $\mathbf{x} \in \mathcal{X} \setminus \mathbf{x}^*$ we have that Problem (P) satisfies $w(\mathbf{x}) \geq \delta(\mathcal{X}) \langle \nabla f(\mathbf{x}), \mathbf{x} - \mathbf{x}^* \rangle / \|\mathbf{x} - \mathbf{x}^*\|$.

To implement a parameter-free variant of a projection-based accelerated algorithm, we will rely on a second measure of optimality, the *gradient mapping*. Recall that we do *not* assume the availability of projections onto the polytope \mathcal{X} ; instead, our algorithm will only rely on low-complexity projections onto simplices spanned by a small number of vertices of \mathcal{X} (see Diakonikolas et al. (2020) for a more detailed discussion). In the following, given a convex set \mathcal{C} , we denote the projection of $\mathbf{x} \in \mathbb{R}^n$ onto \mathcal{C} by $P_{\mathcal{C}}(\mathbf{x})$.

Definition 2.3 (Gradient mapping). Given a convex set $\mathcal{C} \subseteq \mathbb{R}^n$, a differentiable function $f : \mathcal{C} \rightarrow \mathbb{R}$, and a scalar $\eta > 0$, the gradient mapping of f w.r.t. η, \mathcal{C} is defined by:

$$G_{\eta}(\mathbf{x}) = \eta \left(\mathbf{x} - P_{\mathcal{C}} \left(\mathbf{x} - \frac{1}{\eta} \nabla f(\mathbf{x}) \right) \right).$$

The gradient mapping is a generalization of the gradient to constrained sets: when $\mathcal{C} \equiv \mathbb{R}^n$, $G_{\eta}(\mathbf{x}) = \nabla f(\mathbf{x})$. The norm of the gradient mapping can also be used as a measure of optimality: the gradient mapping at a point \mathbf{x} is zero if and only if \mathbf{x} minimizes f over \mathcal{C} ; more generally, a small gradient mapping norm for a smooth function implies a small optimality gap. See Beck (2017, Chapter 10) and Appendix B for more useful properties.

2.1. Assumptions

We make two key assumptions. The first one, the *strict complementarity assumption* (Assumption 2.4) is key to proving the convergence of the iterates to $\mathcal{F}(\mathbf{x}^*)$, and is a common assumption in the Frank-Wolfe literature (Guélat & Marcotte, 1986; Garber, 2020), which is related to the stability of the solution with respect to noise, and rules out degenerate instances.

Assumption 2.4 (Strict complementarity). We have that $\langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle = 0$ if and only if $\mathbf{x} \in \mathcal{F}(\mathbf{x}^*)$. Or stated equivalently, there exists a $\tau > 0$ such that $\langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \geq \tau$ for all $\mathbf{x} \in \text{vert}(\mathcal{X}) \setminus \mathcal{F}(\mathbf{x}^*)$.

Lastly, to achieve local acceleration, similar to Diakonikolas et al. (2020), we require that the optimal solution \mathbf{x}^* is *sufficiently deep* in the relative interior of a face of \mathcal{X} . We make the following assumption about the problem, which covers all cases of interest.

Assumption 2.5 (Location of \mathbf{x}^*). The optimum satisfies $\mathbf{x}^* \notin \text{vert}(\mathcal{X})$, or conversely, there exists an $r > 0$ such that \mathbf{x}^* is r -deep in a face \mathcal{F} of \mathcal{X} .

Note that the entire polytope \mathcal{X} is an n -dimensional face of itself, and thus Assumption 2.5 allows $\mathbf{x}^* \in \text{rel.int}(\mathcal{X})$.

Note that if $\mathbf{x}^* \in \text{vert}(\mathcal{X})$ and the strict complementarity assumption is satisfied (Assumption 2.4), then the projection-free algorithm that will be presented in later sections will reach \mathbf{x}^* in a finite number of iterations (see Garber (2020)), and so there is no need for acceleration. For computational feasibility, we assume that r is bounded away from zero and much larger than the desired accuracy $\epsilon > 0$.

3. Parameter-Free Local Acceleration

This section provides our main result: a parameter-free locally-accelerated CG method (PF-LaCG). Before delving into the technical details, we first describe the core ideas driving our algorithm and its analysis. Due to space constraints, most of the proofs are omitted and are instead provided in the supplementary material.

3.1. Overview of Main Technical Ideas

A standard idea for achieving acceleration in smooth and strongly convex setups where m is unknown is to use a restart-based strategy, which can be described as follows: run an accelerated method for smooth (non-strongly) convex minimization and restart it every time some measure of optimality is reduced by a constant factor. The measures of optimality used in these strategies are either $f(\mathbf{x}) - f(\mathbf{x}^*)$ (Roulet & d’Aspremont, 2020) or $\|G_{\eta}(\mathbf{x})\|$ (Ito & Fukuda, 2019; Nesterov, 2013).

Neither of these two optimality measures can be applied directly to our setting, as neither can be evaluated: (i) it is rarely the case that $f(\mathbf{x}^*)$ is known, which is needed for evaluating $f(\mathbf{x}) - f(\mathbf{x}^*)$, and we make no such assumption here; and (ii) the gradient mapping norm $\|G_{\eta}(\mathbf{x})\|$ is a valid optimality measure only for the entire feasible set and requires computing projections onto it; and we do not assume availability of a projection operator onto \mathcal{X} .

Even though our algorithm utilizes a restarting-based strategy, the restarts are not used for local acceleration, but for the coupling of a CG method and a projection-based accelerated method. This idea is similar to Diakonikolas et al. (2020); however, there are important technical differences. First, the restarts in Diakonikolas et al. (2020) are scheduled and parameter-based; as such, they cannot be utilized in our parameter-free setting. Our idea is to instead use $w(\mathbf{x}, \mathcal{S})$ as a measure of optimality over the polytope, which is observable naturally in many CG algorithms (see Definition 2.1), where \mathcal{S} , dubbed the *active set* of the CG algorithm, is a proper support of \mathbf{x} . We perform restarts each time $w(\mathbf{x}, \mathcal{S})$ is halved. The idea of using $w(\mathbf{x}, \mathcal{S})$ as a measure of optimality comes from Kerdreux et al. (2019); however in the aforementioned paper $w(\mathbf{x}, \mathcal{S})$ was not used to obtain a locally accelerated algorithm.

The aforementioned idea of coupling an active-set-based

CG method and an accelerated projection-based method can be summarized as follows. Because the objective function is strongly convex, any convergent algorithm (under any global optimality measure) will be reducing the distance between its iterates and the optimal solution \mathbf{x}^* . The role of the CG method is to ensure such convergence without requiring projections onto the feasible set \mathcal{X} . When \mathbf{x}^* is contained in the interior of a face of \mathcal{X} , it can be argued that after a finite burn-in phase whose length is independent of the target accuracy ϵ , every active set of the utilized CG method will contain \mathbf{x}^* in its convex hull. As it is possible to keep the active sets reasonably small (and their size can never be larger than the current iteration count), projections onto the convex hull of an active set can be performed efficiently, using low-complexity projections onto a probability simplex. Thus, after the burn-in phase, we could switch to a projection-based accelerated algorithm that uses the convex hull of the active set as its feasible set and attains an accelerated convergence rate from then onwards.

There are, of course, several technical challenges related to implementing such a coupling strategy. To begin with, there is no computationally feasible approach we are aware of that could allow detecting the end of the burn-in phase. Thus any reasonable locally accelerated algorithm needs to work without this information, i.e., we do not know when \mathbf{x}^* is contained in the convex hull of the active set. In [Diakonikolas et al. \(2020\)](#), this is achieved by using a parameter-based accelerated algorithm that monotonically decreases the objective value, running this accelerated algorithm and a CG method (AFW or the *Pairwise-step Frank-Wolfe* (PFW)) in parallel, and, on restarts, updating the iterate of the accelerated method and the active set to the ones from the coupled CG method whenever the iterate of CG provides a lower function value. Thus, after the burn-in phase, if the feasible set of the accelerated method does not contain \mathbf{x}^* (or its close approximation), CG eventually constructs a point with the lower function value, after which the accelerated algorithm takes over, leading to local acceleration.

We can neither rely on the scheduled restarts nor the accelerated algorithm used in [Diakonikolas et al. \(2020\)](#), as both are parameter-based. Instead, our monotonic progress is w.r.t. $w(\mathbf{x}, \mathcal{S})$ (i.e., upon restarts we pick a point and the active set with the lower value of $w(\mathbf{x}, \mathcal{S})$) and we rely on a parameter-free accelerated method. As mentioned before, even using a parameter-free projection-based acceleration requires new results, as we need to rely on inexact projections (and inexact gradient mappings). Further, for our argument to work, it is required that after a burn-in phase the accelerated method contracts $w(\mathbf{x}, \mathcal{S})$ at an accelerated rate. Although this may seem like a minor point, we note that it is not true in general, as $w(\mathbf{x}, \mathcal{S})$ can be related to other notions of optimality only when the algorithm iterates are contained in $\mathcal{F}(\mathbf{x}^*)$ and the primal gap is sufficiently

small. This is where the strict complementarity assumption (Assumption 2.4) comes into play, as it allows us to show that after a burn-in phase (independent of ϵ), we can upper bound $w(\mathbf{x}, \mathcal{S})$ using $f(\mathbf{x}) - f(\mathbf{x}^*)$ (Theorem 3.4).

3.2. Burn-in Phase

The variant of CG used in our work is the *Away Frank-Wolfe* (AFW) method ([Guélat & Marcotte, 1986](#); [Lacoste-Julien & Jaggi, 2015](#)) with a stopping criterion based on halving the Frank-Wolfe gap ([Kerdreux et al., 2019](#)), shown in Algorithm 1. For completeness, the useful technical results from [Kerdreux et al. \(2019\)](#) utilized in our analysis are provided in Appendix A.

Algorithm 1 Away-Step Frank-Wolfe Algorithm: AFW($\mathbf{x}_0, \mathcal{S}_0$)

```

1:  $k := 0$ 
2: while  $w(\mathbf{x}_k, \mathcal{S}_k) > w(\mathbf{x}_0, \mathcal{S}_0)/2$  do
3:    $\mathbf{v}_k := \operatorname{argmin}_{\mathbf{u} \in \mathcal{X}} \langle \nabla f(\mathbf{x}_k), \mathbf{u} \rangle$ ,  $\mathbf{d}_k^{\text{FW}} := \mathbf{v}_k - \mathbf{x}_k$ 
4:    $\mathbf{s}_k := \operatorname{argmax}_{\mathbf{u} \in \mathcal{S}_k} \langle \nabla f(\mathbf{x}_k), \mathbf{u} \rangle$ ,  $\mathbf{d}_k^{\text{Away}} := \mathbf{x}_k - \mathbf{s}_k$ 

5:   if  $-\langle \nabla f(\mathbf{x}_k), \mathbf{d}_k^{\text{FW}} \rangle \geq -\langle \nabla f(\mathbf{x}_k), \mathbf{d}_k^{\text{Away}} \rangle$  then
6:      $\mathbf{d}_k := \mathbf{d}_k^{\text{FW}}$  with  $\lambda_{\max} := 1$ 
7:   else
8:      $\mathbf{d}_k := \mathbf{d}_k^{\text{Away}}$  with  $\lambda_{\max} := \frac{\alpha_k^{\mathbf{s}_k}}{1 - \alpha_k^{\mathbf{s}_k}}$ 
9:   end if
10:   $\mathbf{x}_{k+1} := \mathbf{x}_k + \lambda_k \mathbf{d}_k$  with  $\lambda_k \in [0, \lambda_{\max}]$  via line-search
11:  Update active set  $\mathcal{S}_{k+1}$  and  $\{\alpha_{k+1}^{\mathbf{v}}\}_{\mathbf{v} \in \mathcal{S}_{k+1}}$ 
12:   $k := k + 1$ 
13: end while
14: return  $(\mathbf{x}_k, \mathcal{S}_k, w(\mathbf{x}_k, \mathcal{S}_k))$ 

```

We now briefly outline how after a finite number of iterations T we can guarantee that $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$ and $\mathbf{x}^* \in \operatorname{co}(\mathcal{S}_k)$ for $k \geq K_0$. In particular, using Assumption 2.4, if the primal gap is made sufficiently small and the iterate \mathbf{x}_k is not contained in $\mathcal{F}(\mathbf{x}^*)$, then the AFW algorithm will continuously drop those vertices in \mathcal{S}_k that are not in $\mathcal{F}(\mathbf{x}^*)$, until the iterates reach $\mathcal{F}(\mathbf{x}^*)$ (see [Garber \(2020, Theorem 5\)](#), or Theorem A.5 in Appendix A.1, included for completeness).

Theorem 3.1. *If the strict complementarity assumption is satisfied (Assumption 2.4) and the primal gap satisfies $f(\mathbf{x}_k) - f(\mathbf{x}^*) < 1/2 \min \left\{ (\tau / (2D(\sqrt{L/m} + 1)))^2 / L, \tau, LD^2 \right\}$ then the following holds for the AFW algorithm (Algorithm 5):*

1. If $\mathbf{x}_k \notin \mathcal{F}(\mathbf{x}^*)$, AFW will perform an away step that drops a vertex $\mathbf{s}_k \in \operatorname{vert}(\mathcal{X}) \setminus \mathcal{F}(\mathbf{x}^*)$.
2. If $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, AFW will either perform a Frank-Wolfe step with a vertex $\mathbf{v}_k \in \operatorname{vert}(\mathcal{F}(\mathbf{x}^*))$ or an

away-step with a vertex $\mathbf{s}_k \in \text{vert}(\mathcal{F}(\mathbf{x}^*))$. Regardless of which step is chosen, the iterate will satisfy:

$$w(\mathbf{x}_k, \mathcal{S}_k) \leq LD\sqrt{2(f(\mathbf{x}_k) - f(\mathbf{x}^*))}/m.$$

Assuming that $\mathbf{x}_0 \in \text{vert}(\mathcal{X})$ in the AFW algorithm, and using the primal gap convergence gap guarantee in [Lacoste-Julien & Jaggi \(2015, Theorem 1\)](#), we can bound the number of iterations until $f(\mathbf{x}_k) - f(\mathbf{x}^*)$ satisfies the requirement in [Theorem 3.1](#). Using this bound, and the fact that the AFW algorithm can pick up at most one vertex per iteration, we can bound the number of iterations until $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$. Note that by the second claim in [Theorem 3.1](#), this means that when $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, then the iterates will not leave $\mathcal{F}(\mathbf{x}^*)$. Furthermore, once the iterates are inside the optimal face, there are two options: if $\mathbf{x}^* = \mathcal{F}(\mathbf{x}^*)$, then the AFW algorithm will exit once $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, as $w(\mathbf{x}_k, \mathcal{S}_k) = 0$, otherwise if $\mathbf{x}^* \notin \text{vert}(\mathcal{X})$ (the case of interest in our setting, by [Assumption 2.5](#)), then we need to prove that after a given number of iterations the active set will satisfy $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$. We prove the former using [Fact 3.2](#) (a variation of [Diakonikolas et al. \(2020, Fact B.3\)](#)).

Fact 3.2 (Critical strong Wolfe gap). There exists a $w_c > 0$ such that for any subset $\mathcal{S} \subseteq \text{vert}(\mathcal{F}(\mathbf{x}^*))$ and point $\mathbf{x} \in \mathcal{F}(\mathbf{x}^*)$ with $\mathbf{x} \in \text{co}(\mathcal{S})$ and $w(\mathbf{x}, \mathcal{S}) \leq w_c$ it follows that $\mathbf{x}^* \in \text{co}(\mathcal{S})$.

Remark 3.3. The critical strong Wolfe gap in [Fact 3.2](#), is a crucial parameter in the coming proofs. However, like the strict complementarity parameter τ ([Guélat & Marcotte, 1986; Garber, 2020](#)) and the critical radius r_c defined in [Diakonikolas et al. \(2020\)](#), the critical strong Wolfe gap can be arbitrarily small for some problems. Fortunately, as we will show in the proofs to come, it only affects the length of the burn-in phase of the accelerated algorithm, and moreover this dependence is logarithmic. In [Remark A.8](#) in the Appendix we show a simple example for which one can compute w_c exactly, and we give a lower bound and an upper bound for w_c for Problem (P).

With these tools at hand, we have the bound shown in [Theorem 3.4](#) (see [Appendix A](#) for the proof).

Theorem 3.4. Assume that the AFW algorithm ([Algorithm 5](#)) is run starting with $\mathbf{x}_0 \in \text{vert}(\mathcal{X})$. If the strict complementarity assumption ([Assumption 2.4](#)) is satisfied and $\mathbf{x}^* \notin \text{vert}(\mathcal{X})$, then for $k \geq K_0$ with

$$K_0 = \frac{32L}{m \ln 2} \left(\frac{D}{\delta(\mathcal{X})} \right)^2 \cdot \log \left(\frac{2w(\mathbf{x}_0, \mathcal{S}_0)}{\min \left\{ \frac{1}{L} \left(\frac{\tau}{(2D(\sqrt{L/m+1}))} \right)^2, \tau, LD^2, 2w_c \right\}} \right),$$

where $\delta(\mathcal{X})$ is the pyramidal width from [Definition 2.2](#) and $w_c > 0$ is the critical strong Wolfe gap from [Fact 3.2](#), we

have that $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$. Moreover:

$$w(\mathbf{x}_k, \mathcal{S}_k) \leq LD\sqrt{2(f(\mathbf{x}_k) - f(\mathbf{x}^*))}/m.$$

3.3. Parameter-free Projection-based Acceleration

The main idea for obtaining parameter-free projection-based acceleration is to use the regularization trick of [Nesterov \(2012\)](#) to obtain a near-optimal method for minimizing $\|G_\eta(\mathbf{x})\|$ for a smooth convex function. Then a near-optimal method for smooth strongly convex minimization is obtained by restarting this method every time $\|G_\eta(\mathbf{x})\|$ is halved.

The restarting approach is important here, as it removes the requirement of knowing the parameter m . However, there are a few technical challenges in implementing the near-optimal method for minimizing $\|G_\eta(\mathbf{x})\|$ without the knowledge of the parameter L or the distance to \mathbf{x}^* , which is needed for setting the value of the regularization parameter. Some of these challenges were addressed in [Ito & Fukuda \(2019\); Nesterov \(2013\)](#). However, as discussed before, the approaches from [Ito & Fukuda \(2019\)](#) and [Nesterov \(2013\)](#) are insufficient for our purposes, as they assume access to exact projections onto the feasible set (and thus, exact evaluations of the gradient mapping).

In the following, we first present a near-optimal method for minimizing $\|G_\eta(\mathbf{x})\|$ for a smooth convex function that does not require knowledge of L and works with inexact projections. We then show how to couple this method with adaptive tuning of the regularization parameter and restarts to obtain an overall near-optimal and parameter-free method for smooth strongly convex minimization. This subsection can be read independently from the rest of the paper.

3.3.1. SMALL GRADIENT MAPPING OF SMOOTH FUNCTIONS

Let $\mathcal{C} \subseteq \mathbb{R}^n$ be a closed, convex, nonempty set and assume that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is L -smooth on \mathcal{C} . The rough idea of the regularization trick is the following: instead of working directly with f , use a method for smooth strongly convex minimization to minimize $f_\sigma(\mathbf{x}) = f(\mathbf{x}) + \frac{\sigma}{2}\|\mathbf{x} - \mathbf{x}_0\|^2$ for some sufficiently small $\sigma > 0$ (for accuracy $\epsilon > 0$, $\sigma = \Theta\left(\frac{\epsilon}{\|\mathbf{x}_0^* - \mathbf{x}_0\|}\right)$ suffices, where $\mathbf{x}_0^* \in \arg\min_{\mathbf{x} \in \mathcal{C}} f(\mathbf{x})$). As we select σ ourselves, the method can assume knowledge of the strong convexity parameter σ of $f_\sigma(\mathbf{x})$.

The method that we employ here is a variant of $\mu\text{AGD+}$ from [Cohen et al. \(2018\)](#), which is similar to [Diakonikolas et al. \(2020\)](#). However, unlike [Cohen et al. \(2018\); Diakonikolas et al. \(2020\)](#), this method is adapted to work with an unknown smoothness parameter and to provide convergence guarantees on $\|G_\eta(\mathbf{x})\|$. One iteration of the algorithm is provided in [AGD-Iter \(Algorithm 2\)](#). In the

algorithm statement, we use $\tilde{\sim}$ argmin to indicate that the function that follows the argmin is minimized to additive error ϵ . The main result is summarized in the following lemma, while the complete analysis is deferred to Appendix B.

Algorithm 2 AGD-Iter($\mathbf{y}_{k-1}, \mathbf{v}_{k-1}, \mathbf{z}_{k-1}, A_{k-1}, \eta_k, \sigma, \epsilon_0, \eta_0$)

```

1:  $\eta_k = \eta_k/2$ 
2: repeat
3:    $\eta_k = 2\eta_k$ 
4:    $\theta_k = \sqrt{\frac{\sigma}{2(\eta_k + \sigma)}}$ ,  $a_k = \frac{\theta_k}{1-\theta_k} A_{k-1}$ 
5:    $\epsilon_k^\ell = \theta_k \epsilon_0/4$ ,  $\epsilon_k^M = a_k \epsilon_0/4$ 
6:    $\mathbf{x}_k = \frac{1}{1+\theta_k} \mathbf{y}_{k-1} + \frac{\theta_k}{1+\theta_k} \mathbf{v}_{k-1}$ 
7:    $\mathbf{z}_k = \mathbf{z}_{k-1} - a_k \nabla f_\sigma(\mathbf{x}_k) + \sigma a_k \mathbf{x}_k$ 
8:    $\mathbf{v}_k \stackrel{\epsilon_k^M}{\sim} \operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} M_k(\mathbf{u})$ , where  $M_k(\mathbf{u}) = -\langle \mathbf{z}_k, \mathbf{u} \rangle + \frac{\sigma A_{k-1} + \eta_0}{2} \|\mathbf{u}\|^2$ 
9:    $\hat{\mathbf{y}}_k = (1 - \theta_k) \mathbf{y}_{k-1} + \theta_k \mathbf{v}_k$ 
10:   $\mathbf{y}_k \stackrel{\epsilon_k^\ell}{\sim} \operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} \ell_k(\mathbf{u})$ , where  $\ell_k(\mathbf{u}) = \langle \nabla f_\sigma(\hat{\mathbf{y}}_k), \mathbf{u} - \hat{\mathbf{y}}_k \rangle + \frac{\eta_k + \sigma}{2} \|\mathbf{u} - \hat{\mathbf{y}}_k\|^2$ 
11:  until  $f(\hat{\mathbf{y}}_k) \leq f(\mathbf{x}_k) + \langle \nabla f(\mathbf{x}_k), \hat{\mathbf{y}}_k - \mathbf{x}_k \rangle + \frac{\eta_k}{2} \|\hat{\mathbf{y}}_k - \mathbf{x}_k\|^2$  and  $f(\mathbf{y}_k) \leq f(\hat{\mathbf{y}}_k) + \langle \nabla f(\hat{\mathbf{y}}_k), \mathbf{y}_k - \hat{\mathbf{y}}_k \rangle + \frac{\eta_k}{2} \|\mathbf{y}_k - \hat{\mathbf{y}}_k\|^2$ 
12:   $A_k = A_{k-1} + a_k$ ,  $\tilde{G}_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k) = (\eta_k + \sigma)(\hat{\mathbf{y}}_k - \mathbf{y}_k)$ 
13: return  $\eta_k, A_k, \mathbf{z}_k, \mathbf{v}_k, \hat{\mathbf{y}}_k, \mathbf{y}_k, \tilde{G}_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)$ 

```

Lemma 3.5. Let $\mathcal{C} \subseteq \mathbb{R}^n$ be a closed convex set and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be an L -smooth function on \mathcal{C} . Let $\mathbf{x}_0 \in \mathcal{C}$ be an arbitrary initial point, and, given $\sigma > 0$, define $f_\sigma(\mathbf{x}) = f(\mathbf{x}) + \frac{\sigma}{2} \|\mathbf{x} - \mathbf{x}_0\|^2$, $\mathbf{x}_\sigma^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{C}} f_\sigma(\mathbf{x})$.

Let $\mathbf{z}_0 = (\eta_0 + \sigma)\mathbf{x}_0 - \nabla f(\mathbf{x}_0)$, $\mathbf{y}_0 = \mathbf{v}_0 = \hat{\mathbf{y}}_0 \stackrel{\epsilon_0^M}{\sim} \operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} M_0(\mathbf{u})$, where $\epsilon_0^M > 0$, $M_0(\mathbf{u})$ is defined as in Algorithm 2, and the estimate η_0 is doubled until $f(\mathbf{y}_0) \leq f(\mathbf{x}_0) + \langle \nabla f(\mathbf{x}_0), \mathbf{y}_0 - \mathbf{x}_0 \rangle + \frac{\eta_0}{2} \|\mathbf{y}_0 - \mathbf{x}_0\|^2$, same as in Algorithm 2. Given $\eta_0 > 0$, sequence $\{a_k\}_{k \geq 0}$, $A_k = \sum_{i=0}^k a_i$, $\theta_k = \frac{a_k}{A_k}$, and the sequences of errors $\{\epsilon_k^M\}_{k \geq 0}$, $\{\epsilon_k^\ell\}_{k \geq 0}$, let the sequences of points $\{\hat{\mathbf{y}}_k, \mathbf{y}_k\}_{k \geq 0}$ evolve according to Algorithm 2 for $k \geq 1$.

If $a_0 = A_0 = 1$ and $\theta_k = \frac{a_k}{A_k} \leq \sqrt{\frac{\sigma}{2(\eta_k + \sigma)}}$ for $k \geq 1$, then for all $k \geq 1$

$$\frac{\|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2}{\eta_k + \sigma} \leq \frac{\eta_0}{A_k} \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2 + \frac{2}{A_k} \sum_{i=0}^k (2\epsilon_i^M + \epsilon_i^\ell),$$

where $G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)$ is the gradient mapping w.r.t. f_σ , at $\hat{\mathbf{y}}_k$.

In particular, if $a_0 = A_0 = 1$, $\theta_k = \frac{a_k}{A_k} = \sqrt{\frac{\sigma}{2(\eta_k + \sigma)}}$ for $k \geq 1$, $\epsilon_k^M = \frac{a_k \epsilon^2}{8}$, and $\epsilon_k^\ell \leq \frac{a_k}{A_k} \frac{\epsilon^2}{8}$, then $\frac{1}{\eta_k + \sigma} \|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2 \leq \epsilon^2$ after at most

$$k = O\left(\sqrt{\frac{L}{\sigma}} \log\left(\frac{L \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|}{\epsilon}\right)\right)$$

iterations. Further, the total number of first-order queries to f and oracle queries to inexact projections is at most

$$k' = k + 2 \log\left(\frac{2L}{\eta_0}\right).$$

3.3.2. ADAPTIVE REGULARIZATION AND RESTARTS

We now provide the full details for the updates of the accelerated sequence. There are two main questions that need to be addressed: (1) how to adaptively adjust the regularization parameter σ and (2) how to perform adaptive restarts based on inexact evaluations of the gradient mapping. For the former, note that, to obtain a near-optimal algorithm, σ should not be allowed to decrease too much (and in fact, needs to satisfy $\sigma = \Omega(m)$; see the proof of Theorem 3.6). For the latter, we can rely on the strong convexity of the function $\ell_k(\mathbf{u}) = \langle \nabla f_\sigma(\hat{\mathbf{y}}_k), \mathbf{u} - \mathbf{y}_k \rangle + \frac{\eta_k + \sigma}{2} \|\mathbf{u} - \hat{\mathbf{y}}_k\|^2$ to relate the exact and the inexact gradient mapping at $\hat{\mathbf{y}}_k$.

We state the main convergence bound of the ACC algorithm (Algorithm 3) here, while the detailed description of the algorithm and its analysis are deferred to Appendix B, for space considerations. Note that the algorithm name stems from *ACCeleration*.

Algorithm 3 ACC($\mathbf{x}_0, \eta_0, \sigma$)

```

1:  $\sigma = 2\sigma$ 
2: repeat
3:    $\sigma = \sigma/2$ 
4:   Run a minimization procedure for  $\ell_0(\mathbf{u}) = \langle \nabla f(\mathbf{x}_0), \mathbf{u} - \mathbf{x}_0 \rangle + \frac{\eta_0 + \sigma}{2} \|\mathbf{u} - \mathbf{x}_0\|^2$ . Halt when the current iterate  $\mathbf{y}$  of the procedure satisfies  $\ell_0(\mathbf{y}) - \min_{\mathbf{u} \in \mathcal{C}} \ell_0(\mathbf{u}) \leq \epsilon_0$ , where  $\epsilon_0 = \frac{\eta_0 + \sigma}{32} \|\mathbf{y}_0 - \mathbf{x}_0\|^2 = \frac{1}{32(\eta_0 + \sigma)} \|\tilde{G}_{\eta_0 + \sigma}^\sigma(\mathbf{x}_0)\|^2$ .
5:   Set  $\hat{\mathbf{y}}_0 = \mathbf{v}_0 = \mathbf{y}_0$ ;  $\mathbf{z}_0 = (\eta_0 + \sigma)\mathbf{x}_0 - \nabla f(\mathbf{x}_0)$ 
6:    $a_0 = A_0 = 1$ 
7:   repeat
8:      $k = k + 1$ 
9:      $\eta_k, A_k, \mathbf{z}_k, \mathbf{v}_k, \hat{\mathbf{y}}_k, \mathbf{y}_k, \tilde{G}_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k) = \text{AGD-Iter}(\mathbf{y}_{k-1}, \mathbf{v}_{k-1}, \mathbf{z}_{k-1}, A_{k-1}, \eta_{k-1}, \sigma, \epsilon_0, \eta_0)$ 
10:    until  $\frac{1}{\eta_k + \sigma} \|\tilde{G}_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2 \leq \frac{9\epsilon_0}{4}$ 
11:    until  $\frac{\sigma}{\sqrt{\eta_k + \sigma}} \|\hat{\mathbf{y}}_k - \mathbf{x}_0\| \leq \sqrt{\epsilon_0}$ 
12:  return  $\hat{\mathbf{y}}_k, \eta_k \sigma$ 

```

Theorem 3.6. Let $\mathcal{C} \subseteq \mathbb{R}^n$ be a closed convex set, and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a function that is L -smooth and m -strongly convex on \mathcal{C} . Let $\mathbf{x}_0 \in \mathcal{C}$ be an arbitrary initial point, and let σ_0, η_0 be such that $m \leq \sigma_0 \leq \eta_0 \leq L$. Let $\mathbf{x}_0^{\text{out}} = \mathbf{x}_0$, and for $k \geq 0$, consider the following updates:

$$\mathbf{x}_{k+1}^{\text{out}}, \eta_{k+1}, \sigma_{k+1} = \text{ACC}(\mathbf{x}_k^{\text{out}}, \eta_k, \sigma_k),$$

where ACC is provided in Algorithm 3. Let $G_{\eta + \sigma}$ denote the gradient mapping of f on \mathcal{C} with parameter

$\eta + \sigma$. Then, for any $\epsilon > 0$, $\|G_{\eta_k + \sigma_k}(\mathbf{x}_k^{\text{out}})\| \leq \epsilon$ for $k = O\left(\log\left(\frac{L\|G_{\eta_0 + \sigma_0}(\mathbf{x}_0)\|}{m\epsilon}\right)\right)$. The algorithm for computing $\mathbf{x}_k^{\text{out}}$ utilizes a total number of

$$K = \left(\sqrt{\frac{L}{m}} \log\left(\frac{L}{m}\right) \log\left(\frac{L}{m} \frac{\|G_{\eta_0 + \sigma_0}(\mathbf{x}_0)\|}{\epsilon}\right)\right)$$

queries to the FOO for f and an inexact and efficiently computable projection oracle for \mathcal{C} , without knowledge of any of the problem parameters.

Observe that, up to the $\log(L/m)$ factor, the bound in Theorem 3.6 is optimal for the class of smooth strongly convex functions, under the local oracle model (which includes the FOO model as a special case).

3.4. Coupling of Methods and Local Acceleration

We now show how to couple the AFW algorithm (Algorithm 5 in Appendix A) with the algorithm described in the previous subsection to achieve local acceleration. The key observation is that after a burn-in phase whose length is independent of ϵ , every active set that AFW constructs contains $\mathbf{x}^* = \arg\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$. Thus, minimizing $f(\mathbf{x})$ over \mathcal{X} becomes equivalent to minimizing $f(\mathbf{x})$ over the convex hull of the active set of AFW. Of course, the algorithm needs to adapt to this setting without knowledge of when the burn-in phase has ended. The pseudocode for the resulting algorithm Parameter-Free Locally Accelerated Conditional Gradients (PF-LaCG) is provided in Algorithm 4. The coupling between AFW and ACC in PF-LaCG is illustrated on a simple example in Fig. 1.

For simplicity, Algorithm 4 is stated with the accelerated sequence started with $\mathbf{x}_0 \in \text{vert}(\mathcal{X})$ and with the active set $\mathcal{S}_0 = \{\mathbf{x}_0\}$. As \mathcal{S}_0 contains only one vertex, there is no need to run the accelerated algorithm until a later restart (triggered by the condition in Line 6) where the active set contains more than one vertex occurs. Additionally, for the bound on the number of oracle queries in Theorem 3.7 to hold, we need the iterations of ACC (Algorithm 6) and AFW (Algorithm 5) to be aligned, in the sense that one iteration of ACC (Algorithm 2) occurs in parallel to one iteration of AFW. This is only needed for the theoretical bound on oracle complexity. In practice, the two algorithms can be run in parallel without aligning the iterations, so that the overall execution time (modulo coordination at restarts) is never higher than for running AFW alone.

We are now ready to state our main result. The complete proof is deferred to Appendix C.

Theorem 3.7. *Let $\mathcal{X} \subseteq \mathbb{R}^n$ be a closed convex polytope of diameter D , and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a function that is L -smooth and m -strongly convex on \mathcal{X} . Let Assumptions 2.2 and 2.4 be satisfied for f, \mathcal{X} . Denote*




Figure 1. An example of coupling between AFW and ACC in PF-LaCG on a tetrahedron as the feasible set, starting from initial point \mathbf{x}_0 with the base of the tetrahedron as its support \mathcal{S}_0 . The two algorithms are run in parallel from \mathbf{x}_0 : AFW optimizes over the entire tetrahedron, allowing it to add and remove vertices, while ACC optimizes over the base of the tetrahedron only and it cannot converge to the optimal point \mathbf{x}^* , as $\mathbf{x}^* \notin \text{co}(\mathcal{S}_0)$. After several iterations, once the restart criterion for AFW is triggered, PF-LaCG chooses the output point of AFW over that of ACC, as $w^{\text{AFW}} \leq \min\{w^{\text{ACC}}, w_{\text{prev}}^{\text{ACC}}/2\}$, hence a PF-LaCG restart occurs at \mathbf{x}^R . For ease of exposition we assume that the point outputted by AFW is contained in $\mathcal{F}(\mathbf{x}^*)$ after a single halving of $w(\mathbf{x}, \mathcal{S})$, although in practice several restarts may be needed for AFW to reach $\mathcal{F}(\mathbf{x}^*)$. Since \mathbf{x}^R is on the optimal face $\mathcal{F}(\mathbf{x}^*)$, PF-LaCG has completed the burn-in phase. The two algorithms again run in parallel from \mathbf{x}^R after the restart. However, ACC converges to the optimal \mathbf{x}^* at an accelerated rate, much faster than AFW. Hence, local acceleration is achieved by PF-LaCG while being at least as fast as vanilla AFW.

$\mathbf{x}^* = \arg\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$. Given $\epsilon > 0$, let \mathbf{x}^{out} be the output point of PF-LaCG (Algorithm 4), initialized at an arbitrary vertex \mathbf{x}_0 of \mathcal{X} . Then $w(\mathbf{x}^{\text{out}}, \mathcal{S}^{\text{out}}) \leq \epsilon$ and PF-LaCG uses a total of at most

$$K = O\left(\min\left\{\log\left(\frac{w(\mathbf{x}_0, \mathcal{S}_0)}{LD^2}\right) + \frac{LD^2}{m\delta^2} \log\left(\frac{w(\mathbf{x}_0, \mathcal{S}_0)}{\epsilon}\right), K_0 + K_1 + \sqrt{\frac{L}{m}} \log\left(\frac{L}{m}\right) \log\left(\frac{LD}{m\delta}\right) \log\left(\frac{LD}{\epsilon}\right)\right\}\right)$$

queries to the FOO for f and the LMO for \mathcal{X} , where

$$K_0 = \frac{32L}{m \ln 2} \left(\frac{D}{\delta(\mathcal{X})}\right)^2 \cdot \log\left(\frac{2w(\mathbf{x}_0, \mathcal{S}_0)}{\min\left\{\frac{1}{L} \left(\frac{\tau}{2D(\sqrt{L/m}+1)}\right)^2, \tau, LD^2, 2w_c\right\}}\right),$$

$$\text{and } K_1 = \frac{128LD^2}{m\delta^2}.$$

Strong Wolfe Gap as an Upper Bound For Primal Gap.

Note that while in Theorem 3.7 we show a convergence rate for $w(\mathbf{x}, \mathcal{S})$, this translates to the same convergence rate in primal gap, as the aforementioned quantity is an upper bound on the strong Wolfe gap and hence an upper bound on the primal gap (Kerdreux et al., 2019). Furthermore, we

Algorithm 4 PF-LaCG($\mathbf{x}_0 \in \text{vert}(\mathcal{X}), \epsilon > 0$)

```

1:  $\mathcal{S}^{\text{out}} = \mathcal{S}^{\text{AFW}} = \mathcal{S}^{\text{ACC}} = \{\mathbf{x}_0\}$ 
2:  $\mathbf{x}^{\text{out}} = \mathbf{x}^{\text{AFW}} = \mathbf{x}^{\text{ACC}} = \mathbf{x}_0$ 
3:  $w^{\text{out}} = w_{\text{prev}}^{\text{AFW}} = w(\mathbf{x}_0, \mathcal{S}_0)$ 
4: while  $w^{\text{out}} > \epsilon$  do
5:   Run AFW and restarted ACC (Theorem 3.6) in parallel, and let  $(\mathbf{x}^{\text{AFW}}, \mathcal{S}^{\text{AFW}})$  and  $(\mathbf{x}^{\text{ACC}}, \mathcal{S}^{\text{ACC}})$  denote their respective most recent output points and active sets; note that ACC is run on  $\mathcal{C} = \text{co}(\mathcal{S}^{\text{ACC}})$ 
6:   if  $w^{\text{AFW}} = w(\mathbf{x}^{\text{AFW}}, \mathcal{S}^{\text{AFW}}) \leq \frac{1}{2}w_{\text{prev}}^{\text{AFW}}$  then
7:      $w_{\text{prev}}^{\text{AFW}} = w^{\text{AFW}}, w_{\text{prev}}^{\text{ACC}} = w^{\text{ACC}},$ 
8:     Compute  $w^{\text{ACC}} = w(\mathbf{x}^{\text{ACC}}, \mathcal{S}^{\text{ACC}})$ 
9:     if  $w^{\text{AFW}} \leq \min\{w^{\text{ACC}}, w_{\text{prev}}^{\text{ACC}}/2\}$  then
10:       $\mathbf{x}^{\text{ACC}} = \mathbf{x}^{\text{AFW}}, \mathcal{S}^{\text{ACC}} = \mathcal{S}^{\text{AFW}}$ 
11:       $\mathbf{x}^{\text{out}} = \mathbf{x}^{\text{AFW}}, \mathcal{S}^{\text{out}} = \mathcal{S}^{\text{AFW}}$ 
12:       $w^{\text{out}} = w^{\text{AFW}}$ 
13:     else
14:       if  $|\mathcal{S}^{\text{ACC}}| \leq |\mathcal{S}^{\text{AFW}}|$  then
15:          $\mathbf{x}^{\text{AFW}} = \mathbf{x}^{\text{ACC}}, \mathcal{S}^{\text{AFW}} = \mathcal{S}^{\text{ACC}}$ 
16:          $w^{\text{AFW}} = w^{\text{ACC}}$ 
17:       end if
18:        $\mathbf{x}^{\text{out}} = \mathbf{x}^{\text{ACC}}, \mathcal{S}^{\text{out}} = \mathcal{S}^{\text{ACC}}, w^{\text{out}} = w^{\text{ACC}}$ 
19:     end if
20:   end if
21: end while
22: return  $\mathbf{x}^{\text{out}}$ 

```

remark that $w(\mathbf{x}, \mathcal{S}) = 0$ if and only if $\mathbf{x} = \mathbf{x}^*$ when $\mathcal{S} \subseteq \text{vert}(\mathcal{X})$ is a proper support of \mathbf{x} with respect to the polytope \mathcal{X} . In the next section we illustrate our computational results with respect to the primal gap $f(\mathbf{x}) - f(\mathbf{x}^*)$, as this quantity is more common in the optimization literature.

4. Computational Experiments

We numerically demonstrate that PF-LaCG (Algorithm 4) when implemented in Python 3 outperforms other parameter-free CG methods in both iteration count and in wall-clock time. Most notably, we compare against the AFW, PFW, *Decomposition-Invariant CG* (DICG) (in the case of the probability simplex, as this algorithm only applies to 0 – 1 polytopes (Garber & Meshi, 2016)) and the *Lazy AFW* (AFW (Lazy)) (Braun et al., 2017) algorithms. As predicted by our theoretical results, the improvement is observed locally, once the iterates of the algorithm reach the optimal face. Further, we empirically observe on the considered examples that the active sets do not become too large, which is important for the accelerated algorithm to have an edge over standard CG updates in the overall wall-clock time. Our code can be found at <https://github.com/ericlincc/Parameter-free-LaCG>.

Similar to Diakonikolas et al. (2020), we solve the min-

imization subproblems (projections onto the convex hull of the active sets) within ACC (Algorithm 6) using Nesterov’s accelerated gradient descent (Nesterov, 2018) with $\mathcal{O}(n \log n)$ projections onto the simplex described in Duchi et al. (2008, Algorithm 1). Even though in our analysis we consider coupling accelerated sequence with AFW, we note that PF-LaCG can be coupled with any CG variant that maintains an active set, such as standard AFW or PFW.

Since the execution of local acceleration within PF-LaCG is completely independent of the execution of the coupled CG variant, it is possible to run the locally accelerated algorithm in parallel on a separate core within one machine or even on a secondary machine, allowing us to utilize more computational power with the goal of solving large-scale problems to high accuracy in less computing time. In our experiments, we implemented the former approach where we run each process using one CPU core, and we run the accelerated algorithm of PF-LaCG on a separate process. This approach enables us to guarantee that PF-LaCG is not slower than the CG variant such as AFW and PFW barring negligible process creation overhead and inter-process communication while achieving much faster convergence once we have reached the optimal face.

Probability Simplex. The unit probability simplex, although a toy example, can give us insight into the behaviour of PF-LaCG. We know that after a restart in which the AFW active set satisfies $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$, we should expect to see accelerated convergence from the iterates computed by the ACC algorithm. Given the structure of the probability simplex, this is easy to check in the experiments once we have a high-accuracy solution to the minimization problem. The function being minimized in this example is $f(\mathbf{x}) = \mathbf{x}^T (M^T M + \alpha \mathbf{1}_n) \mathbf{x} / 2 + \mathbf{b}^T \mathbf{x}$, where $M \in \mathbb{R}^{n \times n}$ and $\mathbf{b} \in \mathbb{R}^n$ have entries sampled uniformly at random between 0 and 1 and $n = 10000$. The parameter $\alpha = 500$ is set so that the objective function satisfies $m \approx 500$. The resulting condition number is $L/m = 50000$, and the number of nonzero elements in \mathbf{x}^* is around 320. The AFW algorithm in PF-LaCG (AFW) satisfies that $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$ around iteration 400, consequently we achieve the accelerated convergence rate from then onwards. The same can be said regarding PF-LaCG (PFW) around iteration 350.

Structured LASSO Regression. The LASSO is an extremely popular sparse regression analysis method in statistics, and has become an important tool in ML. In many applications, we can impose additional structure on the solution being learnt through the addition of linear constraints; this is useful, for example, when learning sparse physics dynamics from data (see Carderera et al. (2021)), in which we can impose additional linear equality constraints on the

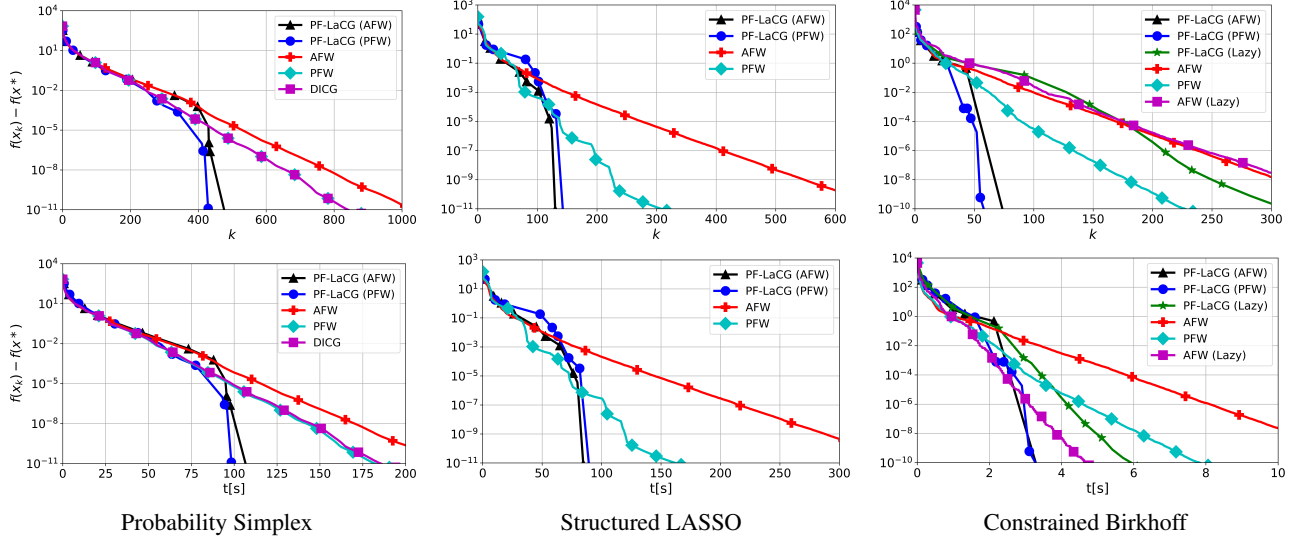


Figure 2. **Numerical performance of PF-LaCG:** Top row depicts primal gap convergence in terms of iteration count, while bottom row depicts primal gap convergence in terms of time. Left-most column shows the results over the probability simplex, center column shows the results for the structured Lasso problem, and right-most column shows the results for the constrained Birkhoff polytope.

problem to reflect the symmetries present in the physical problem. This allows us to learn dynamics that are consistent with the underlying physics, and which potentially generalize better when predicting on unseen data. The objective function is the same as in the last section, except we now have $n = 1000$, $\alpha = 100$, and we choose the elements of \mathbf{b} uniformly from 0 to 100. This results in a condition number of $L/m = 250000$. The feasible region is the intersection of the ℓ_1 unit ball and a series of equality constraints. To generate the additional equality constraints, we sample 125 pairs of distinct integers (i, j) from $1 \leq i, j \leq n$ without replacement, and we set $x_i = x_j$ for each pair, adding 125 linear constraints.

Constrained Birkhoff Polytope. We also solve a matching with the same quadratic function as in the previous section with $\alpha = 1$, and where we have scaled the matrix $M^T M$ to have a maximum eigenvalue of 100000. This results in an objective function that has a condition number of $L/m = 100000$. The matching problem is solved over the Birkhoff polytope with $n = 400$ where we have imposed additional linear constraints. We sample 80 integers i from $1 \leq i \leq n$ without replacement, and we set $x_i = 0$ for the first 40 integers (to represent that certain matchings are not possible), and $x_i \leq 0.5$ for the remaining 40 integers to represent a maximum fractional matching. As in the LaCG algorithm (Diakonikolas et al., 2020), our approach is compatible with the lazification technique for AFW (Braun et al., 2017). In this last example we couple our algorithm with the AFW (Lazy) algorithm, resulting in the PF-LaCG (Lazy) algorithm. Moreover, we also benchmark against the AFW (Lazy) algorithm for reference.

5. Discussion

We have introduced a novel projection-free PF-LaCG algorithm for minimizing smooth and strongly convex functions over polytopes. This algorithm is parameter-free and locally accelerated. In particular, we have shown that after a finite burn-in phase (independent of the target error ϵ) PF-LaCG achieves a near-optimal accelerated convergence rate without knowledge of any of the problem parameters (such as the smoothness or strong convexity of the function). As mentioned in the introduction, global acceleration is generally not possible for CG-type methods. We have also demonstrated the improved locally accelerated convergence rate of PF-LaCG using numerical experiments. Some interesting questions for future research remain. For example, it is an interesting and practically relevant question whether local acceleration is possible for CG methods that are not active set-based, such as, e.g., DICG (Garber & Meshi, 2016), which would possibly lead to even faster algorithms.

Acknowledgments

This research was partially funded by NSF grant CCF-2007757, by the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin–Madison with funding from the Wisconsin Alumni Research Foundation, and by the Deutsche Forschungsgemeinschaft (DFG) through the DFG Cluster of Excellence MATH+ and the Research Campus Modal funded by the German Federal Ministry of Education and Research (fund numbers 05M14ZAM, 05M20ZBM).

References

- Bashiri, M. A. and Zhang, X. Decomposition-invariant conditional gradient for general polytopes with line search. In *NIPS*, pp. 2690–2700, 2017.
- Beck, A. *First-order methods in optimization*. MOS-SIAM, 2017.
- Beck, A. and Shtern, S. Linearly convergent away-step conditional gradient for non-strongly convex functions. *Mathematical Programming*, 164(1-2):1–27, 2017.
- Braun, G., Pokutta, S., and Zink, D. Lazifying conditional gradient algorithms. In *Proc. ICML’2017*, 2017.
- Braun, G., Pokutta, S., Tu, D., and Wright, S. Blended conditional gradients: The unconditioning of conditional gradients. In *Proc. ICML’19*, 2019.
- Carderera, A. and Pokutta, S. Second-order conditional gradient sliding. *arXiv preprint arXiv:2002.08907*, 2020.
- Carderera, A., Pokutta, S., Schütte, C., and Weiser, M. CINDy: Conditional gradient-based identification of non-linear dynamics—noise-robust recovery. *arXiv preprint arXiv:2101.02630*, 2021.
- Chen, S. S., Donoho, D. L., and Saunders, M. A. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.
- Cohen, M. B., Diakonikolas, J., and Orecchia, L. On acceleration with noise-corrupted gradients. In *Proc. ICML’18*, 2018.
- Combettes, C. W. and Pokutta, S. Complexity of linear minimization and projection on some sets. *arXiv preprint arXiv:2101.10040*, 2021.
- Combettes, C. W., Spiegel, C., and Pokutta, S. Projection-free adaptive gradients for large-scale optimization. *arXiv preprint arXiv:2009.14114*, 2020.
- Condat, L. Fast projection onto the simplex and the ℓ_1 ball. *Mathematical Programming*, 158(1):575–585, 2016.
- Diakonikolas, J. and Orecchia, L. The approximate duality gap technique: A unified theory of first-order methods. *SIAM Journal on Optimization*, 29(1):660–689, 2019.
- Diakonikolas, J., Carderera, A., and Pokutta, S. Locally accelerated conditional gradients. In *Proc. AISTATS’20*, 2020.
- Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *Proc. NIPS’08*, 2008.
- Dvurechensky, P., Ostroukhov, P., Safin, K., Shtern, S., and Staudigl, M. Self-concordant analysis of Frank-Wolfe algorithms. In *Proc. ICML’20*, 2020.
- Frank, M. and Wolfe, P. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- Garber, D. Faster projection-free convex optimization over the spectrahedron. In *Proc. NIPS’16*, 2016.
- Garber, D. Revisiting Frank-Wolfe for polytopes: Strict complementarity and sparsity. In *Proc. NeurIPS’20*, 2020.
- Garber, D. and Meshi, O. Linear-memory and decomposition-invariant linearly convergent conditional gradient algorithm for structured polytopes. In *Proc. NIPS’16*, 2016.
- Guélat, J. and Marcotte, P. Some comments on Wolfe’s ‘away step’. *Mathematical Programming*, 35(1):110–119, 1986.
- Gutman, D. H. and Peña, J. F. The condition of a function relative to a polytope. *arXiv preprint arXiv:1802.00271*, 2018.
- Hazan, E. and Luo, H. Variance-reduced and projection-free stochastic optimization. In *Proc. ICML’16*, 2016.
- Held, M., Wolfe, P., and Crowder, H. P. Validation of subgradient optimization. *Mathematical programming*, 6(1):62–88, 1974.
- Ito, M. and Fukuda, M. Nearly optimal first-order methods for convex optimization under gradient norm measure: An adaptive regularization approach. *arXiv preprint arXiv:1912.12004*, 2019.
- Jaggi, M. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *Proc. ICML’13*, 2013.
- Kerdreux, T., d’Aspremont, A., and Pokutta, S. Restarting Frank-Wolfe: Faster rates under Hölderian error bounds. In *Proc. AISTATS’19*, 2019.
- Kerdreux, T., d’Aspremont, A., and Pokutta, S. Projection-free optimization on uniformly convex sets. In *Proc. AISTATS’21*, 2021.
- Lacoste-Julien, S. and Jaggi, M. On the global linear convergence of Frank-Wolfe optimization variants. In *Proc. NIPS’15*, 2015.
- Lam, S. K., Pitrou, A., and Seibert, S. Numba: A LLVM-based Python JIT compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pp. 1–6, 2015.

- Lan, G. The complexity of large-scale convex programming under a linear optimization oracle. *arXiv preprint arXiv:1309.5550*, 2013.
- Lei, Q., Zhuo, J., Caramanis, C., Dhillon, I. S., and Dimakis, A. G. Primal-dual block generalized Frank-Wolfe. *Proc. NeurIPS'19*, 2019.
- Levitin, E. S. and Polyak, B. T. Constrained minimization methods. *USSR Computational mathematics and mathematical physics*, 6(5):1–50, 1966.
- Négier, G., Dresdner, G., Tsai, A., El Ghaoui, L., Locatello, F., Freund, R., and Pedregosa, F. Stochastic Frank-Wolfe for constrained finite-sum minimization. In *Proc. ICML'20*, 2020.
- Nesterov, Y. How to make the gradients small. *Optima. Mathematical Optimization Society Newsletter*, (88):10–11, 2012.
- Nesterov, Y. Gradient methods for minimizing composite functions. *Mathematical Programming*, (140(1)):125–161, 2013.
- Nesterov, Y. *Lectures on Convex Optimization*. Springer, 2018.
- Peña, J. and Rodríguez, D. Polytope conditioning and linear convergence of the Frank–Wolfe algorithm. *Mathematics of Operations Research*, 44(1):1–18, 2019.
- Pedregosa, F., Négier, G., Askari, A., and Jaggi, M. Linearly convergent Frank–Wolfe with backtracking line-search. In *Proc. AISTATS'20*, 2020.
- Roulet, V. and d’Aspremont, A. Sharpness, restart, and acceleration. *SIAM Journal on Optimization*, 30(1):262–289, 2020.
- Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Tsiligkaridis, T. and Roberts, J. On Frank-Wolfe optimization for adversarial robustness and interpretability. *arXiv preprint arXiv:2012.12368*, 2020.
- Zhang, M., Shen, Z., Mokhtari, A., Hassani, H., and Karbasi, A. One sample stochastic Frank-Wolfe. In *Proc. AISTATS'20*, 2020.
- Zhou, S., Gupta, S., and Udell, M. Limited memory Kelley’s method converges for composite convex and submodular objectives. *arXiv preprint arXiv:1807.07531*, 2018.

Parameter-free Locally Accelerated Conditional Gradients

Supplementary material

Outline. The appendix of the paper is organized as follows:

- Section A presents the proofs related to the AFW algorithm used in the main body of the paper.
- Section B presents the proofs related to the accelerated algorithm with inexact gradient mappings that is used in this work, as well as some useful properties of the gradient mapping.
- Section C shows how we can couple the AFW algorithm and the accelerated algorithm presented in Section B to achieve a parameter-free accelerated CG variant, dubbed *Parameter-free Locally Accelerated Conditional Gradients* (PF-LaCG).
- Section D presents the full details of the computational experiments performed in the paper.

A. Away-Step Frank-Wolfe (AFW)

We use a modified version of the *Away-Step Frank-Wolfe* (AFW) algorithm (Guélat & Marcotte, 1986; Lacoste-Julien & Jaggi, 2015) that is run until the Frank-Wolfe gap is halved, in a similar manner as in (Kerdreux et al., 2019). The only difference between the AFW algorithm presented in Algorithm 5 and the one from (Kerdreux et al., 2019) is that we substitute the condition that chooses between Frank-Wolfe steps (Line 6) and away-steps (Line 8) for the condition used in the classical Away-Step Frank Wolfe (AFW) algorithm, shown in Line 5. In the following, we will say that a step is a *full-progress step* if it is either a Frank-Wolfe Step (Line 6) or an away-step (Line 8) that is not a drop step, i.e., when $\lambda_k < \alpha_k^{\mathbf{s}_k} / (1 - \alpha_k^{\mathbf{s}_k})$, where $\alpha_k^{\mathbf{s}_k}$ is the barycentric coordinate of \mathbf{s}_k with respect to the current active set. Here we summarize the results that are utilized in our analysis.

The following definition is introduced for completeness, to state the results from (Kerdreux et al., 2019).

Definition A.1 (Away curvature). Given a problem (P), the away curvature C_f^A is defined as

$$C_f^A := \sup_{\substack{\mathbf{x}, \mathbf{s}, \mathbf{v} \in \mathcal{X} \\ \rho \in [0,1] \\ \mathbf{y} = \mathbf{x} + \rho(\mathbf{s} - \mathbf{v})}} \frac{2}{\rho^2} (f(\mathbf{y}) - f(\mathbf{x}) - \rho \langle \nabla f(\mathbf{x}), \mathbf{s} - \mathbf{v} \rangle).$$

As f is L -smooth, we have $C_f^A \leq LD^2$, where $D = \max_{\mathbf{x}, \mathbf{y} \in \mathcal{X}} \|\mathbf{y} - \mathbf{x}\|$ denotes the diameter of \mathcal{X} .

Fact A.2 ((Kerdreux et al., 2019, Lemma 3.6)). When f is m -strongly convex over \mathcal{X} and \mathcal{X} is a polytope satisfying δ -scaling inequality with $\delta > 0$, then we have for all $\mathbf{x} \in \mathcal{X}$:

$$f(\mathbf{x}) - \min_{\mathbf{y} \in \mathcal{X}} f(\mathbf{y}) \leq \frac{2w(\mathbf{x})^2}{m\delta^2}.$$

Proposition A.3 (Iterations needed to halve $w(\mathbf{x}, \mathcal{S})$ (Kerdreux et al., 2019, Proposition 4.1)). Let f be an L -smooth, m -strongly convex function with away curvature C_f^A and let \mathcal{X} be a polytope satisfying the δ -scaling inequality with $\delta > 0$. Assume that $\mathbf{x}_0 \in \mathcal{X}$ is such that $w(\mathbf{x}_0)/4 \leq C_f^A$. Then AFW (Algorithm 5) outputs an iterate $\mathbf{x}_K \in \mathcal{X}$ such that

$$w(\mathbf{x}_K, \mathcal{S}_K) \leq w(\mathbf{x}_0, \mathcal{S}_0)/2$$

after at most

$$K \leq |\mathcal{S}_0| - |\mathcal{S}_K| + \frac{128C_f^A}{m\delta^2}$$

iterations, where \mathcal{S}_0 and \mathcal{S}_K are the initial active set and the active set at iteration K respectively.

Algorithm 5 Away-Step Frank-Wolfe Algorithm: AFW($\mathbf{x}_0, \mathcal{S}_0$)

```

1:  $k := 0$ 
2: while  $w(\mathbf{x}_k, \mathcal{S}_k) > w(\mathbf{x}_0, \mathcal{S}_0)/2$  do
3:    $\mathbf{v}_k := \operatorname{argmin}_{\mathbf{u} \in \mathcal{X}} \langle \nabla f(\mathbf{x}_k), \mathbf{u} \rangle$  and  $\mathbf{d}_k^{\text{FW}} := \mathbf{v}_k - \mathbf{x}_k$ 
4:    $\mathbf{s}_k := \operatorname{argmin}_{\mathbf{u} \in \mathcal{S}_k} \langle -\nabla f(\mathbf{x}_k), \mathbf{u} \rangle$  with  $\mathcal{S}_k$  current active set and  $\mathbf{d}_k^{\text{Away}} := \mathbf{x}_k - \mathbf{s}_k$ 
5:   if  $-\langle \nabla f(\mathbf{x}_k), \mathbf{d}_k^{\text{FW}} \rangle \geq -\langle \nabla f(\mathbf{x}_k), \mathbf{d}_k^{\text{Away}} \rangle$  then
6:      $\mathbf{d}_k := \mathbf{d}_k^{\text{FW}}$  with  $\lambda_{\max} := 1$ 
7:   else
8:      $\mathbf{d}_k := \mathbf{d}_k^{\text{Away}}$  with  $\lambda_{\max} := \frac{\alpha_k^{\mathbf{s}_k}}{1 - \alpha_k^{\mathbf{s}_k}}$ 
9:   end if
10:   $\mathbf{x}_{k+1} := \mathbf{x}_k + \lambda_k \mathbf{d}_k$  with  $\lambda_k \in [0, \lambda_{\max}]$  via line-search
11:  Update active set  $\mathcal{S}_{k+1}$  and coefficients  $\{\alpha_{k+1}^{\mathbf{v}}\}_{\mathbf{v} \in \mathcal{S}_{k+1}}$ 
12:   $k := k + 1$ 
13: end while
14: return  $(\mathbf{x}_k, \mathcal{S}_k, w(\mathbf{x}_k, \mathcal{S}_k))$  where  $\mathbf{x}_k \in \mathcal{X}$  and  $w(\mathbf{x}_k, \mathcal{S}_k) \leq w(\mathbf{x}_0, \mathcal{S}_0)/2$ 
    
```

Proof. Note that by Fact A.2 we have that the 2-strong Wolfe primal bound holds. We only explain the part of the proof that differs from the proof in (Kerdreux et al., 2019), due to the modified condition in Line 5. This condition only affects how we bound $w(\mathbf{x}_0, \mathcal{S}_0)$, depending on whether we take a Frank-Wolfe step or an away-step, as follows:

$$w(\mathbf{x}_0, \mathcal{S}_0)/2 < w(\mathbf{x}_k, \mathcal{S}_k) = -\langle \nabla f(\mathbf{x}_k), \mathbf{d}_k^{\text{FW}} \rangle - \langle \nabla f(\mathbf{x}_k), \mathbf{d}_k^{\text{Away}} \rangle \leq -2 \langle \nabla f(\mathbf{x}_k), \mathbf{d}_k \rangle,$$

where the first inequality simply states that we have not yet halved $w(\mathbf{x}, \mathcal{S})$ (see Line 2), and the second inequality stems from the condition with which we choose the steps in Line 5. This allows us to claim that in case a Frank-Wolfe step is chosen we have that $w(\mathbf{x}_0, \mathcal{S}_0)/4 < \langle \nabla f(\mathbf{x}_k), \mathbf{d}_k \rangle$. Using this fact, we can proceed with the same arguments as in the proof of Kerdreux et al. (2019, Proposition 4.1), which is omitted. \square

Notice that Proposition A.3 requires an assumption that $w(\mathbf{x}_0)/4 \leq C_f^A$. This assumption is satisfied after a small number of iterations, as summarized in the following proposition.

Proposition A.4 ((Kerdreux et al., 2019, Proposition 4.2)). *Assume the AFW algorithm (Algorithm 5) is run repeatedly until it outputs a point $\mathbf{x} \in \mathcal{X}$ such that $w(\mathbf{x}, \mathcal{S})/4 \leq C_f^A$ where \mathcal{S} is a proper support for \mathbf{x} . This will happen after at most*

$$T_0 = \frac{16}{\log 2} \log \frac{w(\mathbf{x}_0, \mathcal{S}_0)}{2C_f^A} + |\mathcal{S}_0|$$

FOO and LMO calls.

Since $w(\mathbf{x}) \leq w(\mathbf{x}, \mathcal{S})$ for all $\mathbf{x} \in \mathcal{X}$ and any support \mathcal{S} for \mathbf{x} , Proposition A.4 implies that Proposition A.3 applies after at most T_0 initial iterations. The proof is omitted and can instead be found in Kerdreux et al. (2019).

A.1. Implications of Strict Complementarity

We reproduce here for completeness the relevant results from Garber (2020, Theorem 1), where it is shown that if the iterates are not in $\mathcal{F}(\mathbf{x}^*)$ and the primal gap is below a given tolerance, then all subsequent steps will be drop steps that drop vertices in the current active set that are not in $\mathcal{F}(\mathbf{x}^*)$.

Theorem A.5. *If the strict complementarity assumption is satisfied (Assumption 2.4) and the primal gap satisfies $f(\mathbf{x}_k) - f(\mathbf{x}^*) < 1/2 \min \left\{ (\tau/(2D(\sqrt{L/m} + 1)))^2/L, \tau, LD^2 \right\}$ then the following holds for the AFW algorithm (Algorithm 5):*

1. *If $\mathbf{x}_k \notin \mathcal{F}(\mathbf{x}^*)$, AFW will perform an away step that drops a vertex $\mathbf{s}_k \in \operatorname{vert}(\mathcal{X}) \setminus \mathcal{F}(\mathbf{x}^*)$.*

2. If $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, AFW will either perform a Frank-Wolfe step with a vertex $\mathbf{v}_k \in \text{vert}(\mathcal{F}(\mathbf{x}^*))$ or an away-step with a vertex $\mathbf{s}_k \in \text{vert}(\mathcal{F}(\mathbf{x}^*))$. Regardless of which step is chosen, the iterate will satisfy:

$$w(\mathbf{x}_k, \mathcal{S}_k) \leq \frac{LD\sqrt{2}}{\sqrt{m}} \sqrt{f(\mathbf{x}_k) - f(\mathbf{x}^*)}.$$

Proof. We first prove the first claim. Consider two vertices $\mathbf{v} \in \text{vert}(\mathcal{X}) \cap \mathcal{F}(\mathbf{x}^*)$ and $\mathbf{s} \in \text{vert}(\mathcal{X}) \setminus \mathcal{F}(\mathbf{x}^*)$, then we have that:

$$\begin{aligned} \langle \mathbf{s} - \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle &= \langle \mathbf{s} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{x}^* - \mathbf{v}, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{s} - \mathbf{v}, \nabla f(\mathbf{x}_k) - \nabla f(\mathbf{x}^*) \rangle \\ &\geq \tau - \|\mathbf{s} - \mathbf{v}\| \|\nabla f(\mathbf{x}_k) - \nabla f(\mathbf{x}^*)\| \\ &\geq \tau - LD \|\mathbf{x}_k - \mathbf{x}^*\| \\ &\geq \tau - LD\sqrt{2(f(\mathbf{x}_k) - f(\mathbf{x}^*))}/m \\ &\geq \tau/2, \end{aligned} \tag{A.1}$$

where the first inequality comes from the strict complementarity assumption, as $\mathbf{v} \in \text{vert}(\mathcal{X}) \cap \mathcal{F}(\mathbf{x}^*)$ and $\mathbf{s} \in \text{vert}(\mathcal{X}) \setminus \mathcal{F}(\mathbf{x}^*)$, and the Cauchy-Schwarz inequality, the second from L -smoothness and the fact that $\|\mathbf{s} - \mathbf{v}\| \leq D$, and the third inequality from m -strong convexity. Note that the last inequality comes from the fact that we assume that the primal gap satisfies $f(\mathbf{x}_k) - f(\mathbf{x}^*) < (\tau/(2D(\sqrt{L/m} + 1)))^2/(2L) \leq (\tau/(2LD))^2 m/2$. This allows us to claim that $\langle \mathbf{s}, \nabla f(\mathbf{x}_k) \rangle > \langle \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle$ for any $\mathbf{v} \in \text{vert}(\mathcal{X}) \cap \mathcal{F}(\mathbf{x}^*)$ and $\mathbf{s} \in \text{vert}(\mathcal{X}) \setminus \mathcal{F}(\mathbf{x}^*)$, which means that the Frank-Wolfe vertex in Line 3 satisfies that $\mathbf{v}_k \in \text{vert}(\mathcal{X}) \cap \mathcal{F}(\mathbf{x}^*)$. Alternatively, if we have that $\mathbf{x}_k \notin \mathcal{F}(\mathbf{x}^*)$ then $\mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$ is nonempty, which also means that the away-vertex chosen in Line 4 will be such that $\mathbf{s}_k \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$. The proof proceeds by showing that the AFW algorithm will chose to perform an away-step in Line 5 of Algorithm 5, as opposed to a Frank-Wolfe step. Let $\mathbf{s} \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$, using arguments that are similar to the ones in the previous chain of inequalities, then:

$$\begin{aligned} \langle \mathbf{s} - \mathbf{x}_k, \nabla f(\mathbf{x}_k) \rangle &= \langle \mathbf{s} - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{s} - \mathbf{x}^*, \nabla f(\mathbf{x}_k) - \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{x}^* - \mathbf{x}_k, \nabla f(\mathbf{x}_k) \rangle \\ &\geq \tau - \|\mathbf{s} - \mathbf{x}^*\| \|\nabla f(\mathbf{x}_k) - \nabla f(\mathbf{x}^*)\| - \max_{\mathbf{v} \in \mathcal{X}} \langle \mathbf{x}_k - \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle \\ &\geq \tau - LD \|\mathbf{x}_k - \mathbf{x}^*\| - \max_{\mathbf{v} \in \mathcal{X}} \langle \mathbf{x}_k - \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle \\ &\geq \tau - LD\sqrt{2(f(\mathbf{x}_k) - f(\mathbf{x}^*))}/m - D\sqrt{2L(f(\mathbf{x}_k) - f(\mathbf{x}^*))} \\ &= \tau - D\sqrt{2L(f(\mathbf{x}_k) - f(\mathbf{x}^*))} \left(\sqrt{L/m} + 1 \right) \\ &\geq \tau/2, \end{aligned} \tag{A.2}$$

where the first inequality stems from the strict complementarity assumption, which applies to $\mathbf{s} \in \text{vert}(\mathcal{X}) \setminus \mathcal{F}(\mathbf{x}^*)$, the Cauchy-Schwarz inequality, and the fact that $\langle \mathbf{x}_k - \mathbf{x}^*, \nabla f(\mathbf{x}_k) \rangle \leq \max_{\mathbf{v} \in \mathcal{X}} \langle \mathbf{x}_k - \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle$. The second inequality uses the fact that $\|\mathbf{s} - \mathbf{x}^*\| \leq D$ and L -smoothness. Note that from Theorem 2 in [Lacoste-Julien & Jaggi \(2015\)](#), we know that since the function is L -smooth then if we have that $f(\mathbf{x}_k) - f(\mathbf{x}^*) < LD^2/2$, then we also have that $\max_{\mathbf{v} \in \mathcal{X}} \langle \mathbf{x}_k - \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle \leq D\sqrt{2L(f(\mathbf{x}_k) - f(\mathbf{x}^*))}$. The third inequality uses the aforementioned bound, along with m -strong convexity for the term that contains $\|\mathbf{x}_k - \mathbf{x}^*\|$. The last inequality again uses the assumption that the primal gap satisfies $f(\mathbf{x}_k) - f(\mathbf{x}^*) < (\tau/(D(\sqrt{L/m} + 1)))^2/(8L)$. Moving on to the bound on the Frank-Wolfe gap, using the bound on Theorem 2 in [Lacoste-Julien & Jaggi \(2015\)](#) and our assumption on the primal gap allows us to conclude that:

$$\max_{\mathbf{v} \in \mathcal{X}} \langle \mathbf{x}_k - \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle \leq D\sqrt{2L(f(\mathbf{x}_k) - f(\mathbf{x}^*))} \leq \frac{\tau}{2(\sqrt{L/m} + 1)} \leq \frac{\tau}{2}, \tag{A.3}$$

where the last inequality simply stems from $L/m \geq 1$. Putting together Eqs. (A.2) and (A.3) we can conclude that for any $\mathbf{s} \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$ we have that $\langle \mathbf{s} - \mathbf{x}_k, \nabla f(\mathbf{x}_k) \rangle \geq \max_{\mathbf{v} \in \mathcal{X}} \langle \mathbf{x}_k - \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle$. Consequently, we have that $\max_{\mathbf{s} \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)} \langle \mathbf{s} - \mathbf{x}_k, \nabla f(\mathbf{x}_k) \rangle \geq \max_{\mathbf{v} \in \mathcal{X}} \langle \mathbf{x}_k - \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle$. This means that the AFW algorithm will perform an away step, moving away from a vertex $\mathbf{s}_k \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$. The last step of the proof proceeds to show that the step size chosen in the away-step corresponds to the largest possible step size λ_{\max} . To do this, assume for contradiction that this is not the case, and assume that $\lambda_k = \arg\min_{\lambda \in [0, \lambda_{\max}]} f(\mathbf{x} + \lambda(\mathbf{x}_k - \mathbf{s}_k)) < \lambda_{\max}$, by the optimality of the line search we

must have that $\langle \mathbf{s}_k - \mathbf{x}_k, \nabla f(\mathbf{x}_{k+1}) \rangle = 0$. On the other hand, we can write:

$$\begin{aligned} \langle \mathbf{s}_k - \mathbf{x}_k, \nabla f(\mathbf{x}_{k+1}) \rangle &= \langle \mathbf{s}_k - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{x}^* - \mathbf{x}_k, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{s}_k - \mathbf{x}_k, \nabla f(\mathbf{x}_{k+1}) - \nabla f(\mathbf{x}^*) \rangle, \\ &\geq \tau - (f(\mathbf{x}_k) - f(\mathbf{x}^*)) + \langle \mathbf{s}_k - \mathbf{x}_k, \nabla f(\mathbf{x}_{k+1}) - \nabla f(\mathbf{x}^*) \rangle \end{aligned} \quad (\text{A.4})$$

$$\begin{aligned} &> \tau/2 - \|\mathbf{s}_k - \mathbf{x}_k\| \|\nabla f(\mathbf{x}_{k+1}) - \nabla f(\mathbf{x}^*)\| \\ &\geq \tau/2 - LD \|\mathbf{x}_{k+1} - \mathbf{x}^*\| \\ &\geq \tau/2 - LD \sqrt{2(f(\mathbf{x}_k) - f(\mathbf{x}^*))}/m \\ &> 0, \end{aligned} \quad (\text{A.5})$$

where the inequality in Eq. (A.4) follows from the strict complementarity assumption, which means that $\langle \mathbf{s}_k - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq \tau$, and from convexity, which implies that $\langle \mathbf{x}^* - \mathbf{x}_k, \nabla f(\mathbf{x}^*) \rangle \geq -(f(\mathbf{x}_k) - f(\mathbf{x}^*))$. The following inequality, shown in Eq. (A.5), follows from the Cauchy-Schwarz inequality and from the bound on the primal gap $f(\mathbf{x}_k) - f(\mathbf{x}^*) < \tau/2$. The following inequalities follow from the application of the L -smoothness and m -strong convexity of the objective function and the bound on the primal gap $f(\mathbf{x}_k) - f(\mathbf{x}^*) < (\tau/(2D(\sqrt{L/m}+1)))^2/(2L) \leq (\tau/(2LD))^2 m/2$. This proves the desired contradiction, and so we must have that $\lambda_k = \lambda_{\max}$. This means that at iteration k we have performed an away step that has dropped a vertex $\mathbf{s}_k \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$. This proves the first claim.

The first part of the second claim follows by noting that if $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, then the away-vertex chosen in Line 4 will satisfy $\mathbf{s}_k \in \mathcal{F}(\mathbf{x}^*)$, as $\mathcal{S}_k \subseteq \mathcal{F}(\mathbf{x}^*)$. Moreover by Eq. (A.1) we know that the Frank-Wolfe vertex chosen in Line 3 will satisfy $\mathbf{v}_k \in \mathcal{F}(\mathbf{x}^*)$. This proves the first part of the second claim. Moving on to the second part of the second claim. If we denote the active set of \mathbf{x}_k at iteration k by \mathcal{S}_k it follows that $\mathcal{S}_k \subseteq \mathcal{F}(\mathbf{x}^*)$ from the fact that $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$ (as otherwise we would have $\mathbf{x}_k \notin \mathcal{F}(\mathbf{x}^*)$). Moreover from the first part of the second claim we have that $\mathbf{v}_k = \operatorname{argmin}_{\mathbf{u} \in \mathcal{X}} \langle \nabla f(\mathbf{x}_k), \mathbf{u} \rangle \in \mathcal{F}(\mathbf{x}^*)$ and $\mathbf{s}_k = \operatorname{argmin}_{\mathbf{u} \in \mathcal{S}_k} \langle -\nabla f(\mathbf{x}_k), \mathbf{u} \rangle \in \mathcal{F}(\mathbf{x}^*)$. From the definition of $w(\mathbf{x}_k, \mathcal{S}_k)$ we have that:

$$\begin{aligned} w(\mathbf{x}_k, \mathcal{S}_k) &= \langle \mathbf{s}_k - \mathbf{v}_k, \nabla f(\mathbf{x}_k) \rangle \\ &= \langle \mathbf{s}_k - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{x}^* - \mathbf{v}_k, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{s}_k - \mathbf{v}_k, \nabla f(\mathbf{x}_k) - \nabla f(\mathbf{x}^*) \rangle \\ &= \langle \mathbf{s}_k - \mathbf{v}_k, \nabla f(\mathbf{x}_k) - \nabla f(\mathbf{x}^*) \rangle \\ &\leq \|\mathbf{s}_k - \mathbf{v}_k\| \|\nabla f(\mathbf{x}_k) - \nabla f(\mathbf{x}^*)\| \\ &\leq LD \|\mathbf{x}_k - \mathbf{x}^*\| \\ &\leq \frac{LD\sqrt{2}}{\sqrt{m}} \sqrt{f(\mathbf{x}_k) - f(\mathbf{x}^*)}. \end{aligned}$$

Where the third equality follows from the strict complementarity assumption (Assumption 2.4), and the fact that as $\mathbf{s}_k, \mathbf{v}_k \in \mathcal{F}(\mathbf{x}^*)$ we have that $\langle \mathbf{x}^* - \mathbf{v}_k, \nabla f(\mathbf{x}^*) \rangle = 0$ and $\langle \mathbf{s}_k - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle = 0$. The first inequality follows from the Cauchy-Schwarz inequality, the second from the L -smoothness of $f(x)$ and $\|\mathbf{s}_k - \mathbf{v}_k\| \leq D$, and the last one from the m -strong convexity of the $f(x)$. This completes the proof. \square

Corollary A.6. *If the strict complementarity assumption is satisfied (Assumption 2.4) and the primal gap satisfies $f(\mathbf{x}_k) - f(\mathbf{x}^*) < 1/2 \min \left\{ (\tau/(2D(\sqrt{L/m}+1)))^2/L, \tau, LD^2 \right\}$ then the claims in Theorem A.5 hold for the Pairwise-Step Frank-Wolfe (PFW) algorithm, the natural extension of the AFW algorithm to the use of pairwise-steps.*

Proof. The proof uses many of the same techniques and concepts as Theorem A.5, so we only give a brief proof-sketch here. Note that chain of inequalities in Eq. (A.1) is independent of the algorithm being used. Therefore if the primal gap satisfies $f(\mathbf{x}_k) - f(\mathbf{x}^*) < (\tau/(2D(\sqrt{L/m}+1)))^2/(2L) \leq (\tau/(2LD))^2 m/2$, then we can claim that $\langle \mathbf{s}, \nabla f(\mathbf{x}_k) \rangle > \langle \mathbf{v}, \nabla f(\mathbf{x}_k) \rangle$ for any $\mathbf{v} \in \operatorname{vert}(\mathcal{X}) \cap \mathcal{F}(\mathbf{x}^*)$ and $\mathbf{s} \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$, which means that the Frank-Wolfe vertex in the PFW algorithm satisfies that $\mathbf{v}_k \in \operatorname{vert}(\mathcal{X}) \cap \mathcal{F}(\mathbf{x}^*)$. Alternatively, this also means that the away-vertex chosen in the PFW algorithm will be such that $\mathbf{s}_k \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$. Note that the directions towards which the iterates move in the PFW algorithm are given by $\mathbf{s}_k - \mathbf{x}_k$, with a maximum step size of $\lambda_{\max} = \alpha_k^{\mathbf{s}_k}$. If the maximum step size is chosen this means that the vertex $\mathbf{s}_k \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$ has been dropped, and so to proceed we will show that this maximum step size is chosen. To do this, assume for contradiction that this is not the case, and assume that $\lambda_k = \operatorname{argmin}_{\lambda \in [0, \lambda_{\max}]} f(\mathbf{x} + \lambda(\mathbf{x}_k - \mathbf{s}_k)) < \lambda_{\max}$, by

the optimality of the line search we must have that $\langle \mathbf{s}_k - \mathbf{x}_k, \nabla f(\mathbf{x}_{k+1}) \rangle = 0$. On the other hand, we can write:

$$\begin{aligned}
 \langle \mathbf{s}_k - \mathbf{v}_k, \nabla f(\mathbf{x}_{k+1}) \rangle &= \langle \mathbf{s}_k - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{x}^* - \mathbf{v}_k, \nabla f(\mathbf{x}^*) \rangle + \langle \mathbf{s}_k - \mathbf{v}_k, \nabla f(\mathbf{x}_{k+1}) - \nabla f(\mathbf{x}^*) \rangle, \\
 &> \tau - \|\mathbf{s}_k - \mathbf{v}_k\| \|\nabla f(\mathbf{x}_{k+1}) - \nabla f(\mathbf{x}^*)\| \\
 &\geq \tau - LD \|\mathbf{x}_{k+1} - \mathbf{x}^*\| \\
 &\geq \tau - LD \sqrt{2(f(\mathbf{x}_k) - f(\mathbf{x}^*)) / m} \\
 &> 0,
 \end{aligned} \tag{A.6}$$

where the inequality in Eq. (A.6) follows from the Cauchy-Schwarz inequality and the strict complementarity assumption for \mathbf{s}_k and \mathbf{v}_k , which means that $\langle \mathbf{s}_k - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle \geq \tau$ and $\langle \mathbf{v}_k - \mathbf{x}^*, \nabla f(\mathbf{x}^*) \rangle = 0$. The following inequalities follow from the application of the L -smoothness and m -strong convexity of the objective function and the bound on the primal gap $f(\mathbf{x}_k) - f(\mathbf{x}^*) < (\tau / (2D(\sqrt{L/m} + 1)))^2 / (2L) \leq (\tau / (2LD))^2 m / 2$. This proves the desired contradiction, and so we must have that $\lambda_k = \lambda_{\max}$. This means that at iteration k we have dropped a vertex $\mathbf{s}_k \in \mathcal{S}_k \setminus \mathcal{F}(\mathbf{x}^*)$ from the active set \mathcal{S}_k . This proves the first claim. The proof of the second claim and the bound on $w(\mathbf{x}_k, \mathcal{S}_k)$ can be repeated word for word from Theorem A.5. \square

Using Theorem A.5 (or Corollary A.6 for the Pairwise-Step Frank-Wolfe (PFW)), and assuming that $\mathbf{x}_0 \in \text{vert}(\mathcal{X})$, and using the primal gap convergence gap guarantee in Lacoste-Julien & Jaggi (2015, Theorem 1), we can bound the number of iterations until $f(\mathbf{x}_k) - f(\mathbf{x}^*)$ satisfies the requirement in Theorem A.5. Using this bound, and the fact that the AFW algorithm can pick up at most one vertex per iteration, we can bound the number of iterations until $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$ (there exist other bounds for the PFW algorithm). Note that by the second claim in Theorem 3.1, this means that when $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, then the iterates will not leave $\mathcal{F}(\mathbf{x}^*)$. Furthermore, once the iterates are inside the optimal face, there are two options: if $\mathbf{x}^* \in \mathcal{F}(\mathbf{x}^*)$, then the AFW algorithm will exit once $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, as $w(\mathbf{x}_k, \mathcal{S}_k) = 0$, otherwise if $\mathbf{x}^* \notin \text{vert}(\mathcal{X})$ (the case of interest in our setting, by Assumption 2.5), then we need to prove that after a given number of iterations the active set will satisfy $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$. We prove the former using Fact A.7 (a variation of Diakonikolas et al. (2020, Fact B.3)).

Fact A.7 (Critical strong Wolfe gap). There exists an $w_c > 0$ such that for any subset $\mathcal{S} \subseteq \text{vert}(\mathcal{F}(\mathbf{x}^*))$ and point $\mathbf{x} \in \mathcal{F}(\mathbf{x}^*)$ with $\mathbf{x} \in \text{co}(\mathcal{S})$ and $w(\mathbf{x}, \mathcal{S}) \leq w_c$ it follows that $\mathbf{x}^* \in \text{co}(\mathcal{S})$.

Remark A.8. The critical strong Wolfe gap in Fact A.7, is a crucial parameter in the coming proofs. However, like the strict complementarity parameter τ (Guélat & Marcotte, 1986; Garber, 2020) and the critical radius defined in (Diakonikolas et al., 2020), the critical strong Wolfe gap can be arbitrarily small. Fortunately, as we will show in the proofs to come, it only affects the length of the burn-in phase of the accelerated algorithm, and moreover this dependence is logarithmically. Below, we sketch a simple example for which one can compute an exact expression for w_c . We then give an upper and a lower bound on w_c for any general problem of the form shown in Problem (P).

Consider minimizing $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|^2$ over the unit probability simplex, which we denote by Δ^n . The function is smooth and strongly convex, and the minimizer of this function over the feasible region is given by $\mathbf{x}^* = \mathbf{1}/n$, where n is the dimensionality of the problem and $\mathcal{F}(\mathbf{x}^*) = \Delta^n$. We can compute the exact expression for w_c as:

$$w_c = \inf_{\substack{\mathcal{S} \subseteq \text{vert}(\Delta^n) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \max_{\substack{\mathbf{u} \in \mathcal{S} \\ \mathbf{v} \in \Delta^n}} \langle \nabla f(\mathbf{x}), \mathbf{u} - \mathbf{v} \rangle \tag{A.7}$$

$$= \inf_{\substack{\mathcal{S} \subseteq \text{vert}(\Delta^n) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \left(\max_{\mathbf{u} \in \mathcal{S}} \langle \mathbf{x}, \mathbf{u} \rangle - \min_{\mathbf{v} \in \Delta^n} \langle \mathbf{x}, \mathbf{v} \rangle \right) \tag{A.8}$$

$$= \inf_{\substack{\mathcal{S} \subseteq \text{vert}(\Delta^n) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \max_{\mathbf{u} \in \mathcal{S}} \langle \mathbf{x}, \mathbf{u} \rangle \tag{A.9}$$

$$= \inf_{\substack{\mathcal{S} \subseteq \text{vert}(\Delta^n) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \|\mathbf{x}\|_\infty \tag{A.10}$$

$$= \frac{1}{n-1}. \tag{A.11}$$

Where the equality in Equation (A.8) follows from $\nabla f(\mathbf{x}) = \mathbf{x}$ and Equation (A.9) from the fact that as $\mathcal{S} \subset \text{vert}(\Delta^n)$, $\mathbf{x}^* \notin \text{co}(\mathcal{S})$ and $\mathbf{x} \in \text{co}(\mathcal{S})$, then one or more of the components of \mathbf{x} are zero, and as the rest are positive, it follows that $\min_{\mathbf{v} \in \Delta^n} \langle \mathbf{x}, \mathbf{v} \rangle = 0$. The equality in Equation (A.10) follows from the fact that the components of \mathbf{x} are non-negative and the elements of \mathcal{S} are the standard orthogonal basis vectors, and so $\max_{\mathbf{u} \in \mathcal{S}} \langle \mathbf{x}, \mathbf{u} \rangle = \|\mathbf{x}\|_\infty$. The last equality simply follows from the fact that the infimum is achieved for the set \mathcal{S} of largest cardinality that does not contain \mathbf{x}^* , and for a point $\mathbf{x}^* \in \text{co}(\mathcal{S})$ that minimizes the infinity norm. We can now prove a lower bound on w_c for the general case as follows:

$$w_c = \inf_{\substack{\mathcal{S} \subset \text{vert}(\mathcal{F}(\mathbf{x}^*)) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} w(\mathbf{x}, \mathcal{S}) \geq \inf_{\substack{\mathcal{S} \subset \text{vert}(\mathcal{F}(\mathbf{x}^*)) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \frac{m}{2} \|\mathbf{x}^* - \mathbf{x}\|^2 \geq \inf_{\substack{\mathcal{S} \subset \text{vert}(\mathcal{X}) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \frac{m}{2} \|\mathbf{x}^* - \mathbf{x}\|^2 = \frac{\mu}{2} r_c^2.$$

The first inequality stems from the fact as the function is m -strongly convex then $w(\mathbf{x}, \mathcal{S}) \geq f(\mathbf{x}) - f(\mathbf{x}^*) \geq \frac{m}{2} \|\mathbf{x} - \mathbf{x}^*\|^2$. The second inequality stems from the fact that $\text{vert}(\mathcal{F}(\mathbf{x}^*)) \subset \text{vert}(\mathcal{X})$. The last equality comes from the definition of r_c in (Diakonikolas et al., 2020). Regarding the upper bound, for convenience we denote away-vertex and the Frank-Wolfe vertex as $\mathbf{u} = \arg\max_{\mathbf{y} \in \mathcal{S}} \langle \nabla f(\mathbf{x}), \mathbf{y} \rangle$ and $\mathbf{v} = \arg\min_{\mathbf{y} \in \mathcal{X}} \langle \nabla f(\mathbf{x}), \mathbf{y} \rangle$, respectively. This allows us to write:

$$\begin{aligned} w_c &= \inf_{\substack{\mathcal{S} \subset \text{vert}(\mathcal{F}(\mathbf{x}^*)) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \langle \nabla f(\mathbf{x}^*), \mathbf{u} - \mathbf{x}^* \rangle + \langle \nabla f(\mathbf{x}^*), \mathbf{x}^* - \mathbf{v} \rangle + \langle \nabla f(\mathbf{x}) - \nabla f(\mathbf{x}^*), \mathbf{u} - \mathbf{v} \rangle \\ &= \inf_{\substack{\mathcal{S} \subset \text{vert}(\mathcal{F}(\mathbf{x}^*)) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \langle \nabla f(\mathbf{x}) - \nabla f(\mathbf{x}^*), \mathbf{u} - \mathbf{v} \rangle \\ &\leq \inf_{\substack{\mathcal{S} \subset \text{vert}(\mathcal{F}(\mathbf{x}^*)) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \|\nabla f(\mathbf{x}) - \nabla f(\mathbf{x}^*)\| \|\mathbf{u} - \mathbf{v}\| \\ &\leq LD \inf_{\substack{\mathcal{S} \subset \text{vert}(\mathcal{F}(\mathbf{x}^*)) \\ \mathbf{x} \in \text{co}(\mathcal{S}) \\ \mathbf{x}^* \notin \text{co}(\mathcal{S})}} \|\mathbf{x} - \mathbf{x}^*\|. \end{aligned}$$

Where the first equality is simply obtained from rewriting $\langle \nabla f(\mathbf{x}), \mathbf{u} - \mathbf{v} \rangle$, and the second inequality comes from the fact that $\langle \nabla f(\mathbf{x}^*), \mathbf{u} - \mathbf{x}^* \rangle = 0$ from the strict complementarity assumption and $\langle \nabla f(\mathbf{x}^*), \mathbf{v} - \mathbf{x}^* \rangle \geq 0$ from the optimality conditions for \mathbf{x}^* . The first inequality in the previous chain comes from the application of the Cauchy-Schwarz inequality, and the second inequality from the L -smoothness of f . To see why this upper bound means that w_c can be arbitrarily small for certain problems, consider the setup in Figure 3.




Figure 3. Example problem setup.

The figure shows a two dimensional example where $\mathcal{F}(\mathbf{x}^*) = \mathcal{X} = \text{co}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4)$. Moreover, the set $\mathcal{S} = \text{co}(\mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4)$ shown in light gray is contained in the optimal face, and does not contain \mathbf{x}^* (which is on the diagonal defined by \mathbf{v}_1 and \mathbf{v}_3 and just off the diagonal defined by \mathbf{v}_2 and \mathbf{v}_4). The distance from \mathcal{S} to \mathbf{x}^* is precisely $\inf_{\mathcal{S} \subset \text{vert}(\mathcal{F}(\mathbf{x}^*)), \mathbf{x} \in \text{co}(\mathcal{S}), \mathbf{x}^* \notin \text{co}(\mathcal{S})} \|\mathbf{x} - \mathbf{x}^*\|$, and it can be made arbitrarily small as \mathbf{x}^* approaches the diagonal defined by \mathbf{v}_2 and \mathbf{v}_4 along the diagonal defined by \mathbf{v}_1 and \mathbf{v}_3 , which means that w_c can become arbitrarily small.

With these tools, we are ready to prove the desired convergence bound.

Theorem A.9. *Assume the AFW algorithm (Algorithm 5) is run starting with $\mathbf{x}_0 \in \text{vert}(\mathcal{X})$. If the strict complementarity assumption (Assumption 2.4) is satisfied and $\mathbf{x}^* \notin \text{vert}(\mathcal{X})$, then for $k \geq K_0$ with*

$$K_0 = \frac{32LD^2}{m\delta(\mathcal{X})^2 \ln 2} \log \left(\frac{2w(\mathbf{x}_0, \mathcal{S}_0)}{\min\{(\tau/(2D(\sqrt{L/m} + 1)))^2/L, \tau, LD^2, 2w_c\}} \right),$$

we have that $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$ and moreover:

$$w(\mathbf{x}_k, \mathcal{S}_k) \leq LD\sqrt{2(f(\mathbf{x}_k) - f(\mathbf{x}^*))}/m.$$

Where $\delta(\mathcal{X})$ is the pyramidal width in Definition 2.2, and $w_c > 0$ is the critical strong Wolfe gap in Fact 3.2.

Proof. Using the primal gap bound in Lacoste-Julien & Jaggi (2015, Theorem 1), and assuming that $\mathbf{x}_0 \in \text{vert}(\mathcal{X})$, then for iterations

$$k \geq \frac{8LD^2}{m\delta(\mathcal{X})^2} \log \left(\frac{2w(\mathbf{x}_0, \mathcal{S}_0)}{\min\{(\tau/(2D(\sqrt{L/m} + 1)))^2/L, \tau, LD^2\}} \right),$$

we know that

$$f(\mathbf{x}_k) - f(\mathbf{x}^*) < \frac{1}{2} \min \left\{ (\tau/(2D(\sqrt{L/m} + 1)))^2/L, \tau, LD^2 \right\}.$$

Using the first claim in Theorem 3.1, as the algorithm will at most have picked up one vertex of the polytope \mathcal{X} per iteration, then the AFW algorithm will at most need to drop a number of vertices from its active set equal to the number of iterations that have elapsed before $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$. This means that for

$$k \geq \frac{16LD^2}{m\delta(\mathcal{X})^2} \log \left(\frac{2w(\mathbf{x}_0, \mathcal{S}_0)}{\min\{(\tau/(2D(\sqrt{L/m} + 1)))^2/L, \tau, LD^2\}} \right)$$

we know that $\mathbf{x}_k \in \mathcal{F}(\mathbf{x}^*)$, and, moreover, we have that $w(\mathbf{x}_k, \mathcal{S}_k) \leq LD\sqrt{2/m}\sqrt{f(\mathbf{x}_k) - f(\mathbf{x}^*)}$ by the second claim in Theorem 3.1. It remains to bound the number of iterations until $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$ in the AFW algorithm. Using Fact 3.2, and the convergence guarantees in Kerdreux et al. (2019), we know that it suffices to run $32LD^2/(m\delta(\mathcal{X})^2 \ln 2) \log(w(\mathbf{x}_0, \mathcal{S}_0)/w_c)$ iterations of the AFW algorithm for $w(\mathbf{x}_k, \mathcal{S}_k) \leq w_c$. Consequently, we have that for $k \geq K_0$ it holds $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$. \square

B. Acceleration Analysis

Before delving into the analysis of the accelerated algorithm used in our work, we review some useful properties of gradient mapping and define its inexact variant.

B.1. Gradient Mapping

One of the properties that will be useful for our analysis is summarized in the following two propositions.

Proposition B.1. *Let $\mathcal{C} \subseteq \mathbb{R}^n$ and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be differentiable and m -strongly convex on \mathcal{C} , for some $m > 0$. Then, $\forall \mathbf{x} \in \mathcal{C}, \forall m' \geq m$:*

$$f(\mathbf{x}) - \min_{\mathbf{u} \in \mathcal{C}} f(\mathbf{u}) \leq \frac{1}{2m} \|G_m(\mathbf{x})\|^2 \leq \frac{1}{2m'} \|G_{m'}(\mathbf{x})\|^2.$$

Proof. Fix any $\mathbf{x} \in \mathcal{C}$. The first inequality in the proposition follows by using strong convexity of f , which states that for all $\mathbf{u} \in \mathcal{C}$ we have that $f(\mathbf{u}) - f(\mathbf{x}) \geq \langle \nabla f(\mathbf{x}), \mathbf{u} - \mathbf{x} \rangle + \frac{m}{2} \|\mathbf{u} - \mathbf{x}\|^2$, and minimizing both sides over $\mathbf{u} \in \mathcal{C}$. The second inequality is a consequence of the monotonicity of gradient mapping (see Beck (2017, Theorem 10.9)). \square

Proposition B.2. *Let $\mathcal{C} \subseteq \mathbb{R}^n$ and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be L -smooth on \mathcal{C} , for some $L < \infty$. Then, $\forall \mathbf{x} \in \mathcal{C}, \forall \eta \geq L$:*

$$f(\mathbf{x}) - \min_{\mathbf{u} \in \mathcal{C}} f(\mathbf{u}) \geq \frac{1}{2\eta} \|G_\eta(\mathbf{x})\|^2.$$

The proof is omitted and can instead be found in, e.g., [Beck \(2017, Chapter 10\)](#).

To construct an implementable version of our algorithm, we rely on *inexact* projections, which in turn lead to *inexact* evaluations of the gradient mapping, defined below. This is the primary reason why we cannot rely on black box applications of any of the existing results, and instead require a variant of a parameter-free accelerated method that can work with inexact evaluations of the gradient mapping.

Definition B.3 (Inexact gradient mapping). Let \mathcal{C} be a closed convex set and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be smooth and convex on \mathcal{C} . For any $\mathbf{x} \in \mathcal{C}$ and $\eta > 0$, we say that $\tilde{G}_\eta(\mathbf{x}) = \eta(\mathbf{x} - \mathbf{y})$ is an ϵ^ℓ -approximate gradient mapping at \mathbf{x} for some $\mathbf{x} \in \mathcal{C}$, if $\mathbf{y} \in \mathcal{C}$ and $\ell(\mathbf{y}) \leq \min_{\mathbf{u} \in \mathcal{C}} \ell(\mathbf{u}) + \epsilon^\ell$, where we define $\ell(\mathbf{u}) = \langle \nabla f(\mathbf{x}), \mathbf{u} - \mathbf{x} \rangle + \frac{\eta}{2} \|\mathbf{u} - \mathbf{x}\|_2^2$.

B.2. Proof of Lemma 3.5

In this section, we prove Lemma 3.5. To make it easier to follow the argument, we restate the lemma below. Much of the proof closely follows the proof of Lemma 3.2 in [\(Diakonikolas et al., 2020\)](#). However, there are differences in both the algorithm and the proof itself required to obtain the bound for the gradient mapping, and we fill in the additional arguments that are needed to prove the lemma below.

Lemma 3.5. Let $\mathcal{C} \subseteq \mathbb{R}^n$ be a closed convex set and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be an L -smooth function on \mathcal{C} . Let $\mathbf{x}_0 \in \mathcal{C}$ be an arbitrary initial point, and, given $\sigma > 0$, define $f_\sigma(\mathbf{x}) = f(\mathbf{x}) + \frac{\sigma}{2} \|\mathbf{x} - \mathbf{x}_0\|^2$, $\mathbf{x}_\sigma^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{C}} f_\sigma(\mathbf{x})$. Let $\mathbf{z}_0 = (\eta_0 + \sigma)\mathbf{x}_0 - \nabla f(\mathbf{x}_0)$, $\mathbf{y}_0 = \mathbf{v}_0 = \hat{\mathbf{y}}_0 \stackrel{\epsilon_0^M}{\sim} \operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} M_0(\mathbf{u})$, where $\epsilon_0^M > 0$, $M_0(\mathbf{u})$ is defined as in Algorithm 2, and the estimate η_0 is doubled until $f(\mathbf{y}_0) \leq f(\mathbf{x}_0) + \langle \nabla f(\mathbf{x}_0), \mathbf{y}_0 - \mathbf{x}_0 \rangle + \frac{\eta_0}{2} \|\mathbf{y}_0 - \mathbf{x}_0\|^2$, same as in Algorithm 2. Given $\eta_0 > 0$, sequence $\{a_k\}_{k \geq 0}$, $A_k = \sum_{i=0}^k a_i$, $\theta_k = \frac{a_k}{A_k}$, and the sequences of errors $\{\epsilon_k^M\}_{k \geq 0}$, $\{\epsilon_k^\ell\}_{k \geq 0}$, let the sequences of points $\{\hat{\mathbf{y}}_k, \mathbf{y}_k\}_{k \geq 0}$ evolve according to Algorithm 2 for $k \geq 1$.

If $a_0 = A_0 = 1$ and $\theta_k = \frac{a_k}{A_k} \leq \sqrt{\frac{\sigma}{2(\eta_k + \sigma)}}$ for $k \geq 1$, then for all $k \geq 1$

$$\frac{\|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2}{\eta_k + \sigma} \leq \frac{\eta_0}{A_k} \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2 + \frac{2}{A_k} \sum_{i=0}^k (2\epsilon_i^M + \epsilon_i^\ell),$$

where $G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)$ is the gradient mapping w.r.t. f_σ , at $\hat{\mathbf{y}}_k$.

In particular, if $a_0 = A_0 = 1$, $\theta_k = \frac{a_k}{A_k} = \sqrt{\frac{\sigma}{2(\eta_k + \sigma)}}$ for $k \geq 1$, $\epsilon_k^M = \frac{a_k \epsilon^2}{8}$, and $\epsilon_k^\ell \leq \frac{a_k \epsilon^2}{8}$, then $\frac{1}{\eta_k + \sigma} \|G_{\eta_k}^\sigma(\hat{\mathbf{y}}_k)\|^2 \leq \epsilon^2$ after at most

$$k = O\left(\sqrt{\frac{L}{\sigma}} \log\left(\frac{L \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|}{\epsilon}\right)\right)$$

iterations. Further, the total number of first-order queries to f and oracle queries to inexact projections is at most

$$k' = k + 2 \log\left(\frac{2L}{\eta_0}\right).$$

Proof. The proof follows the Approximate Duality Gap Technique [\(Diakonikolas & Orecchia, 2019\)](#), similar to [\(Diakonikolas et al., 2020\)](#). The main idea is to construct an approximation of the true optimality gap $f_\sigma(\mathbf{y}_k) - f_\sigma(\mathbf{x}_\sigma^*)$, where $\mathbf{x}_\sigma^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{C}} f_\sigma(\mathbf{x})$ and show that it contracts at rate $1/A_k$, where A_k is made as fast growing as possible.

The approximate gap Γ_k is constructed as $\Gamma_k = \Upsilon_k - \Lambda_k$, where $\Upsilon_k = f_\sigma(\mathbf{y}_k)$ and Λ_k is defined by [\(Diakonikolas et al., 2020\)](#)

$$\Lambda_k = \frac{\sum_{i=0}^k a_i f_\sigma(\mathbf{x}_i) + \min_{\mathbf{u} \in \mathcal{C}} m_k(\mathbf{u}) - \frac{\eta_0 \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2}{2}}{A_k}, \quad (\text{B.1})$$

$$m_k(\mathbf{u}) \stackrel{\text{def}}{=} \sum_{i=0}^k a_i \langle \nabla f_\sigma(\mathbf{x}_i), \mathbf{u} - \mathbf{x}_i \rangle + \sum_{i=0}^k a_i \frac{\sigma}{2} \|\mathbf{u} - \mathbf{x}_i\|^2 + \frac{\eta_0}{2} \|\mathbf{u} - \mathbf{x}_0\|^2.$$

By construction, $\Lambda_k \leq f_\sigma(\mathbf{x}_\sigma^*) [\(Diakonikolas et al., 2020\)](#), and, thus, $f_\sigma(\mathbf{x}_k) - f_\sigma(\mathbf{x}_\sigma^*) \leq \Gamma_k$. Further, $\operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} m_k(\mathbf{u}) = \operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} M_k(\mathbf{u})$.$

Let $\mathbf{v}_k^* = \operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} m_k(\mathbf{u})$, and observe that, by construction, $m_k(\mathbf{v}_k) - m_k(\mathbf{v}_k^*) \leq \epsilon_k^M$. Following the same argument as in (Diakonikolas et al., 2020), we have that

$$A_0 \Gamma_0 \leq \frac{\eta_0 \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2}{2} + \epsilon_0^M \quad (\text{B.2})$$

and

$$\begin{aligned} m_k(\mathbf{v}_k^*) - m_{k-1}(\mathbf{v}_{k-1}^*) \\ \geq a_k \langle \nabla f_\sigma(\mathbf{x}_k), \mathbf{v}_k - \mathbf{x}_k \rangle + \frac{\sigma A_k}{4} \|\mathbf{v}_k - (1 - \theta_k) \mathbf{v}_{k-1} - \theta_k \mathbf{x}_k\|^2 - \epsilon_k^M - \epsilon_{k-1}^M. \end{aligned} \quad (\text{B.3})$$

From Algorithm 2, $\hat{\mathbf{y}}_k - \mathbf{x}_k = \theta_k(\mathbf{v}_k - (1 - \theta_k) \mathbf{v}_{k-1} - \theta_k \mathbf{x}_k)$, and, thus, combining with Eq. (B.3) and the definition of the lower bound in Eq. (B.1), it follows that:

$$A_k \Lambda_k - A_{k-1} \Lambda_{k-1} \geq a_k f_\sigma(\mathbf{x}_k) + a_k \langle \nabla f_\sigma(\mathbf{x}_k), \mathbf{v}_k - \mathbf{x}_k \rangle + \frac{\sigma A_k}{4\theta_k^2} \|\hat{\mathbf{y}}_k - \mathbf{x}_k\|^2 - \epsilon_k^M - \epsilon_{k-1}^M. \quad (\text{B.4})$$

On the other hand, from $\Upsilon_k = f(\mathbf{y}_k)$, we have

$$\begin{aligned} A_k \Upsilon_k - A_{k-1} \Upsilon_{k-1} &= A_k f_\sigma(\mathbf{y}_k) - A_{k-1} f_\sigma(\mathbf{y}_{k-1}) \\ &= A_k (f_\sigma(\mathbf{y}_k) - f_\sigma(\hat{\mathbf{y}}_k)) + a_k f_\sigma(\mathbf{x}_k) \\ &\quad + A_k (f_\sigma(\hat{\mathbf{y}}_k) - f_\sigma(\mathbf{x}_k)) + A_{k-1} (f_\sigma(\mathbf{x}_k) - f_\sigma(\mathbf{y}_{k-1})). \end{aligned} \quad (\text{B.5})$$

By the exit criterion of the repeat-until loop in Algorithm 2, we have $f_\sigma(\hat{\mathbf{y}}_k) - f_\sigma(\mathbf{x}_k) \leq \langle \nabla f_\sigma(\mathbf{x}_k), \hat{\mathbf{y}}_k - \mathbf{x}_k \rangle + \frac{\eta_k + \sigma}{2} \|\hat{\mathbf{y}}_k - \mathbf{x}_k\|^2$, while by convexity of f_σ , $f_\sigma(\mathbf{x}_k) - f_\sigma(\mathbf{y}_{k-1}) \leq \langle \nabla f_\sigma(\mathbf{x}_k), \mathbf{x}_k - \mathbf{y}_{k-1} \rangle$. Plugging into Eq. (B.5), we have

$$\begin{aligned} A_k \Upsilon_k - A_{k-1} \Upsilon_{k-1} &\leq A_k (f_\sigma(\mathbf{y}_k) - f_\sigma(\hat{\mathbf{y}}_k)) + a_k f_\sigma(\mathbf{x}_k) \\ &\quad + \langle \nabla f_\sigma(\mathbf{x}_k), A_k \hat{\mathbf{y}}_k - A_{k-1} \mathbf{y}_{k-1} - a_k \mathbf{x}_k \rangle + \frac{A_k (\eta_k + \sigma)}{2} \|\hat{\mathbf{y}}_k - \mathbf{x}_k\|^2. \end{aligned}$$

By definition of $\hat{\mathbf{y}}_k$, $A_k \hat{\mathbf{y}}_k - A_{k-1} \mathbf{y}_{k-1} - a_k \mathbf{x}_k = a_k(\mathbf{v}_k - \mathbf{x}_k)$. Thus, combining the last inequality with Eq. (B.4)

$$\begin{aligned} A_k \Gamma_k - A_{k-1} \Gamma_{k-1} &\leq A_k (f_\sigma(\mathbf{y}_k) - f_\sigma(\hat{\mathbf{y}}_k)) + \frac{A_k}{4} \left(2(\eta_k + \sigma) - \frac{\sigma}{\theta_k^2} \right) \|\hat{\mathbf{y}}_k - \mathbf{x}_k\|^2 + \epsilon_k^M + \epsilon_{k-1}^M \\ &\leq A_k (f_\sigma(\mathbf{y}_k) - f_\sigma(\hat{\mathbf{y}}_k)) + \epsilon_k^M + \epsilon_{k-1}^M, \end{aligned}$$

$$\text{as } \theta_k = \frac{a_k}{A_k} \leq \sqrt{\frac{\sigma}{2(\eta_k + \sigma)}}.$$

To bound $f_\sigma(\mathbf{y}_k) - f_\sigma(\hat{\mathbf{y}}_k)$, observe that, by the exit condition in the repeat-until loop in Algorithm 2, we have

$$f_\sigma(\mathbf{y}_k) - f_\sigma(\hat{\mathbf{y}}_k) \leq \langle \nabla f_\sigma(\hat{\mathbf{y}}_k), \mathbf{y}_k - \hat{\mathbf{y}}_k \rangle + \frac{\eta_k + \sigma}{2} \|\mathbf{y}_k - \hat{\mathbf{y}}_k\|^2 = \ell_k(\mathbf{y}_k). \quad (\text{B.6})$$

Let $\bar{\mathbf{y}}_k = \operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} \ell_k(\mathbf{u})$, and let us argue now that $\ell_k(\bar{\mathbf{y}}_k) \leq -\frac{1}{2(\eta_k + \sigma)} \|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2$. This can be proved using the definitions of $G_{\eta_k}^\sigma(\hat{\mathbf{y}}_k)$ and $\bar{\mathbf{y}}_k$. In particular, this property follows from

$$\begin{aligned} \ell_k(\bar{\mathbf{y}}_k) &= \langle \nabla f_\sigma(\hat{\mathbf{y}}_k), \bar{\mathbf{y}}_k - \hat{\mathbf{y}}_k \rangle + \frac{\eta_k + \sigma}{2} \|\bar{\mathbf{y}}_k - \hat{\mathbf{y}}_k\|^2 \\ &= \langle \nabla f_\sigma(\hat{\mathbf{y}}_k) - G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k), \bar{\mathbf{y}}_k - \hat{\mathbf{y}}_k \rangle - \frac{1}{2(\eta_k + \sigma)} \|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2 \end{aligned}$$

and

$$\langle \nabla f_\sigma(\hat{\mathbf{y}}_k) - G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k), \bar{\mathbf{y}}_k - \hat{\mathbf{y}}_k \rangle = \langle \nabla f_\sigma(\hat{\mathbf{y}}_k) + (\eta_k + \sigma)(\bar{\mathbf{y}}_k - \hat{\mathbf{y}}_k), \bar{\mathbf{y}}_k - \hat{\mathbf{y}}_k \rangle = \langle \nabla \ell_k(\bar{\mathbf{y}}_k), \bar{\mathbf{y}}_k - \hat{\mathbf{y}}_k \rangle \leq 0,$$

where the last inequality follows from the first-order optimality of $\bar{\mathbf{y}}_k = \operatorname{argmin}_{\mathbf{y} \in \mathcal{C}} \ell_k(\mathbf{y})$.

Combining with Eq. (B.6), and using that $\ell_k(\mathbf{y}_k) - \ell_k(\bar{\mathbf{y}}_k) \leq \epsilon_k^\ell$ (by definition of \mathbf{y}_k), we have

$$f_\sigma(\mathbf{y}_k) - f_\sigma(\hat{\mathbf{y}}_k) \leq -\frac{1}{2(\eta_k + \sigma)} \|G_{\eta_k}^\sigma(\hat{\mathbf{y}}_k)\|^2 + \epsilon_k^\ell.$$

Thus, it follows that

$$A_k \Gamma_k - A_{k-1} \Gamma_{k-1} \leq -\frac{A_k}{2(\eta_k + \sigma)} \|G_{\eta_k}^\sigma(\hat{\mathbf{y}}_k)\|^2 + A_k \epsilon_k^\ell + \epsilon_k^M + \epsilon_{k-1}^M. \quad (\text{B.7})$$

Telescoping Eq. (B.7) and using the bound on the initial gap from Eq. (B.2), we have

$$\begin{aligned} A_k \Gamma_k &\leq A_0 \Gamma_0 - \sum_{i=1}^k \frac{A_i}{2(\eta_i + \sigma)} \|G_{\eta_i + \sigma}^\sigma(\hat{\mathbf{y}}_i)\|^2 + \sum_{i=0}^k (2\epsilon_i^M + A_i \epsilon_i^\ell) \\ &\leq -\frac{A_k}{2(\eta_k + \sigma)} \|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2 + \frac{\eta_0 \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2}{2} + \sum_{i=0}^k (2\epsilon_i^M + A_i \epsilon_i^\ell). \end{aligned}$$

As $\Gamma_k \geq f_\sigma(\mathbf{y}_k) - f_\sigma(\mathbf{x}_C^*) \geq 0$, we finally have

$$\frac{1}{\eta_k + \sigma} \|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2 \leq \frac{2}{A_k} \left(\frac{\eta_0 \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2}{2} + \sum_{i=0}^k (2\epsilon_i^M + A_i \epsilon_i^\ell) \right), \quad (\text{B.8})$$

which gives the first part of the lemma.

For the remaining part, as f is L -smooth, it follows that the condition from the repeat-until loop is satisfied for any $\eta_k \geq L$. As η_k gets doubled each time the condition is not satisfied, we have $\eta_k \leq 2L, \forall k$. Thus the total number of times the repeat-until loop is entered is at most k (the total number of iterations) plus $\log(\frac{2L}{\eta_0})$. As each pass through the loop requires two gradient evaluations and two calls to a projection oracle, what remains to be shown is the stated bound on k .

Choosing $A_i \epsilon_i^\ell \leq \epsilon_i^M = \frac{a_i \epsilon^2}{8}$ and the already argued bound on η_k , we have, from Eq. (B.8):

$$\frac{1}{\eta_k + \sigma} \|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2 \leq \frac{2L \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2}{A_k} + \frac{3\epsilon^2}{4}, \quad (\text{B.9})$$

and it remains to bound the number of iterations k until $\frac{2L \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2}{A_k} \leq \frac{\epsilon^2}{4}$. To do so, notice that, as $a_0 = A_0 = 1$ and $\frac{a_k}{A_k} = \sqrt{\frac{\sigma}{2(\eta_k + \sigma)}} \leq \sqrt{\frac{\sigma}{2(2L + \sigma)}}$ for $k \geq 1$, we have that $\frac{A_{k-1}}{A_k} \leq 1 - \sqrt{\frac{\sigma}{2(2L + \sigma)}}$, and thus,

$$\frac{1}{A_k} = \frac{A_0}{A_1} \frac{A_1}{A_2} \dots \frac{A_{k-1}}{A_k} \leq \left(1 - \sqrt{\frac{\sigma}{2(2L + \sigma)}} \right)^k \leq \exp \left(-k \sqrt{\frac{\sigma}{2(2L + \sigma)}} \right).$$

Thus, k is bounded by

$$k \leq \sqrt{\frac{2(2L + \sigma)}{\sigma}} \log \left(\frac{16L \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2}{\epsilon^2} \right) = O \left(\sqrt{\frac{L}{\sigma}} \log \left(\frac{\sqrt{L} \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|}{\epsilon} \right) \right),$$

as claimed. \square

B.3. Proof of Theorem 3.6

We now proceed with proving Theorem 3.6. Before providing the full proof, we first state and prove some supporting claims.

Proposition B.4. *Given $\hat{\mathbf{y}} \in \mathcal{C}$, let $\ell(\mathbf{u}) = \langle \nabla f_\sigma(\hat{\mathbf{y}}), \mathbf{u} - \hat{\mathbf{y}} \rangle + \frac{\eta + \sigma}{2} \|\mathbf{u} - \hat{\mathbf{y}}\|^2$, $\bar{\mathbf{y}} = \operatorname{argmin}_{\mathbf{u} \in \mathcal{C}} \ell(\mathbf{u})$, and let \mathbf{y} be such that $\ell(\mathbf{y}) - \ell(\bar{\mathbf{y}}) \leq \epsilon^\ell$. Then:*

$$\frac{1}{\eta + \sigma} \|\tilde{G}_{\eta + \sigma}^\sigma(\hat{\mathbf{y}}) - G_{\eta + \sigma}^\sigma(\hat{\mathbf{y}})\|^2 \leq 2\epsilon^\ell,$$

where $\tilde{G}_{\eta + \sigma}^\sigma(\hat{\mathbf{y}}) = (\eta + \sigma)(\hat{\mathbf{y}} - \mathbf{y})$, $G_{\eta + \sigma}^\sigma(\hat{\mathbf{y}}) = (\eta + \sigma)(\hat{\mathbf{y}} - \bar{\mathbf{y}})$.

Proof. By definition of $\tilde{G}_{\eta+\sigma}^\sigma(\hat{\mathbf{y}})$, $G_{\eta+\sigma}^\sigma(\hat{\mathbf{y}})$,

$$\|\tilde{G}_{\eta+\sigma}^\sigma(\hat{\mathbf{y}}) - G_{\eta+\sigma}^\sigma(\hat{\mathbf{y}})\| = (\eta + \sigma)\|\mathbf{y} - \bar{\mathbf{y}}\|.$$

As ℓ is strongly convex and by definition of \mathbf{y} ,

$$\frac{\eta + \sigma}{2}\|\mathbf{y} - \bar{\mathbf{y}}\|^2 \leq \ell(\mathbf{y}) - \ell(\bar{\mathbf{y}}) \leq \epsilon^\ell.$$

Thus, $\|\tilde{G}_{\eta+\sigma}^\sigma(\hat{\mathbf{y}}) - G_{\eta+\sigma}^\sigma(\hat{\mathbf{y}})\| \leq (\eta + \sigma)\sqrt{\frac{2\epsilon^\ell}{\eta+\sigma}}$, as claimed. \square

It is always possible to obtain an upper estimate of the strong convexity parameter m of f . To do so, we can pick two arbitrary points \mathbf{x}, \mathbf{y} , $\mathbf{y} \neq \mathbf{x}$, from the feasible set \mathcal{C} . By strong convexity,

$$f(\mathbf{y}) \geq \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{m}{2}\|\mathbf{y} - \mathbf{x}\|^2.$$

Thus, we can use the estimate:

$$\sigma_0 = \frac{2(f(\mathbf{y}) - f(\mathbf{x}) - \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle)}{\|\mathbf{y} - \mathbf{x}\|^2} \geq m. \quad (\text{B.10})$$

Note that by smoothness of f , we also have $\sigma_0 \leq L$. We will assume throughout this section that the algorithm is started with such an estimate.

Further, we assume that the algorithm can be started with points $\mathbf{x}_0, \mathbf{y}_0$ such that

$$\ell_0(\mathbf{y}_0) - \min_{\mathbf{u} \in \mathcal{C}} \ell_0(\mathbf{u}) \leq \frac{1}{8(\eta_0 + \sigma_0)} \|\tilde{G}_{\eta_0+\sigma_0}^{\sigma_0}(\mathbf{x}_0)\|^2 = \frac{\eta_0 + \sigma_0}{8} \|\mathbf{y}_0 - \mathbf{x}_0\|^2.$$

This can be achieved by running a minimization procedure for ℓ_0 and halting it when the estimated optimality gap at its current iterate $\mathbf{y}_{0,k}$ is lower than $\frac{\eta_0 + \sigma_0}{8} \|\mathbf{y}_0^k - \mathbf{x}_0\|^2$, and outputting $\mathbf{y}_0 = \mathbf{y}_{0,k}$. For implementing a minimization procedure for ℓ_0 in our context, see (Diakonikolas et al., 2020). Note also that, as $f_{\sigma_0}(\mathbf{x}_0) = f(\mathbf{x}_0)$, we have $G_{\eta_0+\sigma_0}^{\sigma_0}(\mathbf{x}_0) = G_{\eta_0+\sigma_0}(\mathbf{x}_0)$.

We now present the pseudocode for one call to the accelerated algorithm, between restarts.

Algorithm 6 ACC($\mathbf{x}_0, \eta_0, \sigma$)

```

1:  $\sigma = 2\sigma$ 
2: repeat
3:    $\sigma = \sigma/2$ 
4:   Run a minimization procedure for  $\ell_0(\mathbf{u}) = \langle \nabla f(\mathbf{x}_0), \mathbf{u} - \mathbf{x}_0 \rangle + \frac{\eta_0 + \sigma}{2} \|\mathbf{u} - \mathbf{x}_0\|^2$ . Halt when the current iterate  $\mathbf{y}$  of
     the procedure satisfies  $\ell_0(\mathbf{y}) - \min_{\mathbf{u} \in \mathcal{C}} \ell_0(\mathbf{u}) \leq \epsilon_0$ , where  $\epsilon_0 = \frac{\eta_0 + \sigma}{32} \|\mathbf{y}_0 - \mathbf{x}_0\|^2 = \frac{1}{32(\eta_0 + \sigma)} \|\tilde{G}_{\eta_0+\sigma}(\mathbf{x}_0)\|^2$ .
5:   Set  $\hat{\mathbf{y}}_0 = \mathbf{v}_0 = \mathbf{y}_0$ ;  $\mathbf{z}_0 = (\eta_0 + \sigma)\mathbf{x}_0 - \nabla f(\mathbf{x}_0)$ 
6:    $a_0 = A_0 = 1$ 
7:   repeat
8:      $k = k + 1$ 
9:      $\eta_k, A_k, \mathbf{z}_k, \mathbf{v}_k, \hat{\mathbf{y}}_k, \mathbf{y}_k, \tilde{G}_{\eta_k+\sigma}^\sigma(\hat{\mathbf{y}}_k) = \text{AGD} - \text{Iter}(\mathbf{y}_{k-1}, \mathbf{v}_{k-1}, \mathbf{z}_{k-1}, A_{k-1}, \eta_{k-1}, \sigma, \epsilon_0, \eta_0)$ 
10:    until  $\frac{1}{\eta_k + \sigma} \|\tilde{G}_{\eta_k+\sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2 \leq \frac{9\epsilon_0}{4}$ 
11:    until  $\frac{\sigma}{\sqrt{\eta_k + \sigma}} \|\hat{\mathbf{y}}_k - \mathbf{x}_0\| \leq \sqrt{\epsilon_0}$ 
12:  return  $\hat{\mathbf{y}}_k, \eta_k \sigma$ 
    
```

Theorem 3.6. Let $\mathcal{C} \subseteq \mathbb{R}^n$ be a closed convex set, and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a function that is L -smooth and m -strongly convex on \mathcal{C} . Let $\mathbf{x}_0 \in \mathcal{C}$ be an arbitrary initial point, and let σ_0, η_0 be such that $m \leq \sigma_0 \leq \eta_0 \leq L$. Let $\mathbf{x}_0^{\text{out}} = \mathbf{x}_0$, and for $k \geq 0$, consider the following updates:

$$\mathbf{x}_{k+1}^{\text{out}}, \eta_{k+1}, \sigma_{k+1} = \text{ACC}(\mathbf{x}_k^{\text{out}}, \eta_k, \sigma_k),$$

where ACC is provided in Algorithm 3. Let $G_{\eta+\sigma}$ denote the gradient mapping of f on \mathcal{C} with parameter $\eta + \sigma$. Then, for any $\epsilon > 0$, $\|G_{\eta_k+\sigma_k}(\mathbf{x}_k^{\text{out}})\| \leq \epsilon$ for $k = O\left(\log\left(\frac{L\|G_{\eta_0+\sigma_0}(\mathbf{x}_0)\|}{m\epsilon}\right)\right)$. The algorithm for computing $\mathbf{x}_k^{\text{out}}$ utilizes a total number of

$$K = \left(\sqrt{\frac{L}{m}} \log\left(\frac{L}{m}\right) \log\left(\frac{L\|G_{\eta_0+\sigma_0}(\mathbf{x}_0)\|}{\epsilon}\right)\right)$$

queries to the FOO for f and an inexact and efficiently computable projection oracle for \mathcal{C} , without knowledge of any of the problem parameters.

Proof. The proof outline is as follows. We focus on one call to Algorithm 6 to prove that it halves the value of $\frac{1}{\eta+\sigma}\|G(\mathbf{x}^{\text{out}})\|^2$ within $k = O\left(\sqrt{\frac{L}{\sigma}} \log\left(\frac{L}{m}\right)\right)$ iterations. The stated bound then follows by showing that σ cannot be halved by more than $O(\log(\frac{L}{m}))$ times (cumulatively over the entire restarted algorithm run). Note that we have already argued in Lemma 3.5 that the total number of times that η_k can get doubled is $O(\log(\frac{\eta_k}{\eta_0})) = O(\log(\frac{L}{m}))$, this bound holds cumulatively over the entire (restarted) algorithm run, and is absorbed in the bound on K from the statement of the theorem.

Let us start by bounding the total number of iterations in the inner repeat-until loop of Algorithm 6. To keep the notation simple, we will use the same notation as in the pseudocode for Algorithm 6. As the algorithm is started with $\sigma \leq \eta_0$ and in any iteration η_k can be only increased while σ can only be decreased, we have that $\theta_k = \sqrt{\frac{\sigma}{2(\eta_k+\sigma)}} \leq \frac{1}{2}$. Thus, $\epsilon_k^\ell \leq \frac{\epsilon_0}{8}$. From Proposition B.4,

$$\frac{1}{\sqrt{\eta_k+\sigma}} \|\tilde{G}_{\eta_k+\sigma}^\sigma(\hat{\mathbf{y}}_k) - G_{\eta_k+\sigma}^\sigma(\hat{\mathbf{y}}_k)\| \leq \sqrt{2\epsilon_0^k} \leq \frac{\sqrt{\epsilon_0}}{2}. \quad (\text{B.11})$$

Suppose that $\frac{1}{\eta_k+\sigma} \|G_{\eta_k+\sigma}^\sigma(\hat{\mathbf{y}}_k)\|^2 \leq \epsilon_0$. Then, Eq. (B.11) implies that it must be

$$\frac{1}{\eta_k+\sigma} \|\tilde{G}_{\eta_k+\sigma}^\sigma(\hat{\mathbf{y}}_k)\| \leq \frac{9\epsilon_0}{4},$$

and the inner repeat-until loop reaches its exit condition. Applying Lemma 3.5, this happens within

$$k = O\left(\sqrt{\frac{L}{\sigma}} \log\left(\frac{\sqrt{L}\|\mathbf{x}_\sigma^* - \mathbf{x}_0\|}{\sqrt{\epsilon_0}}\right)\right) \quad (\text{B.12})$$

iterations. To further bound k , we have by Proposition B.1 and $(\sigma + m)$ strong convexity of f_σ that

$$\frac{\sigma + m}{2} \|\mathbf{x}_\sigma^* - \mathbf{x}_0\|^2 \leq f_\sigma(\mathbf{x}_0) - f_\sigma(\mathbf{x}_\sigma^*) \leq \frac{1}{2(\sigma + m)} \|G_{m+\sigma}^\sigma(\mathbf{x}_0)\|^2 \leq \frac{1}{2(\sigma + m)} \|G_{\eta_0+\sigma}(\mathbf{x}_0)\|^2,$$

where the last inequality follows by $\eta_0 \geq m$ and $G_{\eta_0+\sigma}^\sigma(\mathbf{x}_0) = G_{\eta_0+\sigma}(\mathbf{x}_0)$. Thus,

$$\|\mathbf{x}_\sigma^* - \mathbf{x}_0\| \leq \frac{1}{\sigma + m} \|G_{\eta_0+\sigma}(\mathbf{x}_0)\|. \quad (\text{B.13})$$

By the definition of ϵ_0 is Algorithm 6 and Proposition B.4,

$$\frac{1}{\sqrt{\eta_0+\sigma}} \|G_{\eta+\sigma}(\mathbf{x}_0) - \tilde{G}_{\eta+\sigma}(\mathbf{x}_0)\| \leq \sqrt{2\epsilon_0}, \quad (\text{B.14})$$

and, thus,

$$\frac{1}{\sqrt{\eta_0+\sigma}} \|G_{\eta+\sigma}(\mathbf{x}_0)\| \leq \frac{1}{\sqrt{\eta_0+\sigma}} \|\tilde{G}_{\eta+\sigma}(\mathbf{x}_0)\| + \sqrt{2\epsilon_0} = 3\sqrt{2\epsilon_0}.$$

Thus, Eq. (B.13) implies

$$\|\mathbf{x}_\sigma^* - \mathbf{x}_0\| \leq \frac{\sqrt{\eta_0+\sigma}}{\sigma+m} 3\sqrt{2\epsilon_0} = O\left(\frac{\sqrt{L}}{m} \sqrt{\epsilon_0}\right), \quad (\text{B.15})$$

which leads to

$$k = O\left(\sqrt{\frac{L}{\sigma}} \log\left(\frac{L}{m}\right)\right).$$

To complete the proof, it remains to:

- Argue that $\sigma \geq cm$ throughout the algorithm run, for some absolute constant $c > 0$; note that as initially $\sigma \in [m, L]$ and σ only gets halved (but never increased), the total number of times that the outer repeat-until loop is accessed beyond the first pass is $O(\log(\frac{L}{m}))$, which, combined with the bound on the number of iterations k above, gets absorbed by the stated bound on the total iteration count;
- Argue that each call to Algorithm 6 reduces the value of $\frac{1}{\eta_k + \sigma} \|G_{\eta_k + \sigma}(\hat{\mathbf{y}}_k)\|^2$ by a constant factor. Then, the total number of calls to Algorithm 6 until $\|G_{\eta_k + \sigma}(\hat{\mathbf{y}}_k)\| \leq \epsilon$ is $O(\log(\frac{\|G_{\eta_0 + \sigma_0}(\mathbf{x}_0)\|}{\epsilon}))$.

To argue about the lower bound on σ , let us bound the value of $\|\hat{\mathbf{y}}_k - \mathbf{x}_0\|$. Applying triangle inequality,

$$\|\hat{\mathbf{y}}_k - \mathbf{x}_0\| \leq \|\hat{\mathbf{y}}_k - \mathbf{x}_\sigma^*\| + \|\mathbf{x}_\sigma^* - \mathbf{x}_0\| \leq \|\hat{\mathbf{y}}_k - \mathbf{x}_\sigma^*\| + \frac{\sqrt{\eta_0 + \sigma}}{\sigma + m} 3\sqrt{2\epsilon_0}. \quad (\text{B.16})$$

Using strong convexity of f_σ and applying Proposition B.1, similar to what we did for bounding $\|\mathbf{x}_\sigma^* - \mathbf{x}_0\|$, we have

$$\|\hat{\mathbf{y}}_k - \mathbf{x}_\sigma^*\| \leq \frac{1}{\sigma + m} \|G_{\eta_k + \sigma}\| \leq \frac{2\|\tilde{G}_{\eta_k + \sigma}(\hat{\mathbf{y}}_k)\| + \sqrt{(\eta_k + m)\epsilon_0}}{2(\sigma + m)} \leq \frac{2\sqrt{(\eta_k + m)\epsilon_0}}{\sigma + m},$$

where the last two inequalities are by Eq. (B.11) and the exit condition of the inner repeat-until loop in Algorithm 6. Combining with Eq. (B.16), we have

$$\|\hat{\mathbf{y}}_k - \mathbf{x}_0\| \leq \frac{(2\sqrt{\eta_k + m} + 3\sqrt{2}\sqrt{\eta_0 + \sigma})}{\sigma + m} \sqrt{\epsilon_0}. \quad (\text{B.17})$$

It follows that, as $\eta_k \geq \eta_0$,

$$\frac{\sigma}{\sqrt{\eta_k + \sigma}} \|\hat{\mathbf{y}}_k - \mathbf{x}_0\| \leq \frac{5\sqrt{2}\sigma}{\sigma + m} \sqrt{\epsilon_0},$$

which is bounded above by $\sqrt{\epsilon_0}$ for $\sigma \leq \frac{m}{5\sqrt{2}-1}$, and, thus, satisfies the exit condition in the outer repeat-until loop in Algorithm 6. As σ can never decrease by more than a factor of 2, we have that it always holds that $\sigma \geq \frac{m}{2(5\sqrt{2}-1)}$, completing the proof that σ is bounded below by a constant factor times m .

For the remaining part of the proof, using the definition of a gradient mapping and $f_\sigma(\mathbf{x}) = f(\mathbf{x}) + \frac{\sigma}{2}\|\mathbf{x} - \mathbf{x}_0\|^2$, we have

$$\begin{aligned} \|G_{\eta_k + \sigma}(\hat{\mathbf{y}}_k) - G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\| &= (\eta_k + \sigma) \left\| P_C\left(\hat{\mathbf{y}}_k - \frac{1}{\eta_k + \sigma} \nabla f(\hat{\mathbf{y}}_k)\right) - P_C\left(\hat{\mathbf{y}}_k - \frac{1}{\eta_k + \sigma} \nabla f_\sigma(\hat{\mathbf{y}}_k)\right) \right\| \\ &\leq \|\nabla f(\hat{\mathbf{y}}_k) - \nabla f_\sigma(\hat{\mathbf{y}}_k)\| \\ &= \sigma \|\hat{\mathbf{y}}_k - \mathbf{x}_0\|, \end{aligned}$$

where we have used the fact that the projection operator is non-expansive. Now, using Eq. (B.11) and the exit condition in the inner repeat-until loop in Algorithm 6, we have

$$\begin{aligned} \frac{1}{\sqrt{\eta_k + \sigma}} \|G_{\eta_k + \sigma}(\hat{\mathbf{y}}_k)\| &\leq \frac{1}{\sqrt{\eta_k + \sigma}} \|G_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\| + \frac{\sigma}{\sqrt{\eta_k + \sigma}} \|\hat{\mathbf{y}}_k - \mathbf{x}_0\| \\ &\leq \frac{1}{\sqrt{\eta_k + \sigma}} \|\tilde{G}_{\eta_k + \sigma}^\sigma(\hat{\mathbf{y}}_k)\| + \frac{\sqrt{\epsilon_0}}{2} + \frac{\sigma}{\sqrt{\eta_k + \sigma}} \|\hat{\mathbf{y}}_k - \mathbf{x}_0\| \\ &\leq 3\sqrt{\epsilon_0}, \end{aligned}$$

where the last inequality uses $\frac{\sigma}{\sqrt{\eta_k + \sigma}} \|\hat{\mathbf{y}}_k - \mathbf{x}_0\| \leq \sqrt{\epsilon_0}$, which holds by the exit criterion of the repeat-until loop. Using the definition of ϵ_0 , we further have

$$\frac{1}{\sqrt{\eta_k + \sigma}} \|G_{\eta_k + \sigma}(\hat{\mathbf{y}}_k)\| \leq \frac{3}{\sqrt{32\sqrt{\eta_0 + \sigma}}} \|\tilde{G}_{\eta_0 + \sigma}(\mathbf{x}_0)\|.$$

Applying Eq. (B.14) and the definition of ϵ_0 , a simple calculation leads to $\|\tilde{G}_{\eta_0+\sigma}(\mathbf{x}_0)\| \leq \frac{1}{1-1/\sqrt{32}} \|G_{\eta_0+\sigma}(\mathbf{x}_0)\|$, and we can finally conclude

$$\frac{1}{\sqrt{\eta_k+\sigma}} \|G_{\eta_k+\sigma}(\hat{\mathbf{y}}_k)\| \leq \frac{3}{\sqrt{32-1\sqrt{\eta_0+\sigma}}} \|G_{\eta_0+\sigma}(\mathbf{x}_0)\| \leq \frac{1}{\sqrt{2}} \|G_{\eta_0+\sigma}(\mathbf{x}_0)\|.$$

Thus, each call to Algorithm 6 halves $\frac{1}{\eta_k+\sigma} \|G_{\eta_k+\sigma}(\mathbf{x}_k^{\text{out}})\|$, and the total number of calls until $\|G_{\eta_k+\sigma}(\mathbf{x}_k^{\text{out}})\| \leq \epsilon$ is bounded by $O(\log(\frac{L}{m} \frac{\|G_{\eta_0+\sigma_0}\|}{\epsilon}))$, leading to the bound from the statement of the theorem. \square

C. Coupling AFW with Acceleration

In this section, we prove how the AFW algorithm (Algorithm 5) can be coupled with the ACC algorithm (Algorithm 6) to achieve an algorithm, dubbed the Parameter-Free Locally accelerated Conditional Gradients (PF-LaCG), that achieves an optimal convergence rate in primal gap (up to poly-logarithmic factors), this is formalized in Theorem C.1.

Theorem C.1. *Let $\mathcal{X} \in \mathbb{R}^n$ be a closed convex polytope of diameter D , and let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a function that is L -smooth and m -strongly convex on \mathcal{X} . Let Assumptions 2.2 and 2.4 be satisfied for f, \mathcal{X} . Denote $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$. Given $\epsilon > 0$, let \mathbf{x}^{out} be the output point of PF-LaCG (Algorithm 4), initialized at an arbitrary vertex \mathbf{x}_0 of \mathcal{X} . Then $w(\mathbf{x}^{\text{out}}, \mathcal{S}^{\text{out}}) \leq \epsilon$ and PF-LaCG uses a total of at most*

$$K = O\left(\min\left\{\log\left(\frac{w(\mathbf{x}_0, \mathcal{S}_0)}{LD^2}\right) + \frac{LD^2}{m\delta^2} \log\left(\frac{w(\mathbf{x}_0, \mathcal{S}_0)}{\epsilon}\right), K_0 + K_1 + \sqrt{\frac{L}{m}} \log\left(\frac{L}{m}\right) \log\left(\frac{LD}{m\delta}\right) \log\left(\frac{LD}{\epsilon}\right)\right\}\right)$$

queries to the FOO for f and the LMO for \mathcal{X} , where

$$K_0 = \frac{32LD^2}{m\delta(\mathcal{X})^2 \ln 2} \log\left(\frac{2w(\mathbf{x}_0, \mathcal{S}_0)}{\min\{(\tau/(2D(\sqrt{L/m}+1)))^2/L, \tau, LD^2, 2w_c\}}\right),$$

and $K_1 = \frac{128LD^2}{m\delta^2}$.

Proof. As the algorithm is specified so that we always have $w^{\text{out}} = w(\mathbf{x}^{\text{out}}, \mathcal{S}^{\text{out}})$ and it terminates when $w^{\text{out}} \leq \epsilon$, it must be $w(\mathbf{x}^{\text{out}}, \mathcal{S}^{\text{out}}) \leq \epsilon$ when the algorithm terminates.

Observe that the algorithm monotonically decreases $w(\mathbf{x}^{\text{out}}, \mathcal{S}^{\text{out}})$ between restarts (by a factor of 2) and w^{out} is never larger than w^{AFW} at the end of a restart, as the algorithm sets $w^{\text{out}} = \min\{w^{\text{AFW}}, w^{\text{ACC}}\}$. Further, AFW is run almost completely independently of ACC: its running iterate and the active set are updated to those of ACC only if $w^{\text{ACC}} < w^{\text{AFW}}$ and $|\mathcal{S}^{\text{ACC}}| \leq |\mathcal{S}^{\text{AFW}}|$. Thus w^{AFW} decreases at least as fast as guaranteed by AFW, and the same holds for w^{out} , as $w^{\text{out}} \leq w^{\text{AFW}}$. This gives the first term in the bound on K from the theorem statement.

By Theorem 3.4 we know that $w(\mathbf{x}_k, \mathcal{S}) \leq w_c$ for $k \geq K_0$, which means that every active set of AFW contains \mathbf{x}^* in its convex hull. As, due to Proposition A.3, w^{AFW} is halved after at most $K_1 = \frac{128LD^2}{m\delta^2}$ additional iterations, between iterations K_0 and $K_0 + K_1$, PF-LaCG enters the if branch from Line 6. Then if $w^{\text{ACC}} \leq w_c$, we know that $\mathbf{x}^* \in \operatorname{co}(\mathcal{S}^{\text{ACC}})$. Otherwise, if $w^{\text{ACC}} > w_c$, then the AFW algorithm will obtain a point such that $w^{\text{AFW}} < w^{\text{ACC}}$ and the ACC algorithm will get updated with $\mathbf{x}^{\text{ACC}} = \mathbf{x}^{\text{AFW}}$ and $\mathcal{S}^{\text{ACC}} = \mathcal{S}^{\text{AFW}}$. From this point on, it must be the case that $\mathbf{x}^* \in \operatorname{co}(\mathcal{S}^{\text{ACC}})$, as PF-LaCG can only update \mathcal{S}^{ACC} to \mathcal{S}^{AFW} , and for $t \geq K_0$, we have $\mathbf{x}^* \in \operatorname{co}(\mathcal{S}^{\text{AFW}})$. It remains to argue that in the remaining iterations w^{out} is reduced at an accelerated rate.

Let us consider what happens between two successive restarts of PF-LaCG. Suppose first that there are r calls to ACC (Algorithm 6) in this time frame, and, to keep the notation simple, let \mathbf{x} and \mathbf{x}^+ denote the output points of ACC at the beginning and at the end of the considered restart period of LaCG, and (η, σ) and (η^+, σ^+) denote their respective smoothness and strong convexity parameter estimates from ACC. For simplicity, we let \mathbf{x}^+ coincide with the output point of the r^{th} call to ACC; if this were not the case, we could choose to output the point with the lower value of the gradient mapping between \mathbf{x}^+ and the output point of the r^{th} call to ACC, and the same bound would hold. Theorem 3.6 guarantees that

$$\frac{\|G_{\eta^++\sigma^+}(\mathbf{x}^+)\|^2}{\eta^+ + \sigma^+} \leq \left(\frac{1}{2}\right)^r \frac{\|G_{\eta+\sigma}(\mathbf{x})\|^2}{\eta + \sigma}. \quad (\text{C.1})$$

By Proposition B.2, if $\eta + \sigma \geq L$, then $\frac{1}{2(\eta+\sigma)} \|G_{\eta+\sigma}(\mathbf{x})\|^2 \leq f(\mathbf{x}) - f(\mathbf{x}^*)$. Otherwise, using monotonicity of gradient mapping (see Proposition B.1), $\frac{1}{2(\eta+\sigma)} \|G_{\eta+\sigma}(\mathbf{x})\|^2 \leq \frac{1}{2(\eta+\sigma)} \|G_L(\mathbf{x})\|^2 \leq \frac{L}{\eta+\sigma} (f(\mathbf{x}) - f(\mathbf{x}^*))$. Either way, as $\eta \geq m$:

$$\begin{aligned} \frac{\|G_{\eta+\sigma}(\mathbf{x})\|^2}{2(\eta+\sigma)} &\leq \max\{1, L/(\eta+\sigma)\} (f(\mathbf{x}) - f(\mathbf{x}^*)) \\ &\leq L(f(\mathbf{x}) - f(\mathbf{x}^*)) / m \\ &\leq 2L(w_{\text{prev}}^{\text{ACC}} / (m\delta))^2, \end{aligned} \quad (\text{C.2})$$

where the last inequality is by $w(\mathbf{x}) \leq w(\mathbf{x}, \mathbf{s}) = w_{\text{prev}}^{\text{ACC}}$. On the other hand, as $\eta^+ + \sigma^+ > m$, Proposition B.1 gives $f(\mathbf{x}^+) - f(\mathbf{x}^*) \leq \frac{1}{2m} \|G_{\eta^++\sigma^+}(\mathbf{x}^+)\|^2$. Using the strong Wolfe gap bound in Theorem 3.4 it follows that

$$\begin{aligned} \frac{1}{\eta^+ + \sigma^+} \left(\frac{mw^{\text{ACC}}}{LD} \right)^2 &\leq \frac{2m}{\eta^+ + \sigma^+} (f(\mathbf{x}^+) - f(\mathbf{x}^*)) \\ &\leq \frac{1}{\eta^+ + \sigma^+} \|G_{\eta^++\sigma^+}(\mathbf{x}^+)\|^2. \end{aligned} \quad (\text{C.3})$$

Combining Eqs. (C.1)–(C.3), we have

$$(w^{\text{ACC}})^2 \leq \left(\frac{1}{2} \right)^r \frac{4L^3 D^2 (\eta^+ + \sigma^+)}{m^4 \delta^2} (w_{\text{prev}}^{\text{ACC}})^2.$$

Thus, if $r \geq r^* = \log_2 \left(\frac{16L^3 D^2 (\eta^+ + \sigma^+)}{m^4 \delta^2} \right) = O(\log(\frac{LD}{m\delta}))$, we have that $w^{\text{ACC}} \leq \frac{1}{2} w_{\text{prev}}^{\text{ACC}}$. The total number of iterations in this case is $r \cdot O\left(\sqrt{\frac{L}{m}} \log\left(\frac{L}{m}\right)\right)$, due to Lemma 3.5.

Observe that if $r \geq r^* + p$ for some $p \geq 1$, then

$$(w^{\text{ACC}})^2 \leq \left(\frac{1}{2} \right)^p (w_{\text{prev}}^{\text{ACC}})^2,$$

and past the first r calls to ACC, w^{ACC} halves on every other call to ACC, i.e., it contracts much faster than every r^* iterations.

We now argue that w^{ACC} halves at least as often as every

$$r^* \cdot O\left(\sqrt{\frac{L}{m}} \log\left(\frac{L}{m}\right)\right) = O\left(\sqrt{\frac{L}{m}} \log\left(\frac{L}{m}\right) \log\left(\frac{LD}{m\delta}\right)\right)$$

iterations. To do so, we need to argue that a restart that updates \mathbf{x}^{ACC} to \mathbf{x}^{AFW} does not slow down the overall convergence. Suppose that there is such a restart. Then, by the condition from Line 9 of PF-LaCG, one of the following two situations must occur. If the restart period (number of calls to ACC) was longer than r^* , then w^{ACC} was contracting between restarts at least as fast as if ACC was run independently, and, as $w^{\text{AFW}} \leq w^{\text{ACC}}$ on restart and w^{ACC} is updated to w^{AFW} , w^{ACC} must halve at least as frequently as every r^* calls to ACC. If the restart period was shorter than r^* , then, as $w^{\text{AFW}} \leq w_{\text{prev}}^{\text{ACC}}/2$ and w^{ACC} is updated to w^{AFW} , we get that in this case w^{ACC} is halved in fewer than r^* calls to ACC.

Finally, as $w^{\text{out}} \leq w^{\text{ACC}}$ and each iteration requires a constant number of calls to the FOO for f and an LMO for \mathcal{X} , the claimed bound on K follows. \square

D. Computational Results

In this section we provide a complete overview of the implementation of PF-LaCG and a detailed comparison of the performance of PF-LaCG relative to other state of the art parameter-free algorithms. All the experiments in this paper were run on a Linux machine with an Intel Xeon Processor (Skylake, IBRS) and 64 GB of RAM.

D.1. Running ACC in Parallel with AFW

One of the key computational advantages of PF-LaCG is that it allows local acceleration speedups through parallelism, guaranteeing nearly as much progress as state-of-the-art CG algorithms such as AFW and PFW in terms of wall-clock time. This is possible due to the way PF-LaCG is structured: between restarts where we use $w(\mathbf{x}, \mathcal{S})$ as the measure of optimality, the ACC and the AFW algorithms are executed completely independently of each other while AFW checks whether ACC has made sufficient progress whenever AFW has halved $w(\mathbf{x}, \mathcal{S})$ without interrupting ACC's execution. As such, PF-LaCG has the potential to utilize twice the computational power when compared to AFW by running the locally accelerated algorithm ACC on a secondary connected machine or on a separate process within a single machine. As a proof of concept in our experiments, we implement the parallelism of PF-LaCG using Python's `multiprocessing` library. We utilize its recently developed functionality `shared_memory` to provide efficient inter-process communications whenever a restart happens and to orchestrate iteration synchronization through semaphores and locks. Moreover, in order to simulate the performance comparisons between PF-LaCG and other CG algorithms in the setting where PFLaCG has access to twice the computational power, we limit each process (note that PF-LaCG runs on two processes) to run on one virtualized CPU core through restricting the number of threads used by Intel's high-performance computing library MKL to one. In Section D.6, we also show how PF-LaCG outperforms other CG algorithms even when given equal amount of total computational power.

D.2. Solving Minimization Subproblems in ACC and AGD-Iter

Two of the key steps in the AGD-Iter algorithm are the approximate projections that need to be carried out in Algorithm 2. We give a brief description of how these subproblems can be solved, in a very similar way as how they were solved in [Diakonikolas et al. \(2020, Section C.3\)](#). Both of these subproblems can be written without loss of generality as:

$$\mathbf{v} = \underset{\mathbf{u} \in \mathcal{C}}{\overset{\epsilon}{\operatorname{argmin}}} \langle \mathbf{z}, \mathbf{u} \rangle + \|\mathbf{u}\|^2, \quad (\text{D.1})$$

where \mathcal{C} is the convex hull of a known set of vertices, i.e., $\mathcal{C} = \operatorname{co}(\mathcal{S})$, and $\underset{\mathbf{u} \in \mathcal{C}}{\overset{\epsilon}{\operatorname{argmin}}}$ is used to indicate that the objective function that follows the argmin is minimized to additive error ϵ , as we described in the main body of the paper. We can write the subproblem shown in Eq. (D.1) as an equivalent problem over the unit probability simplex of dimension $|\mathcal{S}|$, which we denote by $\Delta_{|\mathcal{S}|}$, where $|\mathcal{S}|$ is the cardinality of the set \mathcal{S} . This allows us to write $\mathbf{u} = \mathcal{V}\boldsymbol{\lambda}$, where \mathcal{V} is the matrix whose columns are the elements of the set \mathcal{S} and $\boldsymbol{\lambda} \in \Delta_{|\mathcal{S}|}$. This leads to $\mathbf{v} = \mathcal{V}\boldsymbol{\lambda}_{\mathbf{v}}$, where:

$$\boldsymbol{\lambda}_{\mathbf{v}} = \underset{\boldsymbol{\lambda} \in \Delta_{|\mathcal{S}|}}{\overset{\epsilon}{\operatorname{argmin}}} \langle \mathbf{z}, \mathcal{V}\boldsymbol{\lambda} \rangle + \|\mathcal{V}\boldsymbol{\lambda}\|^2. \quad (\text{D.2})$$

As noted earlier, because Euclidean projections onto $\Delta_{|\mathcal{S}|}$ can be computed in closed-form with reasonable complexity, we can use accelerated projection-based methods to compute an ϵ -optimal solution to Problem (D.2) efficiently. In the code, we use the projections onto the probability simplex based on the `quicksort` algorithm with a worst-case complexity of $\mathcal{O}(|\mathcal{S}| \log(|\mathcal{S}|))$ ([Held et al., 1974](#); [Duchi et al., 2008](#)). Note that there exist projections onto the simplex with worst-case complexity of $\mathcal{O}(|\mathcal{S}|)$ using a variation of the `quicksort` algorithm that uses median-pivot partitioning ([Condat, 2016](#)). However, we have used the projections onto the simplex using the standard `quicksort` algorithm due to its simplicity and the fact that there are fast and reliable implementations of the aforementioned sorting algorithm in Python.

In order to ensure that we reach an ϵ -optimal solution to Problem (D.2) we use one of the following two criteria:

Frank-Wolfe gap stopping criterion. In order to compute the approximate solutions to Problem (D.2) one could use the Frank-Wolfe gap as a stopping criterion, that is, we stop running the accelerated projection-based algorithm when $\boldsymbol{\lambda}_{\mathbf{v}}$ satisfies:

$$\max_{\boldsymbol{\lambda} \in \Delta_{|\mathcal{S}|}} \langle \mathcal{V}^T \mathbf{z} + 2\mathcal{V}^T \mathcal{V} \boldsymbol{\lambda}_{\mathbf{v}}, \boldsymbol{\lambda}_{\mathbf{v}} - \boldsymbol{\lambda} \rangle \leq \epsilon \quad (\text{D.3})$$

Note that Problem (D.2) is convex, and so the Frank-Wolfe gap provides a useful upper bound on the primal gap. Ensuring that the Frank-Wolfe gap is below the tolerance ϵ ensures that the primal gap is below the tolerance ϵ too. Moreover, note that the quantity $\mathcal{V}^T \mathbf{z} + 2\mathcal{V}^T \mathcal{V} \boldsymbol{\lambda}_{\mathbf{v}}$ is readily computed at each iteration, as it constitutes the gradient of the objective function at $\boldsymbol{\lambda}_{\mathbf{v}}$, and is used in the projection-based accelerated algorithm. Furthermore, solving a linear optimization problem like the one shown in Eq. (D.3) over the probability simplex has complexity $|\mathcal{S}|$, and so using the Frank-Wolfe as a stopping criterion does not add a noticeable overhead to the resolution of the subproblems.

Gradient mapping stopping criterion. Alternatively, one could use the norm of the gradient mapping for Problem (D.2) to bound above the primal gap in the case where the objective function in the subproblem is strongly-convex. That is, if we denote the objective function being minimized in Eq. (D.2) as $g(\lambda)$, and we have that the smallest eigenvalue of the Hessian of $g(\lambda)$, which we denote by m for simplicity, is greater than zero, this means that if:

$$\frac{m}{2} \left\| \lambda_v - P_{\Delta_{|S|}} \left(\lambda_v - \frac{1}{m} \nabla g(\lambda_v) \right) \right\|^2 \leq \epsilon,$$

then the primal gap at λ_v for Problem (D.2) is also smaller than ϵ . To compute this stopping criterion we require knowledge of the smallest eigenvalue of $\nabla^2 g(\lambda) = 2\mathcal{V}^T \mathcal{V}$, which is a quantity that is already computed in the projection-based accelerated gradient descent algorithms, as it is used to set the step size of the accelerated algorithms. Lastly, note that as we have mentioned above, there are efficient ways to compute closed-form projections onto the probability simplex $\Delta_{|S|}$, so computing $P_{\Delta_{|S|}} \left(\lambda_v - \frac{1}{m} \nabla g(\lambda_v) \right)$ does not pose a high cost.

Remark D.1 (LLVM-enhanced subproblem solver). We also use Numba (Lam et al., 2015), a *Just-In-Time* Python compiler that uses the LLVM compiler library that transforms the accelerated projection-based optimization algorithm to machine code, in order to more efficiently solve the subproblems from Eq. (D.2).

Remark D.2 (Computing LMO). For the problems without closed form solutions for LMO (such as structured LASSO and constrained matching), we use Scipy’s `linprog` function to compute the LMO.

D.3. PF-LaCG over the Probability Simplex

The structure of the unit probability simplex can give us deep insight into how the PF-LaCG algorithm works. Although it can be considered as a toy example, as projections onto this feasible region can be computed with a complexity equal to that of solving a linear program, it allows us to know exactly when the acceleration should kick in. Assume we are minimizing a smooth and strongly convex function $f(\mathbf{x})$ over the unit probability simplex Δ_n in \mathbb{R}^n . As it is immediate to compute a proper support for any point $\mathbf{x} \in \Delta_n$ (it suffices to find the non-zero elements in the vector \mathbf{x}) we can easily find $\mathcal{F}(\mathbf{x}^*)$ if we know \mathbf{x}^* , or a high accuracy solution to the optimization problem. This enables us to pinpoint when the active set of the AFW algorithm is equal to the vertices of the optimal face, that is $\mathcal{S}_k = \text{vert}(\mathcal{F}(\mathbf{x}^*))$. We know from the proof of Theorem C.1 that immediately following the restart after the iteration where we have that $\mathcal{S}_k = \text{vert}(\mathcal{F}(\mathbf{x}^*))$ we should observe that the ACC algorithm converges at an accelerated rate.

We present in Figure 4 a comparison of several CG algorithms minimizing a quadratic function over the unit probability simplex. The function being minimized in this example is $f(\mathbf{x}) = \mathbf{x}^T (M^T M + \alpha \mathbf{1}_n) \mathbf{x} / 2 + \mathbf{b}^T \mathbf{x}$, where $M \in \mathbb{R}^{n \times n}$ and $\mathbf{b} \in \mathbb{R}^n$ have entries sampled uniformly at random between 0 and 1 and $n = 10000$. The parameter $\alpha = 500$ is set so that the objective function satisfies $m \approx 500$. The resulting condition number is $L/m = 50000$, and the number of nonzero elements in \mathbf{x}^* is around 320. The AFW algorithm in PF-LaCG (AFW) satisfies that $\mathbf{x}^* \in \text{co}(\mathcal{S}_k)$ around iteration 400, consequently we achieve the accelerated convergence rate from then onwards. The same can be said regarding PF-LaCG (PFW) around iteration 350.

Remark D.3 (Using the structure of Δ_n). The structure of the unit probability simplex allows us to simplify the CG-variants employed in the comparison, as well as the PF-LaCG algorithm. Note that at each iteration the AFW and PFW algorithms (as well as the AFW algorithm) require maintaining an active set $\mathcal{S}_k \subseteq \text{vert}(\mathcal{X})$ that contains the vertices that give rise to the iterate \mathbf{x}_k as a convex combination. This active set is used to compute the *away vertices* in the algorithms, denoted by \mathbf{s}_k . In order to compute an away vertex one has to solve the linear optimization problem $\mathbf{s}_k = \arg\max_{\mathbf{v} \in \mathcal{S}_k} \langle \nabla f(\mathbf{x}_k), \mathbf{v} \rangle$, which typically requires looping through all the elements in \mathcal{S}_k . This can become computationally expensive as $|\mathcal{S}_k|$ grows. However, computing \mathbf{s}_k is extremely easy when $\mathbf{x}_k \in \Delta_n$, as one need only look at the non-zero elements in \mathbf{x}_k to solve the linear program. Therefore no CG algorithm needs to explicitly maintain an active set when solving a problem over the probability simplex. This significantly reduces the running time of all the CG algorithms in our implementation. Note that the absence of an active set is one of the main advantages of the DICG algorithm (Garber & Meshi, 2016) (aside from the notable fact that it is able to achieve a primal gap linear convergence guarantee that depends on the dimensionality of $\mathcal{F}(\mathbf{x}^*)$, as opposed to n), which often allows it to outperform other CG variants when the feasible region \mathcal{X} is a $0 - 1$ polytope. In this example over the probability simplex, the PFW and DICG algorithms are equivalent, and we do not observe a significant advantage from using the DICG algorithm.




Figure 4. **Strongly convex and smooth problem over the probability simplex:** Algorithm convergence in terms of $f(\mathbf{x}_k) - f(\mathbf{x}^*)$ and $w(\mathbf{x}_k, \mathcal{S}_k)$ versus iteration count k and versus wall-clock time in seconds.

Moreover, if we focus on the problem shown in Eq. (D.1) we can make use of the fact that:

$$\begin{aligned}
 \operatorname{argmin}_{\boldsymbol{\lambda} \in \Delta_{|\mathcal{S}|}} \langle \mathbf{z}, \mathcal{V} \boldsymbol{\lambda} \rangle + \|\mathcal{V} \boldsymbol{\lambda}\|^2 &= \operatorname{argmin}_{\boldsymbol{\lambda} \in \Delta_{|\mathcal{S}|}} \langle \mathcal{V}^T \mathbf{z}, \boldsymbol{\lambda} \rangle + \|\boldsymbol{\lambda}\|^2 \\
 &= \operatorname{argmin}_{\boldsymbol{\lambda} \in \Delta_{|\mathcal{S}|}} -2 \left\langle -\frac{\mathcal{V}^T \mathbf{z}}{2}, \boldsymbol{\lambda} \right\rangle + \|\boldsymbol{\lambda}\|^2 + \left\| \frac{\mathcal{V}^T \mathbf{z}}{2} \right\|^2 \\
 &= \operatorname{argmin}_{\boldsymbol{\lambda} \in \Delta_{|\mathcal{S}|}} \left\| -\frac{\mathcal{V}^T \mathbf{z}}{2} - \boldsymbol{\lambda} \right\|^2.
 \end{aligned}$$

The last expression in the chain of equalities is nothing but the Euclidean projection of $-\frac{\mathcal{V}^T \mathbf{z}}{2}$ onto $\Delta_{|\mathcal{S}|}$, which as we have

stated before, can be computed in closed-form efficiently. Moreover, we can recover the active set of any iterate by simply looking at the non-zero elements of the vector, allowing us to easily recover \mathcal{S} given \mathbf{x} . We make use of this fact in our implementation of PF-LaCG over the probability simplex. This means that as the subproblems are solved to optimality in Algorithm 2 we do not require an accelerated algorithm, or a stopping criterion to solve the problems in Eq. D.2

D.4. PF-LaCG over Structured LASSO Regression Problems

We give a brief motivation for the structured LASSO regression problem solved in Section 4. The *Least Absolute Shrinkage and Selection Operator* (LASSO) (Tibshirani, 1996) is an immensely popular regression analysis that simultaneously performs variable selection and regularization to solve a linear regression problem. The formulation is intimately related to the *Basis Pursuit Denoising* (BPD) (Chen et al., 1998) problem in the signal processing community. One of the attractive properties of the LASSO is its ability to return sparse solutions that capture the variables that contribute the most towards producing an output. One of the domains in which CG-type algorithms have received attention is in sparse regression problems in physics (Carderera & Pokutta, 2020).

In many situations we can describe a physical system by the differential equations that govern the phenomenon. These equations allow us to compactly describe the current state of a physical system, or to predict its future state. However, in many situations we do not have any physics-informed differential equation models to describe a natural phenomenon, and we only have access to the state of the system at various times. Our goal then is to find this system of differential equations given some training data. That is, if we denote the state of the system at time t by $\mathbf{x}(t)$, we want to find $d\mathbf{x}(t)/dt = F(\mathbf{x}(t))$, where we assume that $F(\mathbf{x}(t))$ can be expressed as a linear combination of simple *ansatz* functions (like polynomials) that belong to a dictionary $\mathcal{D} = \{\psi_i \mid i \in [1, m]\}$, with $\psi_i : \mathbb{R}^n \rightarrow \mathbb{R}$. This allows us to write $F(\mathbf{x}(t)) = \Xi^T \psi(\mathbf{x}(t))$ where $\Xi \in \mathbb{R}^{m \times n}$ is a sparse matrix $\Xi = [\xi_1, \dots, \xi_n]$ formed by column vectors $\xi_i \in \mathbb{R}^m$ for $i \in [1, n]$ and $\psi(\mathbf{x}(t)) = [\psi_1(\mathbf{x}(t)), \dots, \psi_m(\mathbf{x}(t))]^T \in \mathbb{R}^m$. Therefore the sparse matrix Ξ allows us to reconstruct the differential equations that govern the system. If we are given a series of data points $\{\mathbf{x}(t_i), d\mathbf{x}(t_i)/dt\}_{i=1}^r$, in the absence of noise we will have:

$$\begin{bmatrix} d\mathbf{x}(t_1)/dt & \cdots & d\mathbf{x}(t_r)/dt \\ \vdots & & \vdots \end{bmatrix} = \begin{bmatrix} \xi_1 & \cdots & \xi_n \\ \vdots & & \vdots \\ \xi_n & \cdots & \xi_n \end{bmatrix} \begin{bmatrix} \psi(\mathbf{x}(t_1)) & \cdots & \psi(\mathbf{x}(t_r)) \\ \vdots & & \vdots \end{bmatrix}.$$

However, we are typically only given access to noise-corrupted measurements of $\{\mathbf{x}(t_i)\}_{i=1}^r$, which means that we have access to some noisy $\{\mathbf{y}(t_i)\}_{i=1}^r$, from which we have to estimate the derivatives $\{d\mathbf{y}(t_i)/dt\}_{i=1}^r$. In the presence of this noise, and in pursuit of a sparse matrix Ξ , we can attempt to use the LASSO regression analysis for some $\alpha > 0$, resulting in:

$$\underset{\substack{\|\Omega\|_{1,1} \leq \tau \\ \Omega \in \mathbb{R}^{m \times n}}}{\operatorname{argmin}} \left\| \dot{\mathbf{Y}} - \Omega^T \Psi(Y) \right\|_F^2. \quad (\text{D.4})$$

Where $\|\cdot\|_{1,1}$ and $\|\cdot\|_F$ represent the $\ell_{1,1}$ and Frobenius norm of a matrix, and we have collected the data into matrices $\dot{\mathbf{Y}} = [d\mathbf{y}(t_1)/dt, \dots, d\mathbf{y}(t_r)/dt] \in \mathbb{R}^{n \times r}$ and $\Psi(Y) = [\psi(\mathbf{x}(t_1)), \dots, \psi(\mathbf{x}(t_r))] \in \mathbb{R}^{m \times r}$. Moreover, we can also try to impose physics-informed linear constraints on the problem shown in Eq. (D.4) to reflect symmetries or conservation properties in the system, in the hope that the learnt dynamics will generalize better to unseen data. This transforms the problem to:

$$\underset{\Omega \in \mathcal{X}}{\operatorname{argmin}} \left\| \dot{\mathbf{Y}} - \Omega^T \Psi(Y) \right\|_F^2. \quad (\text{D.5})$$

with $\mathcal{X} = \left\{ \Omega \in \mathbb{R}^{m \times n} \mid \|\Omega\|_{1,1} \leq \tau, \operatorname{trace}(A_l^T \Omega) \leq b_l, l \in [1, L] \right\}$, $A_l \in \mathbb{R}^{m \times n}$ and $b_l \in \mathbb{R}$ for all $l \in [1, L]$, and where we have added L additional linear constraints to the problem in Eq. (D.4) to reflect the underlying structure of the dynamical system that we want to impose.

We solve a stylized version of the problem in Eq. (D.5), where the objective function is a quadratic $f(\mathbf{x}) = \mathbf{x}^T (M^T M + \alpha \mathbf{1}_n) \mathbf{x} / 2 + \mathbf{b}^T \mathbf{x}$, where $M \in \mathbb{R}^{n \times n}$ has entries sampled uniformly at random between 0 and 1, $\mathbf{b} \in \mathbb{R}^n$ has entries sampled uniformly at random from 0 to 100, $n = 1000$ and $\alpha = 100$. The resulting condition number of the function




Figure 5. **Strongly convex and smooth problem over a structured LASSO domain:** Algorithm convergence in terms of $f(\mathbf{x}_k) - f(\mathbf{x}^*)$ and $w(\mathbf{x}_k, \mathcal{S}_k)$ versus iteration count k and versus wall-clock time in seconds.

is $L/m = 250000$. The additional linear constraints we impose on the system are very similar to those used in Carderera et al. (2021), with the exception here that we generate them at random. To generate the additional equality constraints, we sample 125 pairs of distinct integers (i, j) from $1 \leq i, j \leq n$ without replacement, and we set $x_i = x_j$ for each pair, adding 125 linear constraints. Lastly, the radius of the ℓ_1 ball is set to $\tau = 1$. In this example, as the polytope is not a $0 - 1$ polytope, we cannot use the DICG algorithm of Garber & Meshi (2016). Despite the fact that we could resort to the more general decomposition invariant CG algorithm in Bashiri & Zhang (2017), we did not find it to be numerically comparable to the remaining CG algorithms tested in this section, and so have not included it in the comparison. The results obtained can be found in Figure 5

D.5. PF-LaCG over Constrained Matching Problems

We also solve a matching-type problem over the intersection of the Birkhoff polytope and a set of additional linear constraints. The Birkhoff polytope in $\mathbb{R}^{n \times n}$, also called the polytope of doubly-stochastic matrices, is the set of all square matrices whose columns and rows all sum up to 1. This polytope, with close ties to graph theory, is often used in matching problems. For example, if we interpret the rows of the matrix as workers, and the columns of the matrix as tasks that need to be completed, we can view the matrix element $A_{i,j}$ on the i^{th} row and the j^{th} column as being either 0 or 1, depending on if the j^{th} task has been assigned to the i^{th} worker (if $A_{i,j} = 1$), or if the j^{th} task has not been assigned to the i^{th} worker (if $A_{i,j} = 0$).




Figure 6. **Strongly convex and smooth matching problem over a structured Birkhoff polytope:** Algorithm convergence in terms of $f(\mathbf{x}_k) - f(\mathbf{x}^*)$ and $w(\mathbf{x}_k, \mathcal{S}_k)$ versus iteration count k and versus wall-clock time in seconds.

We minimize a quadratic cost function over a Birkhoff polytope of dimension 400, that is with $n = 20$, where the objective function has the form $f(\mathbf{x}) = \mathbf{x}^T (M^T M + \alpha \mathbf{1}_n) \mathbf{x} / 2 + \mathbf{b}^T \mathbf{x}$ with $\alpha = 1$, and where $M \in \mathbb{R}^{n \times n}$ and $\mathbf{b} \in \mathbb{R}^n$ have

entries sampled uniformly at random between 0 and 1. This results in an objective function with a condition number of $L/m = 100000$. In order to make the problem more challenging to solve, we impose an additional set of linear constraints on the learning problem. Otherwise if we were solving the problem over the Birkhoff polytope in $\mathbb{R}^{n \times n}$ we could efficiently solve linear minimization problems over the aforementioned polytope using the Hungarian algorithm, with complexity $\mathcal{O}(n^3)$. The additional constraints that we impose represent either worker-task assignments that are not permitted, or capacity constraints that represent the maximum fractional matching that we can have between a given task and a worker. In order to generate these extra constraints we sample 80 integers i from $1 \leq i \leq n^2$ without replacement, and we set $x_i = 0$ for the first 40 integers (to represent that certain matchings are not possible), and $x_i \leq 0.5$ for the remaining 40 integers to represent a maximum fractional matching. As in the previous example, we did not find the algorithm in Bashiri & Zhang (2017) to have numerical performance comparable to the other algorithms tested, and so have not included the algorithm in the comparison. The results from the comparison can be seen in Figure 6.

D.6. Performance Comparison with 2 Cores

In Figure 7 we show a performance comparison for the different algorithms where we use 2 CPU cores for the CG-variants, and 2 cores for the PFLaCG algorithms (1 core for the AGD algorithm, and 1 core for the AFW algorithm). For the sparse regression problem over the structured LASSO feasible region and the matching problem over the structured Birkhoff polytope we did not observe a large increase in performance for the CG-variants by using 2 cores instead of 1. For example, for the sparse regression problem the PFW algorithm went from reaching a value of $w(\mathbf{x}, \mathcal{S})$ below 10^{-5} in 188 seconds with 1 core, to reaching it in 166 seconds with 2 cores, which constitutes a 10% increase in performance with respect to wall-clock time. Similarly, the matching problem went from reaching a value of $w(\mathbf{x}, \mathcal{S})$ below 10^{-4} in 4.75 seconds with 1 core, to reaching it in 4.63 seconds with 2 cores, which is a meager 2% increase in performance with respect to wall-clock time. On the other hand, for the probability simplex example we can see that the algorithm reaches a value of $w(\mathbf{x}, \mathcal{S})$ below 10^{-5} in 190 seconds with 1 core, and it reaches that value in 96 seconds using 2 cores, which constitutes a 50% increase in performance with respect to wall-clock time.




Figure 7. **Performance comparison using 2-cores for CG-variants:** Comparison of $w(\mathbf{x}_k, \mathcal{S}_k)$ vs time (seconds) for the unit probability simplex experiment in Figure (a), for the structured LASSO problem in Figure (b) and for the structured matching problem in Figure (c).

Regardless of if we use 1 core or 2 cores for the CG variants in the comparison, we still obtain faster convergence in wall-clock time when using the PFLaCG algorithm. Finally, note that the convergence in terms of iteration count is independent on the number of cores used in the experiment.