
Reducing Discretization Error in the Frank-Wolfe Method

Zhaoyue Chen
Stony Brook University

Yifan Sun
Stony Brook University

Abstract

The Frank-Wolfe algorithm is a popular method in structurally constrained machine learning applications, due to its fast per-iteration complexity. However, one major limitation of the method is a slow rate of convergence that is difficult to accelerate due to erratic, zig-zagging step directions, even asymptotically close to the solution. We view this as an artifact of discretization; that is to say, the Frank-Wolfe *flow*, which is its trajectory at asymptotically small step sizes, does not zig-zag, and reducing discretization error will go hand-in-hand in producing a more stabilized method, with better convergence properties. We propose two improvements: a multistep Frank-Wolfe method that directly applies optimized higher-order discretization schemes; and an LMO-averaging scheme with reduced discretization error, and whose local convergence rate over general convex sets accelerates from a rate of $O(1/k)$ to up to $O(1/k^{3/2})$.

1 INTRODUCTION

The Frank Wolfe algorithm (FW) or the conditional gradient algorithm (Levitin and Polyak, 1966) is a popular method in constrained convex optimization. It was first developed in Frank et al. (1956) for maximizing a concave quadratic programming problem with linear inequality constraints, and later extended in Dunn and Harshbarger (1978) to minimizing more general smooth convex objective function on a bounded convex set. More recently, Jaggi (2013) analyzed (FW) over general convex and continuously differentiable objective functions with convex and compact constraint sets, and illustrates that when a sparse structural property is desired, the per-iteration cost can be much cheaper than computing projections. This has spurred a renewed interest

of (FW) to broad applications in machine learning and signal processing (Lacoste-Julien et al., 2013; Krishnan et al., 2015), recommender systems (Freund et al., 2017), image and video co-localization (Joulin et al., 2014), etc.

Specifically, (FW) solves the constrained optimization

$$\underset{x \in \mathcal{D}}{\text{minimize}} \quad f(x) \quad (1)$$

via the repeated iterations

$$\begin{aligned} \mathbf{s}_k &= \underset{\mathbf{s} \in \mathcal{D}}{\text{argmin}} \quad \nabla f(\mathbf{x}_k)^T \mathbf{s}, \\ \mathbf{x}_{k+1} &= \mathbf{x}_k + \gamma_k (\mathbf{s}_k - \mathbf{x}_k). \end{aligned} \quad (\text{FW})$$

The first operation is often referred to as the *linear minimization oracle (LMO)*, and is the support function of \mathcal{D} at $-\nabla f(\mathbf{x})$:

$$\text{LMO}_{\mathcal{D}}(\mathbf{x}) := \underset{\mathbf{s} \in \mathcal{D}}{\text{argmin}} \quad \nabla f(\mathbf{x})^T \mathbf{s}.$$

The iterates \mathbf{s}_k are usually the vertices of \mathcal{D} and are often called *atoms*, as their convex combinations build $\mathbf{x}_k \in \mathcal{D}$.

In particular, computing the LMO over \mathcal{D} is often computationally cheap compared to the corresponding projection, especially when \mathcal{D} is the level set of a sparsifying norm, e.g. the 1-norm or the nuclear norm. However, the tradeoff of the cheap per-iteration rate is that the overall convergence rate, in terms of number of iterations k , is often much slower than that of projected gradient descent (Lacoste-Julien and Jaggi, 2015; Freund and Grigas, 2016), and in particular shows no signs of acceleration even when f is μ -strongly convex or when momentum-based acceleration techniques are employed. While various acceleration schemes (Lacoste-Julien and Jaggi, 2015) have been proposed and several improved rates given under specific problem geometry (Garber and Hazan, 2015), by and large the “vanilla” Frank-Wolfe method, using the “well-studied step size” $\gamma_k = O(1/k)$, can only be shown to reach $O(1/k)$ convergence rate in terms of objective value decrease (Canon and Cullum, 1968; Jaggi, 2013; Freund and Grigas, 2016)

The Zig-Zagging phenomenon. The slowness of the Frank-Wolfe method is often explained as a consequence of a potential “zig-zagging” phenomenon. In particular, when the true solution lies on a low dimensional facet and the

incoming iterate is angled in a particular way, the method will alternate picking up vertices of this facet, causing a “zig-zagging” pattern; e.g. $\mathbf{x}_{k+1} - \mathbf{x}_k$ does not in general point toward \mathbf{x}^* . For this reason, techniques like line search and momentum-based acceleration have little effect. In fact, methods like the Away-Step Frank Wolfe (Lacoste-Julien and Jaggi, 2015) are designed to counter exactly this, by forcing the iterate to change its angle and approach more directly.

Frank Wolfe flow. In this paper, we investigate the *dynamical systems interpretation* of the Frank-Wolfe method. In particular, we study (FWFLOW), whose Euler discretization is (FW):

$$\begin{aligned} s(t) &= \text{LMO}_{\mathcal{D}}(x(t)) \\ \dot{x}(t) &= \gamma(t)(s(t) - x(t)) \end{aligned} \quad (\text{FWFLOW})$$

In particular, (FWFLOW) is an example of Krasovskii regularization (Sanfelice et al., 2008; Krasovskii, 1968) and thus existence of solutions is ensured. This dynamical system first studied in Jacimovic and Geary (1999), and is a part of the construct presented in Diakonikolas and Orecchia (2019). However, neither paper considered the effect of using advanced discretization schemes to better imitate the flow, as a way of improving the method. From an optimization point of view, (FWFLOW) is important because, under the family of decay sequences $\gamma(t) = O(1/t)$, its convergence rate is *arbitrarily close to linear*, under the right parameter choices; in contrast, (FW) is upper bounded by a sublinear $O(1/k)$ convergence rate.

Viewing the excess error in (FW) compared to (FWFLOW) as discretization error, this work looks at characterizing and attacking this source of error, in efforts of proposing an improved Frank-Wolfe method. In particular, we argue that this discretization error is particularly detrimental in almost all useful cases of FW applied to *sparse optimization applications*, e.g. when \mathcal{D} is the level-set of the one-norm or another sparsifying penalty. In this case, improvements found in related works may not apply.

- When the solution \mathbf{x}^* to (1) is in the interior of \mathcal{D} , then using *line search*, (FW) converges with a linear rate if f is strongly-convex (Guélat and Marcotte, 1986). However, this corresponds to a fully dense \mathbf{x}^* , which is not desired in sparse optimization.
- When the set \mathcal{D} is strongly convex and \mathbf{x}^* is on the boundary of \mathcal{D} , or when \mathbf{x}^* is on a minimal facet of \mathcal{D} , then (FW) converges with an $O(1/k^2)$ rate (Garber and Hazan, 2015). However, the 1-norm ball (and other popular choices of \mathcal{D}) are not strongly convex, and the only time when this scenario will apply is if \mathbf{x}^* has exactly one nonzero.

For these reasons, this paper investigates the more general case where \mathbf{x}^* is on a *non-minimal facet or interior* of a *non-strongly convex (but convex)* set \mathcal{D} .

Contributions. In this regime, we offer three main theoretical contributions.

- The *continuous-time Frank-Wolfe method*, e.g. the dynamical system of whose explicit Euler discretization gives (FW), has a fundamentally faster convergence rate, which is arbitrarily close to linear convergence rate. This suggests that the fundamental bottleneck in the convergence speed is *slowly decaying discretization error*.
- While multistep methods reduce discretization error and improve the quality of step directions, overall no finite-window averaging method can fundamentally improve the $O(1/k)$ convergence rate.
- Finally, an infinite-window averaging method is introduced, which aggressively attacks the discretization error, and improves the local convergence rate of (FW) to up to $O(1/k^{3/2})$.

The differentiation between local and global convergence is characterized by the identification of the sparse manifold (\bar{k} , where the LMOs of \mathbf{x}_k will always be contained in the potential LMOs of \mathbf{x}^* for all $k \geq \bar{k}$); this follows the local convergence analysis of Liang et al. (2014); Poon et al. (2018); Liang et al. (2017); Nutini et al. (2022, 2019); Hare and Lewis (2004); Sun et al. (2019) over general optimization methods. Overall, these results suggest improved behavior for sparse optimization applications, which is the primary beneficiary of (FW).

1.1 Related works

Continuous-time optimization Recent years have witnessed a surge of research papers connecting dynamical systems with optimization algorithms, generating more intuitive analyses and proposing accelerations. For example, in Su et al. (2016), the Nesterov accelerated gradient descent and Polyak Heavy Ball schemes are shown to be discretizations of a certain second-order ordinary differential equation (ODE), whose tunable vanishing friction pertains to specific parameter choices in the methods.

Multistep discretization methods Inspired by this analysis, several papers (Zhang et al., 2018; Shi et al., 2019) have proposed improvements using advanced discretization schemes; Zhang et al. (2018) uses Runge-Kutta integration methods to improve accelerated gradient methods, and Shi et al. (2019) shows a generalized Leapfrog acceleration scheme which uses a semi-implicit scheme to achieve a very high resolution approximation of the ODE.

Note that for gradient flow, which has the same order of convergence rates as the gradient method, however, it is not likely that multistep methods can offer rate *improvements*. Thus, a unique angle of this work is that in the case of Frank-Wolfe, we will show that indeed, (FWFLOW) has a fundamentally faster convergence rate than (FW). However, with regards to higher order discretization approaches, this work offers a *negative result*, in that no simple explicit higher order discretization scheme fundamentally improves the method’s overall convergence rate. In a way, this is a cautionary tale, suggesting that even when discretization error accounts for fundamental slowness, simple discretization improvements are usually not sufficient.

Accelerated FW. A notable work that highlights the notorious zig-zagging phenomenon is Lacoste-Julien and Jaggi (2015), where an Away-FW method is proposed that cleverly removes offending atoms and improves the search direction. Using this technique, the method is shown to achieve linear convergence under strong convexity of the objective. The tradeoff, however, is that this method requires keeping past atoms, which may incur an undesired memory cost. A work that is particularly complementary to ours is Garber and Hazan (2015), which show an improved $O(1/k^2)$ rate when the *constraint set* is strongly convex—this reduces zigzagging since solutions cannot lie in low-dimensional “flat facets”. Our work addresses the exact opposite regime, where we take advantage of “flat facets” in sparsifying sets (1-norm ball, simplex, etc). This allows the notion of *manifold identification* as determining when suddenly the method behavior improves.

Averaged FW. Several previous works have investigated *gradient averaging* (Zhang et al., 2021; Abernethy and Wang, 2017). While performance seems promising, the rate was not improved past $O(\frac{1}{k})$. Ding et al. (2020) investigates *oracle averaging* by solving small subproblems at each iteration to achieve optimal weights.

Sparse optimization. Other works that investigate *local convergence behavior* include Liang et al. (2014, 2017); Poon et al. (2018); Sun et al. (2019); Iutzeler and Mallick (2020); Nutini et al. (2019); Hare and Lewis (2004). Here, problems which have these two-stage regimes are described as having *partial smoothness*, which allows for the low-dimensional solution manifold to have significance. In our work, we differentiate a local convergence regime of when this manifold is “identified”, e.g. all future LMOs are drawn from vertices of this specific manifold. After this point, we show that convergence of both the proposed flow and method can be improved with a faster decaying discretization error, which in practice may be fast.

2 THE PROBLEM WITH FRANK-WOLFE

2.1 Small examples

There are several different explanations as to why the convergence rate of (FW) appears fundamentally slow, with the predominant theories focused on the bad (zig-zagging) step directions. Therefore we begin by evaluating the vanilla (FW) with simple acceleration schemes: line search (LS), Nesterov’s extrapolation applied to \mathbf{x}_k direction (Nest. Acc.), and Nesterov’s 3-point acceleration as described in Li et al. (2020) (Li Acc). We do this over 5 cases of \mathcal{D} and \mathbf{x}^* , where f is a strongly convex quadratic.

- **Case 1:** \mathcal{D} is a convex polytope, and \mathbf{x}^* is on a minimal facet of \mathcal{D} ; in other words, \mathbf{x}^* is 1-sparse and contains a mixture of exactly 1 atom. In this surprising case, vanilla methods are at least $O(1/k^2)$, Nesterov’s acceleration seems to achieve additional acceleration, and line search works instantaneously. We emphasize that this is a deliberately trivial case.
- **Case 2:** \mathcal{D} is a convex polytope, and \mathbf{x}^* is on a non-minimal facet of \mathcal{D} ; here, \mathbf{x}^* contains a mixture of atoms, more than 1 but less than n . This is the typical “hard” case that we wish to investigate, where no simple acceleration techniques seem to offer any improvement.
- **Case 3:** \mathcal{D} is a convex polytope, and \mathbf{x}^* is in the relative interior of \mathcal{D} ; \mathbf{x}^* contains a full nontrivial mixture of *all* atoms in \mathcal{D} . This case is covered by Guélat and Marcotte (1986), which suggest linear convergence is possible, *but only when line search is used*. In the absence of line search, all methods *still* do poorly.
- **Case 4:** \mathbf{x}^* is on the boundary of \mathcal{D} , which is a *strongly convex set*. That is, there exists $\epsilon > 0$ where, for all x and y in \mathcal{D} , any $0 \leq \theta \leq 1$, any z such that $\|z\|_2 = 1$,

$$\theta x + (1 - \theta)y + \theta(1 - \theta)\frac{\epsilon}{2}\|x - y\|^2 z \in \mathcal{D}.$$

(Vial, 1982). Here, as Garber and Hazan (2015) has shown, the presence of a strongly convex \mathcal{D} allows the iterates to avoid zig-zagging, and in fact even the poorest methods achieve $O(1/k^2)$ without any special tricks, with further acceleration possible via Nesterov’s extrapolation or line search.

- **Case 5:** \mathcal{D} is a strongly convex set, and \mathbf{x}^* is in the relative interior of \mathcal{D} . In fact, this is exactly the same as Case 3: note that the acceleration due to strongly convex \mathcal{D} is also lost in this case, as the atoms acquired during the local convergence phase do not converge to a small neighborhood.

Figure 1 summarizes these observations visually; from the performance plots, it is clear that the open areas (Case 2, and Cases 3 and 5 without line search) indeed has limited slow convergence.

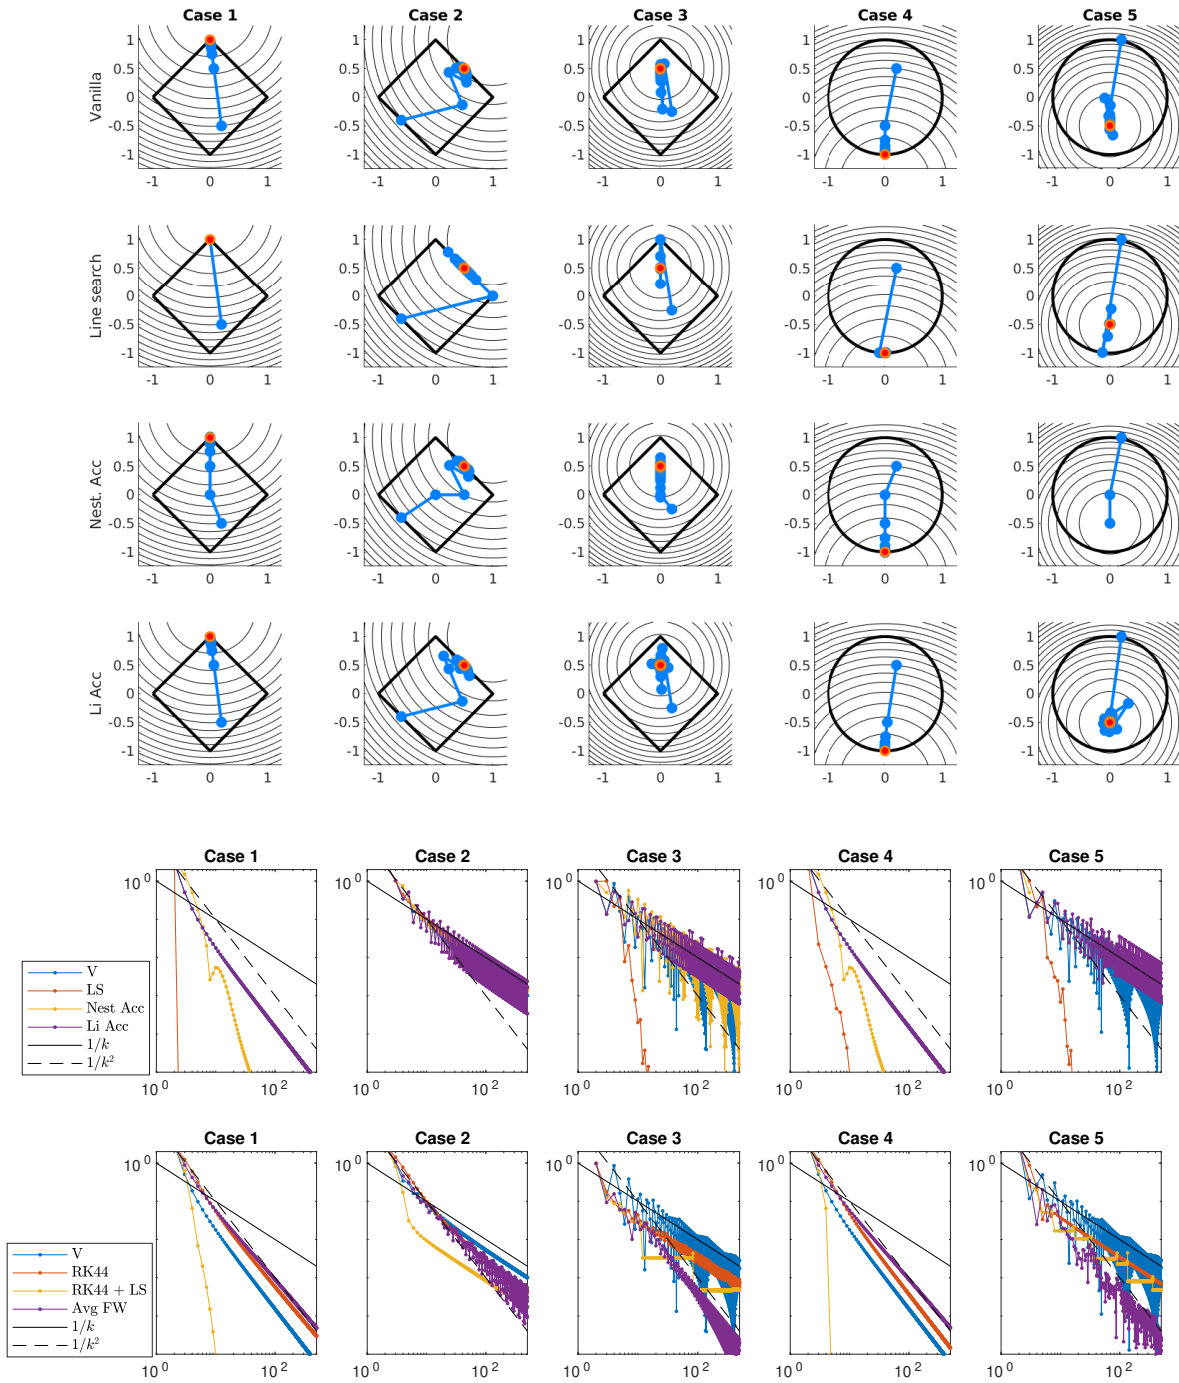


Figure 1: Small case study of 2-D quadratic problem, under different choices of \mathcal{D} . The last 2 rows gives summarizing gap convergence rates (as a function of iterations k) over usual acceleration approaches (upper) and our proposed methods (lower).

2.2 Two different convergence rates

We now investigate these problematic behaviors in terms of the (FWFLOW) convergence behavior, as compared to that of (FW). First, we derive the convergence rate of (FWFLOW). We use the objective value suboptimality as the Lyapunov function:

$$\mathcal{E}(t) = f(x(t)) - f^*, \quad \mathcal{E}_k = f(\mathbf{x}_k) - f^*$$

where $f^* = f(x^*)$ and x^* minimizes (1). Then, using the properties of the LMO and convexity of f ,

$$\begin{aligned} \dot{\mathcal{E}}(t) &= \nabla f(x)^T \dot{x}(t) \\ &= \gamma(t) \underbrace{\nabla f(x)^T (s - x)}_{\text{negative duality gap}} \\ &\leq -\gamma(t)\mathcal{E}(t) \end{aligned} \quad (2)$$

and taking $\gamma(t) = \frac{c}{c+t}$, we have the flow convergence rate

$$\frac{\mathcal{E}(t)}{\mathcal{E}(0)} \leq \frac{c^c}{(c+t)^c} \stackrel{c \rightarrow \pm\infty}{=} \exp(-t).$$

Note that this rate is *arbitrarily close to a linear rate*. In contrast, using L -smoothness, the (FW) method satisfies the recursion

$$f(\mathbf{x}_{k+1}) - f(\mathbf{x}_k) \leq \underbrace{\gamma_k \nabla f(\mathbf{x}_k)^T (\mathbf{x}_k - \mathbf{s}_k)}_{\text{discretization term}} + \underbrace{\frac{L\gamma_k^2}{2} \|\mathbf{x}_k - \mathbf{s}_k\|_2^2}_{=(2D)^2}$$

where $D = \max_{x \in \mathcal{D}} \|x\|_2$. Note that following this line of reasoning, we can at best bound the difference in \mathcal{E}_k as

$$\mathcal{E}_{k+1} - \mathcal{E}_k \leq -\gamma_k \mathcal{E}_k + 2LD\gamma_k^2 \quad (3)$$

which recursively gives a bound of $\mathcal{E}_k = O(\frac{c}{c+k}) = O(1/k)$. Importantly, this analysis is *tight in general*.

Note the key difference in (2) and its analogous terms in (3) is the extra $2LD^2\gamma_k^2 = O(1/k^2)$ term, which throttles the recursion from doing better than $\mathcal{E}_k = O(1/k)$. One may ask if it is possible to bypass this problem by simply picking γ_k decaying more aggressively; however, then such a sequence becomes summable, and then convergence of $\mathbf{x}_k \rightarrow \mathbf{x}^*$ will not be assured. Therefore, we must have a sequence γ_k converging *at least* as slowly as $O(1/k)$. Thus, the primary culprit in this convergence rate tragedy is the bound $\|\mathbf{s}_k - \mathbf{x}_k\|_2 = O(D)$ (nondecaying), which forces the discretization term to decay no faster than $O(1/k)$. As shown in our case studies, this is not a loose bound in general.

In this paper, we investigate methods that push the (FW) method more toward its (FWFLOW) trajectory, using some

form of averaging. First, we use the specific weight updates offered by Runge-Kutta multistep schemes, which reduce discretization errors by constant factors, and additionally seem to improve step direction quality. Second, we propose an averaged LMO method, which gives an overall improved local convergence rate of up to $O(1/k^{3/2})$.

3 RUNGE-KUTTA MULTISTEP METHOD

3.1 The generalized Runge-Kutta family

We now look into multistep methods that better imitate the continuous flow by reducing discretization error. Observe that the standard (FW) algorithm is equivalent to the discretization of (FWFLOW) by the explicit Euler's method with step size $\Delta = 1$. It is well known that the discretization error associated with this scheme is $O(\Delta^q)$ with $q = 1$, e.g. it is a method of order 1.

We now consider Runge-Kutta (RK) methods, a generalized class of higher order methods ($q \geq 1$). Many commonly-used discretization schemes are of the Runge-Kutta family; for example, Dormand–Prince (RK45) is the default ode solver used by MATLAB. Other examples of RK-methods include the Fehlberg (RK23) and Cash–Karp methods. These methods are fully parametrized by some choice of $A \in \mathbb{R}^{q \times q}$, $\beta \in \mathbb{R}^q$, and $\omega \in \mathbb{R}^q$ and at step k can be expressed as (for $i = 1, \dots, q$)

$$\begin{aligned} \xi_i &= \dot{x}((k + \omega_i)\Delta, \mathbf{x}_k + \sum_{j=1}^q A_{ij}\xi_j), \\ \mathbf{x}_{k+1} &= \mathbf{x}_k + \sum_{i=1}^q \beta_i \xi_i. \end{aligned} \quad (4)$$

For consistency, $\sum_i \beta_i = 1$, and to maintain explicit implementations, A is always strictly lower triangular. As a starting point, $\omega_1 = 0$. Referring to the description of $\dot{x}(t)$ as given in (FWFLOW), then given a matrix $A \in \mathbb{R}^{q \times q}$ and vectors $\beta, \omega \in \mathbb{R}^q$, the sequence described in (4) describes the q -stage multistep Frank-Wolfe (MULTFW) method. (See also Appendix C.) Figure 3 compares pictorially the trajectory of the vanilla (FW) method with (MULTFW), which after averaging has a far more controlled and less erratic trajectory. We hope to leverage this into better convergence behavior.

Proposition 3.1 (Feasibility). *For a given q -stage multistep method defined by A , β , and ω , for each given $k \geq 1$, define*

$$\bar{\gamma}_i^{(k)} = \frac{c}{c+k+\omega_i}, \quad \Gamma^{(k)} = \text{diag}(\bar{\gamma}_i^{(k)}),$$

$$\mathbf{P}^{(k)} = \Gamma^{(k)}(I + A^T \Gamma^{(k)})^{-1}, \quad \mathbf{z}^{(k)} = q\mathbf{P}^{(k)}\beta.$$

Then if $0 \leq \mathbf{z}^{(k)} \leq 1$ for all $k \geq 1$, then

$$\mathbf{x}_0 \in \mathcal{D} \Rightarrow \mathbf{x}_k \in \mathcal{D}, \quad \forall k \geq 1.$$

The proof is in Appendix D. This condition can be checked explicitly and is true of almost all RK methods with a notable exception of the midpoint method, where $\mathbf{z}_i^{(k)} < 0$ is possible. A full table of weights for various RK methods is in Appendix B.

3.2 Convergence of (MULTFW)

We first establish that using a generalized Runge-Kutta method cannot *hurt* convergence, as compared to the usual Frank-Wolfe method.

Proposition 3.2. *All Runge-Kutta methods converge at worst with rate $f(\mathbf{x}_k) - f(\mathbf{x}^*) \leq O(1/k)$.*

The proof is in Appendix D.2. That is to say, (MULTFW) cannot be an order *slower* than vanilla (FW).¹

It is left to ask if (MULTFW) can *improve* upon (FW). Recall that (FWFLOW) achieves a rate of $O(1/t^c)$ rate, and taking $c \rightarrow +\infty$ achieves a linear rate. This is also verified numerically in Figure 2; larger c provides a sharper convergence rate. It is hence tempting to think that increasing c can help (FW) methods in general, and in particular by adapting a higher order multistep method, we can overcome the problems caused by discretization errors.

Lower bound derivation: a toy problem. Let us now consider a simple bounded optimization problem over scalar variables $\mathbf{x} \in \mathbb{R}$:

$$\min_{\mathbf{x}} f(\mathbf{x}) \quad \text{s. t.} \quad -1 \leq \mathbf{x} \leq 1 \quad (5)$$

where

$$f(\mathbf{x}) = \begin{cases} \mathbf{x}^2/2 & \text{if } |\mathbf{x}| < \varepsilon \\ \varepsilon\mathbf{x} - \varepsilon^2/2 & \text{if } \mathbf{x} \geq \varepsilon \\ -\varepsilon\mathbf{x} - \varepsilon^2/2 & \text{if } \mathbf{x} \leq -\varepsilon \end{cases}$$

which is a scaled version of the Huber norm applied to scalars. By design, no matter how small ε is, f is L -smooth and 1-Lipschitz. Additionally, $\text{LMO}_{[-1,1]}(\mathbf{x}) = -\text{sign}(\mathbf{x})$.

Proposition 3.3. *Assuming that $0 < q\mathbb{P}^{(k)}\beta < 1$ for all k . Start with $\mathbf{x}_0 = 1$. Suppose the choice of $\beta \in \mathbb{R}^p$ is not “cancellable”; that is, there exist no partition $S_1 \cup S_2 = \{1, \dots, p\}$ where*

$$\sum_{i \in S_1} \beta_i - \sum_{j \in S_2} \beta_j = 0.$$

Then regardless of the order p and choice of A , β , and ω , as long as $\sum_i \beta_i = 1$, then

$$\sup_{k' > k} |\mathbf{x}_k| = \Omega(1/k).$$

¹It should be noted, however, that the convergence rate in terms of k does not account for the extra factor of q gradient calls needed for a q -stage method. While this may be burdensome, it does not increase the *order* of convergence rate.

That is, the tightest upper bound is $O(1/k)$.

The proof is in Appendix D.3. The assumption of a “non-cancellable” choice of β_i may seem strange, but in fact it is true for most of the higher order Runge-Kutta methods. More importantly, the assumption doesn’t matter in practice; even if we force β_i ’s to be all equal, our numerical experiments do not show much performance difference in this toy problem. (Translation: do not design your Runge Kutta method for the β ’s to be cancel-able in hopes of achieving a better rate!)

Proposition 3.3 implies a $\Omega(1/k)$ bound on $|\mathbf{x}_k|$. To extend it to an $\Omega(1/k)$ bound on $f(\mathbf{x}_k) - f^*$, note that whenever $|\mathbf{x}| \geq \varepsilon$,

$$\frac{f(\mathbf{x}_k) - f^*}{f(\mathbf{x}_0) - f^*} \geq \frac{|\mathbf{x}_k|}{2|\mathbf{x}_0|}$$

for ε arbitrarily small.

Corollary 3.4. *Although higher order RK methods may provide constant factor improvements, the worst best case bound for (MULTFW), for any RK method, is of order $O(1/k)$.*

Although this negative result is disappointing, it is valuable to know that simple multistep enhancements of traditional optimization methods may not single-handily produce miraculously better convergence rates.

4 AVERAGED FRANK-WOLFE METHOD

We now propose an infinite-window LMO-averaging Frank-Wolfe (AvgFW) method, by replacing \mathbf{s}_k with an averaged version $\bar{\mathbf{s}}_k$.

$$\begin{aligned} \mathbf{s}_k &= \text{LMO}_{\mathcal{D}}(\mathbf{x}_k) \\ \bar{\mathbf{s}}_k &= \bar{\mathbf{s}}_{k-1} + \beta_k(\mathbf{s}_k - \bar{\mathbf{s}}_{k-1}) \\ \mathbf{x}_{k+1} &= \mathbf{x}_k + \gamma_k(\bar{\mathbf{s}}_k - \mathbf{x}_k) \end{aligned} \quad (\text{AVGFW})$$

where $\beta_k = (\frac{c}{c+k})^p$ for $c \geq 0$ and $0 < p \leq 1$.² Here, the smoothing in $\bar{\mathbf{s}}_k$ has two roles. First, averaging reduces zigzagging; at every k , $\bar{\mathbf{s}}_k$ is a convex combination of past \mathbf{s}_k , and has a smoothing effect that qualitatively also reduces zig-zagging; this is of importance should the user wish to use line search or momentum-based acceleration. Second, this forces the discretization term $\|\mathbf{x}_k - \bar{\mathbf{s}}_k\|_2$ to *decay*, allowing for faster numerical performance. Figure 3 shows a simple 2D example of our (AVGFW) method (right) compared to the usual (FW) method (left) and the Runge-Kutta multistep version (MULTFW) (center). Note that (AVGFWFLOW) is fundamentally different than (FWFLOW), and is better mirrored by (AVGFW).

²Recall that $\gamma_k = c/(c+k)$. While it is possible to use $\beta_k = b/(b+k)$ with $b \neq c$ in practice, our proofs considerably simplify when the constants are the same.




Figure 2: **Continuous vs discrete.** A comparison of the numerical error vs compared with derived rate. The black curve shows the upper bound on the flow rate, compared against simulated method rates for smaller discretization units. The two-stage behavior of the curves is intriguing, as it seems there is a fundamental point where discretization error takes over, and forces the $O(1/k)$ rate to manifest.




Figure 3: Panels showing trajectory behavior of (FW) (left), (MULTFW) (center), and (AVGFW) (right). The method trajectories (blue) and underlying flow (red) are both shown.

Using similar tricks as in Jacimovic and Geary (1999), we may view (AVGFW) as an Euler discretization (with $\Delta = 1$) of the following dynamical system

$$\begin{aligned} s(t) &= \text{LMO}_{\mathcal{D}}(x(t)) \\ \dot{s}(t) &= \beta(t)(s(t) - \bar{s}(t)) \quad , \quad (\text{AVGFWFLOW}) \\ \dot{x}(t) &= \gamma(t)(\bar{s}(t) - x(t)) \end{aligned}$$

where $\gamma(t)$, $\beta(t)$ are such that $\gamma_k = \gamma(k\Delta)$ and $\beta_k = \beta(k\Delta)$.

4.1 Global convergence

We start by showing the method converges, with no assumptions on sparsity or manifold identification. We note that in practice we see faster convergence of (AVGFW) compared to (FW), despite weaker theoretical rates.

Theorem 4.1 (Global rates). *Take $0 < p < 1$, and assume f is L -smooth and μ -strongly convex.*

- For $\beta(t) = \left(\frac{c}{c+t}\right)^p$, the flow (AVGFWFLOW) satisfies $f(x(t)) \leq O\left(\frac{1}{t^{1-p}}\right)$.
- For $\beta_k = \left(\frac{c}{c+k}\right)^p$, the method (AVGFW) satisfies $f(\mathbf{x}_k) \leq O\left(\frac{1}{k^p}\right)$.

The proofs are in Appendix E. Note that the rates are not exactly reciprocal. In terms of analysis, the method is that

allows the term $\beta_k \text{gap}(\mathbf{x}_k)$ to take the weight of an entire step, whereas in the flow, the corresponding term is infinitesimally small. This is a curious disadvantage of the flow analysis; since most of our progress is accumulated in the last step, the method exploits this more readily than the flow, where the step is infinitesimally small.

4.2 Manifold identification

Now that we know that the method converges, we may discuss the local convergence regime, characterized by manifold identification. Specifically, the *manifold is identified* at \bar{k} if any $k > \bar{k}$, any $\text{LMO}_{\mathcal{D}}(\mathbf{x}_k)$ is also an $\text{LMO}_{\mathcal{D}}(\mathbf{x}^*)$. For example, when \mathcal{D} is the one-norm ball, then the manifold is identified when for all $k \geq \bar{k}$, \mathbf{s}_k are 1-hot vectors with nonzero supports contained within the sparsity of \mathbf{x}^* . This is *guaranteed to happen for some finite \bar{k}* under mild degeneracy assumptions, and is the basis for gap-safe screening procedures (Ndiaye et al., 2017; Sun and Bach, 2020).

4.3 Accelerated local convergence

Theorem 4.2 (Local rate). *Assume f is L -smooth and μ -strongly convex. After manifold identification, the flow (AVGFWFLOW) satisfies*

$$f(x(t)) - f(x^*) \leq \frac{\log(t)}{t^c}.$$

The proof is in Appendix E.4.

Theorem 4.3 (Local rate). *Assume that f is μ -strongly convex, and pick $c \geq 3p/2 + 1$. After manifold identification, the method (AVGFW) satisfies*

$$f(\mathbf{x}_k) - f(\mathbf{x}^*) = O(1/k^{3p/2}).$$

The proof is in Appendix E.3. Although the proof requires $p < 1$, in practice, we often use $p = 1$ and observe about a $O(1/k^{3/2})$ rate, regardless of strong convexity.

5 NUMERICAL EXPERIMENTS

We now compare the Frank-Wolfe method (FW), its multistep form (MULTFW), and its averaged form (AVGFW), against two baseline methods:

- the averaged gradient method of Abernethy and Wang (2017),
- and the away-step method of Lacoste-Julien and Jaggi (2015).

In all cases, when appropriate, we also add line search (LS). Note that while the away-step method is theoretically a superior method to those we propose, it requires some computational overhead in maintaining and readjusting a full history of past atoms, and also always requires line search to ensure feasibility. Since MULTFW requires compute additional gradients, our x -axis gives the number of gradient calls, instead of iterations, to ensure fair comparison.

Simulated compressive sensing. In Figure 4 we simulate compressive sensing using a ℓ_1 -norm ball constraint. Given $\mathbf{x}_0 \in \mathbb{R}^m$, a sparse ground truth vector with 10% nonzeros, and given $A \in \mathbb{R}^{n \times m}$ with entries i.i.d. Gaussian, we generate $\mathbf{y} = A\mathbf{x}_0$ (noiseless observation). The problem formulation is

$$\min_{x \in \mathbb{R}^n} \frac{1}{2} \|Ax - y\|_2^2 \quad \text{subject to} \quad \|x\|_1 \leq \alpha. \quad (6)$$




Figure 4: **Compressed sensing.** $n = m = 500$, $c = \beta = 2$, $p = 1$.

Signed compressive sensing. In Fig. 5, under the same parameter distribution as in the previous example, we now use a sampling model where $y_i = \text{sign}(a_i^T \mathbf{x}_0)$, and attempt to recover the sparsity of \mathbf{x}_0 using sparse logistic regression

$$\min_{x \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m \log(1 + \exp(-y_i a_i^T x)) \quad \text{s.t.} \quad \|x\|_1 \leq \alpha. \quad (7)$$

Because signed compressive sensing is a harder task, we use a higher sampling ratio of 50x.

6 CONCLUSION

The main goal of this study is to see if the discretization error, which seems to be the primary cause of slow (FW)




Figure 5: **Signed compressed sensing.** $m = 5000$, $n = 100$, $c = \beta = 2$, $p = 1$.

convergence, can be attacked directly in order to produce a better method. First, we highlight the artifacts and pitfalls of discretization error (largely, bad zig-zagging step directions that are not amenable to momentum-based acceleration) and the regimes where they are most prominent (on the non-minimal boundary of a polytope, or in the interior of a set when line search is not used).

Next, we explore the use of *higher order discretization methods*, which has been used in the past in gradient flow applications with observed success, and which has fundamentally better truncation errors with each growing order. However, we show that though this method gives constant factor improvements in practice, asymptotically the worst case convergence rate does not improve beyond $O(1/k)$. This is consistent with simulations that show initially fast convergence, but a “tapering off” once the residual convergence error is below that of the truncation error order guarantee. While this result is disappointing, it is important to highlight the main flaw of directly adapting this popular dynamical systems tool for optimization improvement.

Finally, we explore an LMO-averaging method, which produces a smoother convergence trajectory and a fundamentally improved convergence rate (from $O(1/k)$ to up to $O(1/k^{3/2})$) with negligible computation and memory overhead. This method is largely inspired from viewing the multistep method as a finite-window averaging method, which achieves finite constant order rate improvement—thus, only an infinite window seems to allow for non-constant order improvement. Our numerical results show that, though our theoretical improvements are only local, the effect of this acceleration appears effective globally; moreover, manifold identification (or at least reduction to a small working set of nonzeros) appears almost immediately in many cases.

Use in gap-safe screening rules. An important related application to these works is gap-safe convergence rates Ndiaye et al. (2017); Sun and Bach (2020), which produce sparsity guarantees given intermediate iterates \mathbf{x}_k , provided the duality gap at \mathbf{x}_k is below some problem-dependent constant. Typically, these guarantees are strong, but slow to realize in practice, especially when gap convergence is slow. In this, (AVGFW) offers a distinct advantage in its faster gap convergence, in providing this guarantee.

References

- Abernethy, J. D. and Wang, J.-K. (2017). On Frank-Wolfe and equilibrium computation. *Advances in Neural Information Processing Systems*, 30.
- Canon, M. and Cullum, C. (1968). A tight upper bound on the rate of convergence of frank-wolfe algorithm. *Siam Journal on Control*, 6:509–516.
- Diakonikolas, J. and Orecchia, L. (2019). The approximate duality gap technique: A unified theory of first-order methods. *SIAM Journal on Optimization*, 29(1):660–689.
- Ding, L., Fan, J., and Udell, M. (2020). *kfw*: A frank-wolfe style algorithm with stronger subproblem oracles. *arXiv preprint arXiv:2006.16142*.
- Dunn, J. C. and Harshbarger, S. (1978). Conditional gradient algorithms with open loop step size rules. *Journal of Mathematical Analysis and Applications*, 62(2):432–444.
- Frank, M., Wolfe, P., et al. (1956). An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110.
- Freund, R., Grigas, P., and Mazumder, R. (2017). An extended frank-wolfe method with “in-face” directions, and its application to low-rank matrix completion. *SIAM J. Optim.*, 27:319–346.
- Freund, R. M. and Grigas, P. (2016). New analysis and results for the frank–wolfe method. *Mathematical Programming*, 155(1-2):199–230.
- Garber, D. and Hazan, E. (2015). Faster rates for the frank-wolfe method over strongly-convex sets. In *International Conference on Machine Learning*, pages 541–549. PMLR.
- Guélat, J. and Marcotte, P. (1986). Some comments on wolfe’s ‘away step’. *Mathematical Programming*, 35(1):110–119.
- Guyon, I., Gunn, S. R., Ben-Hur, A., and Dror, G. (2004). Result analysis of the nips 2003 feature selection challenge. In *NIPS*, volume 4, pages 545–552.
- Hare, W. L. and Lewis, A. S. (2004). Identifying active constraints via partial smoothness and prox-regularity. *Journal of Convex Analysis*, 11(2):251–266.
- Iutzeler, F. and Malick, J. (2020). Nonsmoothness in machine learning: specific structure, proximal identification, and applications. *Set-Valued and Variational Analysis*, 28(4):661–678.
- Jacimovic, M. and Geary, A. (1999). A continuous conditional gradient method. *Yugoslav journal of operations research*, 9(2):169–182.
- Jaggi, M. (2013). Revisiting frank-wolfe: Projection-free sparse convex optimization. In *International Conference on Machine Learning*, pages 427–435. PMLR.
- Joulin, A., Tang, K. D., and Fei-Fei, L. (2014). Efficient image and video co-localization with frank-wolfe algorithm. In *ECCV*.
- Krasovskii, N. N. (1968). Regularization of a problem on the encounter of motions. In *Doklady Akademii Nauk*, volume 179, pages 300–303. Russian Academy of Sciences.
- Krishnan, R., Lacoste-Julien, S., and Sontag, D. (2015). Barrier frank-wolfe for marginal inference. In *NIPS*.
- Lacoste-Julien, S. and Jaggi, M. (2015). On the global linear convergence of frank-wolfe optimization variants. In *NIPS*.
- Lacoste-Julien, S., Jaggi, M., Schmidt, M., and Pletscher, P. (2013). Block-coordinate frank-wolfe optimization for structural svms. In *ICML*.
- Levitin, E. and Polyak, B. (1966). Constrained minimization methods. *USSR Computational Mathematics and Mathematical Physics*, 6(5):1–50.
- Li, B., Coutino, M., Giannakis, G. B., and Leus, G. (2020). How does momentum help frank wolfe? *arXiv preprint arXiv:2006.11116*.
- Liang, J., Fadili, J., and Peyré, G. (2014). Local linear convergence of forward–backward under partial smoothness. *Advances in neural information processing systems*, 27.
- Liang, J., Fadili, J., and Peyré, G. (2017). Local convergence properties of Douglas–Rachford and alternating direction method of multipliers. *Journal of Optimization Theory and Applications*, 172(3):874–913.
- Ndiaye, E., Fercoq, O., Gramfort, A., and Salmon, J. (2017). Gap safe screening rules for sparsity enforcing penalties. *The Journal of Machine Learning Research*, 18(1):4671–4703.
- Nutini, J., Laradji, I., and Schmidt, M. (2022). Let’s make block coordinate descent converge faster: Faster greedy rules, message-passing, active-set complexity, and superlinear convergence. *Journal of Machine Learning Research*, 23(131):1–74.
- Nutini, J., Schmidt, M., and Hare, W. (2019). “Active-set complexity” of proximal gradient: How long does it take to find the sparsity pattern? *Optimization Letters*, 13(4):645–655.
- Poon, C., Liang, J., and Schoenlieb, C. (2018). Local convergence properties of SAGA/Prox-SVRG and acceleration. In *International Conference on Machine Learning*, pages 4124–4132. PMLR.
- Sanfelice, R. G., Goebel, R., and Teel, A. R. (2008). Generalized solutions to hybrid dynamical systems. *ESAIM: Control, Optimisation and Calculus of Variations*, 14(4):699–724.
- Shi, B., Du, S. S., Su, W. J., and Jordan, M. I. (2019). Acceleration via symplectic discretization of high-resolution differential equations. *arXiv preprint arXiv:1902.03694*.

- Su, W., Boyd, S., and Candes, E. J. (2016). A differential equation for modeling nesterov’s accelerated gradient method: Theory and insights. *The Journal of Machine Learning Research*, 17(1):5312–5354.
- Sun, Y. and Bach, F. (2020). Safe screening for the generalized conditional gradient method. *arXiv preprint arXiv:2002.09718*.
- Sun, Y., Jeong, H., Nutini, J., and Schmidt, M. (2019). Are we there yet? manifold identification of gradient-related proximal methods. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1110–1119. PMLR.
- Vial, J.-P. (1982). Strong convexity of sets and functions. *Journal of Mathematical Economics*, 9(1-2):187–205.
- Zhang, J., Mokhtari, A., Sra, S., and Jadbabaie, A. (2018). Direct runge-kutta discretization achieves acceleration. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Zhang, Y., Li, B., and Giannakis, G. B. (2021). Accelerating frank-wolfe with weighted average gradients. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5529–5533.

A ADDITIONAL EXPERIMENTS

We include a few more experiments to expand beyond Runge-Kutta family (see Figure 6). In general, we do not notice much difference in MULTFW with different multistep methods. It is consistent with our theory, which states that MULTFW is still $\Omega(1/k)$.




Figure 6: Additional multistep methods

B EXTRA TABLEAUS OF HIGHER ORDER RUNGE KUTTA METHODS

- Midpoint method

$$A = \begin{bmatrix} 0 & 0 \\ 1/2 & 0 \end{bmatrix}, \quad \beta = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad \omega = \begin{bmatrix} 0 \\ 1/2 \end{bmatrix}, \quad \mathbf{z}^{(1)} \approx \begin{bmatrix} -0.3810 \\ 1.1429 \end{bmatrix}, \quad \mathbf{z}^{(2)} \approx \begin{bmatrix} -0.2222 \\ 0.8889 \end{bmatrix}$$

- Runge Kutta 4th Order Tableau (44)

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad \beta = \begin{bmatrix} 1/6 \\ 1/3 \\ 1/3 \\ 1/6 \end{bmatrix}, \quad \omega = \begin{bmatrix} 0 \\ 1/2 \\ 1/2 \\ 1 \end{bmatrix}, \quad \mathbf{z}^{(1)} \approx \begin{bmatrix} 0.2449 \\ 0.5986 \\ 0.5714 \\ 0.3333 \end{bmatrix}$$

- Runge Kutta 3/8 Rule Tableau (4)

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1/3 & 0 & 0 & 0 \\ -1/3 & 1 & 0 & 0 \\ 1 & -1 & 1 & 0 \end{bmatrix}, \quad \beta = \begin{bmatrix} 1/8 \\ 3/8 \\ 3/8 \\ 1/8 \end{bmatrix}, \quad \omega = \begin{bmatrix} 0 \\ 1/3 \\ 2/3 \\ 1 \end{bmatrix}, \quad \mathbf{z}^{(1)} \approx \begin{bmatrix} 0.1758 \\ 0.6409 \\ 0.6818 \\ 0.2500 \end{bmatrix}$$

- Runge Kutta 5 Tableau

$$A = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 1/4 & 0 & 0 & 0 & 0 & 0 \\ 1/8 & 1/8 & 0 & 0 & 0 & 0 \\ 0 & -1/2 & 1 & 0 & 0 & 0 \\ 3/16 & 0 & 0 & 9/16 & 0 & 0 \\ -3/7 & 2/7 & 12/7 & -12/7 & 8/7 & 0 \end{bmatrix}, \quad \beta = \begin{bmatrix} 7/90 \\ 0 \\ 32/90 \\ 12/90 \\ 32/90 \\ 7/90 \end{bmatrix}, \quad \omega = \begin{bmatrix} 0 \\ 1/4 \\ 1/4 \\ 1/2 \\ 3/4 \\ 1 \end{bmatrix}, \quad \mathbf{z}^{(1)} \approx \begin{bmatrix} 0.1821 \\ 0.0068 \\ 0.8416 \\ 0.3657 \\ 0.9956 \\ 0.2333 \end{bmatrix}$$

In all examples, $\|\mathbf{z}^{(k)}\|_\infty$ monotonically decays with k .




Figure 7: **Zig-zagging on real datasets.** Average deviation of different discretizations of FWFLOW. Top table uses different Δ s and uses the vanilla Euler’s discretization (FW). Bottom uses $\Delta = 1$ and different multistep methods. The two numbers in each box correspond to window sizes 5 / 20.

C EXTRA EXPERIMENTS

Figure 7 quantifies this notion more concretely. We measure “zig-zagging energy” by averaging the deviation of each iterate’s direction across k -step directions, for $k = 1, \dots, W$, for some measurement window W :

$$\mathcal{E}_{\text{zigzag}}(\mathbf{x}^{(k+1)}, \dots, \mathbf{x}^{(k+W)}) = \frac{1}{W-1} \sum_{i=k+1}^{k+W-1} \left\| \underbrace{\left(I - \frac{1}{\|\bar{\mathbf{d}}^{(k)}\|_2} \bar{\mathbf{d}}^{(k)} (\bar{\mathbf{d}}^{(k)})^T \right)}_{\mathbb{Q}} \mathbf{d}^{(i)} \right\|_2,$$

where $\mathbf{d}^{(i)} = \mathbf{x}^{(i+1)} - \mathbf{x}^{(i)}$ is the current iterate direction and $\bar{\mathbf{d}}^{(k)} = \mathbf{x}^{(k+W)} - \mathbf{x}^{(k)}$ a “smoothed” direction. The projection operator \mathbb{Q} removes the component of the current direction in the direction of the smoothed direction, and we measure this “average deviation energy.” We divide the trajectory into these window blocks, and report the average of these measurements $\mathcal{E}_{\text{zigzag}}$ over $T = 100$ time steps (total iteration = T/Δ). Figure 7 (top table) exactly shows this behavior, where the problem is sparse constrained logistic regression minimization over several machine learning classification datasets (Guyon et al., 2004) (Sensing (ours), Gisette ³ and Madelon ⁴) are shown in Fig. 7.

D MULTISTEP FRANK WOLFE PROOFS

D.1 Feasibility

Proposition D.1. For a given q -stage RK method defined by A , β , and ω , for each given $k \geq 1$, define

$$\begin{aligned} \bar{\gamma}_i^{(k)} &= \frac{c}{c+k+\omega_i}, & \Gamma^{(k)} &= \mathbf{diag}(\bar{\gamma}_i^{(k)}), \\ \mathbf{P}^{(k)} &= \Gamma^{(k)}(I + A^T \Gamma^{(k)})^{-1}, & \mathbf{z}^{(k)} &= q\mathbf{P}^{(k)}\beta. \end{aligned}$$

Then if $0 \leq \mathbf{z}^{(k)} \leq 1$ for all $k \geq 1$, then

$$\mathbf{x}_0 \in \mathcal{D} \Rightarrow \mathbf{x}_k \in \mathcal{D}, \quad \forall k \geq 1.$$

Proof of Prop. D.1.

Proof. For a given k , construct additionally

$$\mathbf{Z} = [\xi_1 \quad \xi_2 \quad \cdots \quad \xi_q],$$

³Full dataset available at <https://archive.ics.uci.edu/ml/datasets/Gisette>. We use a subsampling, as given in <https://github.com/cyrillewcombettes/boostfw>.

⁴Dataset: <https://archive.ics.uci.edu/ml/datasets/madelon>

$$\bar{\mathbf{X}} = [\bar{\mathbf{x}}_1 \quad \bar{\mathbf{x}}_2 \quad \cdots \quad \bar{\mathbf{x}}_q], \quad \bar{\mathbf{S}} = [\bar{\mathbf{s}}_1 \quad \bar{\mathbf{s}}_2 \quad \cdots \quad \bar{\mathbf{s}}_q].$$

where

$$\begin{aligned} \bar{\mathbf{x}}_i &= \mathbf{x}_k + \sum_{j=1}^q A_{ij} \xi_j, \\ \bar{\mathbf{s}}_i &= \text{LMO}(\bar{\mathbf{x}}_i). \end{aligned}$$

Then we can rewrite (4) as

$$\begin{aligned} \mathbf{Z} &= (\bar{\mathbf{S}} - \bar{\mathbf{X}})\Gamma = (\bar{\mathbf{S}} - \mathbf{x}_k \mathbf{1}^T - \mathbf{Z}A^T)\Gamma \\ &= (\bar{\mathbf{S}} - \mathbf{x}_k \mathbf{1}^T)\mathbf{P} \end{aligned}$$

for shorthand $\mathbf{P} = \mathbf{P}^{(k)}$ and $\Gamma = \Gamma^{(k)}$. Then

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{x}_k(1 - \mathbf{1}^T \mathbf{P} \beta) + \bar{\mathbf{S}} \mathbf{P} \beta \\ &= \frac{1}{q} \sum_{i=1}^q \underbrace{(1 - \mathbf{z}_i^{(k)}) \mathbf{x}_k + \mathbf{z}_i^{(k)} \bar{\mathbf{s}}_i}_{\hat{\xi}_i} \end{aligned}$$

where $\mathbf{z}_i^{(k)}$ is the i th element of $\mathbf{z}^{(k)}$, and $\beta = (\beta_1, \dots, \beta_q)$. Then if $0 \leq \mathbf{z}_i^{(k)} \leq 1$, then $\hat{\xi}_i$ is a convex combination of \mathbf{x}_k and $\bar{\mathbf{s}}_i$, and $\hat{\xi}_i \in \mathcal{D}$ if $\mathbf{x}_k \in \mathcal{D}$. Moreover, \mathbf{x}_{k+1} is an average of $\hat{\xi}_i$, and thus $\mathbf{x}_{k+1} \in \mathcal{D}$. Thus we have recursively shown that $\mathbf{x}_k \in \mathcal{D}$ for all k . \square

D.2 Positive Runge-Kutta convergence result

Lemma D.2. *After one step, the generalized Runge-Kutta method satisfies*

$$h(\mathbf{x}_{k+1}) - h(\mathbf{x}_k) \leq -\gamma_{k+1} h(\mathbf{x}_k) + D_4 (\gamma_{k+1})^2$$

where $h(\mathbf{x}) = f(\mathbf{x}) - f(\mathbf{x}^*)$ and

$$D_4 = \frac{LD_2^2 + 2LD_2D_3 + 2D_3}{2}, \quad D_2 = c_1D, \quad D_3 = c_2c_1D, \quad c_1 = qp_{\max}, \quad c_2 = q \max_{ij} |A_{ij}|, \quad D = \text{diam}(\mathcal{D}).$$

Proof. For ease of notation, we write $\mathbf{x} = \mathbf{x}_k$ and $\mathbf{x}^+ = \mathbf{x}_{k+1}$. We will use $\gamma = \gamma_k = \frac{c}{c+k}$, and $\tilde{\gamma}_i = \frac{c}{c+k+\omega_i}$. Now consider the generalized RK method

$$\begin{aligned} \bar{\mathbf{x}}_i &= \mathbf{x} + \sum_{j=1}^q A_{ij} \xi_j \\ \xi_i &= \underbrace{\frac{c}{c+k+\omega_i}}_{\tilde{\gamma}_i} (\mathbf{s}_i - \bar{\mathbf{x}}_i) \\ \mathbf{x}^+ &= \mathbf{x} + \sum_{i=1}^q \beta_i \xi_i \end{aligned}$$

where $\mathbf{s}_i = \text{LMO}(\bar{\mathbf{x}}_i)$.

Define $D = \text{diam}(\mathcal{D})$. We use the notation from section 3. Denote the $2, \infty$ -norm as

$$\|A\|_{2, \infty} = \max_j \|a_j\|_2$$

where a_j is the j th column of A . Note that all the element-wise elements in

$$\mathbf{P}^{(k)} = \Gamma^{(k)}(I + A^T \Gamma^{(k)})^{-1}$$

is a decaying function of k , and thus defining $p_{\max} = \|\mathbf{P}^{(1)}\|_{2,\infty}$ we see that

$$\|\bar{\mathbf{Z}}\|_{2,\infty} = \|(\bar{\mathbf{S}} - \mathbf{x}^{(k)} \mathbf{1})\mathbf{P}^{(k)}\|_{2,\infty} \leq qp_{\max}D.$$

Therefore, since $\bar{\mathbf{Z}} = (\bar{\mathbf{S}} - \bar{\mathbf{X}})\Gamma$, and all the diagonal elements of Γ are at most 1,

$$\|\mathbf{s}_i - \bar{\mathbf{x}}_i\|_2 \leq qp_{\max}D =: D_2$$

and

$$\|\mathbf{x} - \bar{\mathbf{x}}_i\|_2 = \left\| \sum_{j=1}^q A_{ij} \gamma_j (\mathbf{s}_j - \bar{\mathbf{x}}_j) \right\|_2 \leq q \max_{ij} |A_{ij}| \gamma D_2 =: D_3 \gamma.$$

Then

$$\begin{aligned} f(\mathbf{x}^+) - f(\mathbf{x}) &\leq \nabla f(\mathbf{x})^T (\mathbf{x}^+ - \mathbf{x}) + \frac{L}{2} \|\mathbf{x}^+ - \mathbf{x}\|_2^2 \\ &= \sum_i \beta_i \tilde{\gamma}_i \nabla f(\mathbf{x})^T (\mathbf{s}_i - \bar{\mathbf{x}}_i) + \frac{L}{2} \underbrace{\left\| \sum_i \beta_i \tilde{\gamma}_i (\mathbf{s}_i - \bar{\mathbf{x}}_i) \right\|_2^2}_{\leq \gamma^2 D_2^2} \\ &= \sum_i \beta_i \tilde{\gamma}_i (\nabla f(\mathbf{x}) - \nabla f(\bar{\mathbf{x}}_i))^T (\mathbf{s}_i - \bar{\mathbf{x}}_i) + \sum_i \beta_i \tilde{\gamma}_i \underbrace{\nabla f(\bar{\mathbf{x}}_i)^T (\mathbf{s}_i - \bar{\mathbf{x}}_i)}_{-\text{gap}(\bar{\mathbf{x}}_i)} + \frac{L\gamma^2 D_2^2}{2} \\ &\leq \sum_i \beta_i \underbrace{\tilde{\gamma}_i}_{\leq \gamma} \underbrace{\|\nabla f(\mathbf{x}) - \nabla f(\bar{\mathbf{x}}_i)\|_2}_{L\|\mathbf{x} - \bar{\mathbf{x}}_i\|_2 = L\gamma D_3} \underbrace{\|\mathbf{s}_i - \bar{\mathbf{x}}_i\|_2}_{\leq D_2} - \sum_i \beta_i \tilde{\gamma}_i \text{gap}(\bar{\mathbf{x}}_i) + \frac{L\gamma^2 D_2^2}{2} \\ &\leq -\sum_i \beta_i \tilde{\gamma}_i \text{gap}(\bar{\mathbf{x}}_i) + \frac{L\gamma^2 D_2^2}{2} + \frac{2L\gamma^2 D_2 D_3}{2} \\ &\leq -\gamma^+ \sum_i \beta_i h(\bar{\mathbf{x}}_i) + \frac{L\gamma^2 D_2 (D_2 + 2D_3)}{2} \end{aligned}$$

where $\gamma = \gamma_k$, and $\gamma^+ = \gamma_{k+1}$. Now assume f is also L_2 -continuous, e.g. $|f(\mathbf{x}_1) - f(\mathbf{x}_2)| \leq L_2 \|\mathbf{x}_1 - \mathbf{x}_2\|_2$. Then, taking $h(\mathbf{x}) = f(\mathbf{x}) - f(\mathbf{x}^*)$,

$$\begin{aligned} h(\mathbf{x}^+) - h(\mathbf{x}) &\leq -\gamma^+ \sum_i \beta_i (h(\bar{\mathbf{x}}_i) - h(\mathbf{x})) - \gamma^+ \underbrace{\sum_i \beta_i h(\mathbf{x})}_{=1} + \frac{L\gamma^2 D_2 (D_2 + 2D_3)}{2} \\ &\leq \gamma \sum_i \beta_i L_2 \underbrace{\|\bar{\mathbf{x}}_i - \mathbf{x}\|_2}_{\leq \gamma D_3} - \gamma^+ h(\mathbf{x}) + \frac{L\gamma^2 D_2 (D_2 + 2D_3)}{2} \\ &\leq -\gamma^+ h(\mathbf{x}) + \frac{\gamma^2 (LD_2^2 + 2LD_2 D_3 + 2D_3)}{2} \\ &\leq -\gamma^+ h(\mathbf{x}) + D_4 (\gamma^+)^2 \end{aligned}$$

where $D_4 = \frac{LD_2^2 + 2LD_2 D_3 + 2D_3}{2}$ and we use $2 \geq (\gamma/\gamma^+)^2$ for all $k \geq 1$.

□

Proof. After establishing Lemma D.2, the rest of the proof is a recursive argument, almost identical to that in Jaggi (2013).

At $k = 0$, we define $h_0 = \max\{h(\mathbf{x}^{(0)}), \frac{D_4 c^2}{c-1}\}$, and it is clear that $h(\mathbf{x}^{(0)}) \leq h_0$.

Now suppose that for some k , $h(\mathbf{x}^{(k)}) \leq \frac{h_0}{k+1}$. Then

$$\begin{aligned}
 h(x_{k+1}) &\leq h(\mathbf{x}_k) - \gamma_{k+1} h(\mathbf{x}^{(k)}) + D_4 \gamma_{k+1}^2 \\
 &\leq \frac{h_0}{k+1} \cdot \frac{k+1}{c+k+1} + D_4 \frac{c^2}{(c+k+1)^2} \\
 &= \frac{h_0}{c+k+1} + D_4 \frac{c^2}{(c+k+1)^2} \\
 &= \left(h_0 + \frac{D_4 c^2}{c+k+1} \right) \left(\frac{k+2}{c+k+1} \right) \frac{1}{k+2} \\
 &\leq h_0 \left(1 + \frac{c-1}{c+k+1} \right) \left(\frac{k+2}{c+k+1} \right) \frac{1}{k+2} \\
 &\leq h_0 \underbrace{\left(\frac{2c+k}{c+k+1} \right) \left(\frac{k+2}{c+k+1} \right)}_{\leq 1} \frac{1}{k+2}.
 \end{aligned}$$

□

D.3 Negative Runge-Kutta convergence result

This section gives the proof for Proposition 3.3.

Lemma D.3 ($O(1/k)$ rate). *Start with $\mathbf{x}_0 = 1$. Then consider the sequence defined by*

$$\mathbf{x}_{k+1} = \left| \mathbf{x}_k - \frac{c_k}{k} \right|$$

where, no matter how large k is, there exist some constant where $C_1 < \max_{k' > k} c_{k'}$. (That is, although c_k can be anything, the smallest upper bound of c_k does not decay.) Then

$$\sup_{k' \geq k} |\mathbf{x}_{k'}| = \Omega(1/k).$$

That is, the smallest upper bound of $|\mathbf{x}_k|$ at least of order $1/k$.

Proof. We will show that the smallest upper bound of $|\mathbf{x}_k|$ is larger than $C_1/(2k)$.

Proof by contradiction. Suppose that at some point K , for all $k \geq K$, $|\mathbf{x}_k| < C_1/(2k)$. Then from that point forward,

$$\mathbf{sign}(\mathbf{x}_k - \frac{c_k}{k}) = -\mathbf{sign}(\mathbf{x}_k)$$

and there exists some $k' > k$ where $c_{k'} > C_1$. Therefore, at that point,

$$|\mathbf{x}_{k'+1}| = \frac{c_{k'}}{k'} - |\mathbf{x}_{k'}| \geq \frac{C_1}{2k'} > \frac{C_1}{2(k'+1)}.$$

This immediately establishes a contradiction. □

Now define the operator

$$T(\mathbf{x}_k) = \mathbf{x}_{k+1} - \mathbf{x}_k$$

and note that

$$|\mathbf{x}_{k+1}| = |\mathbf{x}_k + T(\mathbf{x}_k)| = \left| |\mathbf{x}_k| + \mathbf{sign}(\mathbf{x}_k) T(\mathbf{x}_k) \right|.$$

Thus, if we can show that there exist some ϵ , agnostic to k (but possibly related to Runge Kutta design parameters), and

$$\exists k' \geq k, \quad -\mathbf{sign}(\mathbf{x}_{k'}) T(\mathbf{x}_{k'}) > \frac{\epsilon}{k'}, \quad \forall k, \quad (8)$$

then based on the previous lemma, this shows $\sup_{k' > k} |\mathbf{x}_{k'}| = \Omega(1/k)$ as the smallest possible upper bound.

Lemma D.4. Assuming that $0 < q\mathbb{P}^{(k)}\beta < 1$ then there exists a finite point \bar{k} where for all $k > \bar{k}$,

$$|\mathbf{x}_k| \leq \frac{C_2}{k}$$

for some $C_2 \geq 0$.

Proof. We again use the block matrix notation

$$\mathbf{Z}^{(k)} = (\bar{\mathbf{S}} - \mathbf{x}_k \mathbb{K}^T) \Gamma^{(k)} (I + A^T \Gamma^{(k)})^{-1}$$

where $\Gamma^{(k)} = \mathbf{diag}(\tilde{\gamma}_i^{(k)})$ and each element $\tilde{\gamma}_i^{(k)} \leq \gamma^{(k)}$.

First, note that by construction, since

$$\|\bar{\mathbf{S}} - \mathbf{x}_k \mathbb{K}^T\|_{2,\infty} \leq D_4, \quad \|(I + A^T \Gamma^{(k)})^{-1}\|_2 \leq \|(I + A^T \Gamma^{(0)})^{-1}\|_2$$

are bounded above by constants, then

$$\|\mathbf{Z}^{(k)}\|_\infty \leq \frac{c}{c+k} C_1$$

for $C_1 = D_4 \|(I + A^T \Gamma^{(0)})^{-1}\|_2$.

First find constants C_3, C_4 , and \bar{k} such that

$$\frac{C_3}{k} \leq \mathbb{K}^T \mathbf{P}^{(k)} \beta \leq \frac{C_4}{k}, \quad \forall k > \bar{k}, \quad (9)$$

and such constants always exist, since by assumption, there exists some $a_{\min} > 0$, $a_{\max} < 1$ and some k' where

$$a_{\min} < q\mathbb{P}^{(k')}\beta < a_{\max} \Rightarrow \frac{a_{\min}}{q\gamma_{\max}} \leq (I + A^T \Gamma^{(k')})^{-1} \beta \leq \frac{a_{\max}}{q\gamma_{\min}}$$

where

$$\gamma_{\min} = \min_i \frac{c}{c+k'+\omega_i^{(k')}}, \quad \gamma_{\max} = \frac{c}{c+k'}.$$

Additionally, for all $k > c+1$,

$$\frac{c}{2k} \leq \frac{c}{c+k+1} \leq \Gamma_{ii}^{(k)} \leq \frac{c}{c+k} \leq \frac{c}{k}.$$

Therefore taking

$$C_3 = \frac{ca_{\min}}{2q\gamma_{\max}}, \quad C_4 = \frac{ca_{\max}}{q\gamma_{\min}}, \quad \bar{k} = \max\{k', c+1\}$$

satisfies (9).

Now define

$$C_2 = \max\{|\mathbf{x}_1|, 4cqC_1\|A\|_\infty, 4C_3, 4C_4\}.$$

We will now inductively show that $|\mathbf{x}_k| \leq \frac{C_2}{k}$. From the definition of C_2 , we have the base case for $k=1$:

$$|\mathbf{x}_1| \leq \frac{|\mathbf{x}_1|}{1} \leq \frac{C_2}{k}.$$

Now assume that $|\mathbf{x}_k| \leq \frac{C_2}{k}$. Recall that

$$\mathbf{x}_{k+1} = \mathbf{x}_k (1 - \mathbb{K}^T \mathbf{P}^{(k)} \beta) + \bar{\mathbf{S}} \mathbf{P}^{(k)} \beta, \quad \bar{\mathbf{S}} = [\bar{s}_1, \dots, \bar{s}_q], \quad \mathbf{s}_i = -\mathbf{sign}(\bar{\mathbf{x}}_i)$$

and we denote the composite mixing term $\bar{\gamma}^{(k)} = \mathbb{K}^T \mathbb{P}^{(k)} \beta$. We now look at two cases separately.

- Suppose first that $\bar{\mathbf{S}} = -\mathbf{sign}(\mathbf{x}_k \mathbb{K}^T)$, e.g. $\mathbf{sign}(\bar{\mathbf{x}}_i) = \mathbf{sign}(\mathbf{x}_k)$ for all i . Then

$$\bar{\mathbf{S}} \mathbb{P}^{(k)} \beta = -\mathbf{sign}(\mathbf{x}_k) \bar{\gamma}_k,$$

and

$$\begin{aligned}
 |\mathbf{x}_{k+1}| &= |\mathbf{x}_k(1 - \bar{\gamma}^{(k)}) + \bar{\mathbf{S}}\mathbf{P}^{(k)}\beta| \\
 &= |\mathbf{x}_k(1 - \bar{\gamma}^{(k)}) - \mathbf{sign}(\mathbf{x}_k)\bar{\gamma}^{(k)}| \\
 &= \left| \underbrace{\mathbf{sign}(\mathbf{x}_k)\mathbf{x}_k}_{|\mathbf{x}_k|} (1 - \bar{\gamma}^{(k)}) - \underbrace{\mathbf{sign}(\mathbf{x}_k)\mathbf{sign}(\mathbf{x}_k)}_{=1} \bar{\gamma}^{(k)} \right| \\
 &= ||\mathbf{x}_k|(1 - \bar{\gamma}^{(k)}) - \bar{\gamma}^{(k)}| \\
 &\leq \max\{|\mathbf{x}_k|(1 - \bar{\gamma}^{(k)}) - \bar{\gamma}^{(k)}, \bar{\gamma}^{(k)} - |\mathbf{x}_k|(1 - \bar{\gamma}^{(k)})\} \\
 &\leq \max \left\{ \underbrace{\frac{C_2}{k} \left(1 - \frac{C_3}{k}\right) - \frac{C_3}{k}}_{(*)}, \frac{C_4}{k} \right\}
 \end{aligned}$$

and when $k \geq \frac{C_2}{C_3} \iff C_3 \geq \frac{C_2}{k}$,

$$(*) \leq C_2 \left(\frac{1}{k} - \frac{1}{k^2} \right) \leq \frac{C_2}{k+1}.$$

Taking also $C_4 \leq \frac{C_2}{4}$,

$$|\mathbf{x}_{k+1}| \leq \max \left\{ \frac{C_2}{k+1}, \frac{C_2}{4k} \right\} \leq \frac{C_2}{k+1}$$

for all $k \geq 1$.

- Now suppose that there is some i where $\bar{s}_i = \mathbf{sign}(\mathbf{x}_k \mathbf{A}^T)$. Now since

$$\bar{\mathbf{S}} = -\mathbf{sign}(\mathbf{x}_k \mathbf{A}^T + \mathbf{Z}\mathbf{A}^T)$$

then this implies that $|\mathbf{x}_k| < (\mathbf{Z}\mathbf{A}^T)_i$. But since

$$|(\mathbf{Z}\mathbf{A}^T)_i| \leq \|\mathbf{Z}\|_\infty \|\mathbf{A}\|_\infty q \leq \frac{c}{c+k} (C_1 \|\mathbf{A}\|_\infty q) \leq \frac{C_2}{4k},$$

this implies that

$$|\mathbf{x}_{k+1}| \leq \frac{C_2}{4k} \left(1 - \frac{C_3}{k}\right) + \frac{C_2}{4k} \leq \frac{C_2}{2k} \leq \frac{C_2}{k+1}, \quad \forall k > 1.$$

Thus we have shown the induction step, which completes the proof. \square

Lemma D.5. *There exists a finite point \tilde{k} where for all $k > \tilde{k}$,*

$$\frac{c}{c+k} - \frac{C_4}{k^2} < |\xi_i| < \frac{c}{c+k} + \frac{C_4}{k^2}$$

for some constant $C_4 > 0$.

Proof. Our goal is to show that

$$\gamma^{(k)} - \frac{C_4}{k^2} \leq \|\mathbf{Z}\|_\infty \leq \gamma^{(k)} + \frac{C_4}{k^2}$$

for some $C_4 \geq 0$, and for all $k \geq k'$ for some $k' \geq 0$. Using the Woodbury matrix identity,

$$\Gamma(I + \mathbf{A}^T\Gamma)^{-1} = \Gamma(I - \mathbf{A}^T(\Gamma + \mathbf{A}^T)^{-1}\Gamma)$$

and thus

$$\mathbf{Z}^{(k)} = \bar{\mathbf{S}}\Gamma - \underbrace{(\mathbf{x}_k \mathbf{A}^T \Gamma + (\bar{\mathbf{S}} - \mathbf{x}_k \mathbf{A}^T) \Gamma \mathbf{A}^T (\Gamma + \Gamma \mathbf{A}^T)^{-1} \Gamma)}_{\mathbf{B}}.$$

and thus

$$|\bar{s}_i \tilde{\gamma}_i| - \frac{C_3}{k^2} \leq |\xi_i^{(k)}| \leq |\bar{s}_i \tilde{\gamma}_i| + \frac{C_3}{k^2}$$

where via triangle inequalities and norm decompositions,

$$\frac{C_3}{k^2} = \max_i |\mathbf{B}_i| \leq \underbrace{\|\mathbf{x}_k\|}_{O(1/k)} \gamma_k + D_4 \gamma_k^2 \|A\|_\infty (I + \Gamma^{(0)} A^T)^{-1} = O(1/k^2).$$

Finally, since $\bar{s}_i \in \{-1, 1\}$, then $|\bar{s}_i \tilde{\gamma}_i| = \tilde{\gamma}_i$, and in particular,

$$\frac{c}{c+k+\omega_i} \leq \frac{c}{c+k}$$

and

$$\frac{c}{c+k+\omega_i} \geq \frac{c}{c+k+\omega_{\max}} = \frac{c}{c+k} - \frac{c}{c+k} \frac{\omega_{\max}}{c+k+\omega_{\max}} \geq \frac{c}{c+k} - \frac{c\omega_{\max}}{k^2}$$

Therefore, taking $C_4 = c\omega_{\max} + C_3$ completes the proof. \square

Lemma D.6. *There exists some large enough \tilde{k} where for all $k \geq \tilde{k}$, it must be that*

$$\exists k' \geq k, \quad -\mathbf{sign}(\mathbf{x}_{k'}) T(\mathbf{x}_{k'}) > \frac{\epsilon}{k'}. \quad (10)$$

Proof. Define a partitioning $S_1 \cup S_2 = \{1, \dots, q\}$, where

$$S_1 = \{i : \xi_i > 0\}, \quad S_2 = \{j : \xi_j \leq 0\}.$$

Defining $\bar{\xi} = \frac{c}{c+k}$,

$$\left| \sum_{i=1}^q \beta_i \xi_i \right| = \left| \sum_{i \in S_1} \beta_i |\xi_i| - \sum_{j \in S_2} \beta_j |\xi_j| \right| \geq \left(\bar{\xi} - \frac{C_4}{k^2} \right) \cdot \left| \sum_{i \in S_1} \beta_i - \sum_{j \in S_2} \beta_j \right|.$$

By assumption, there does not exist a combination of β_i where a specific linear combination could cancel them out; that is, suppose that there exists some constant $\bar{\beta}$, where for *every* partition of sets S_1, S_2 ,

$$0 < \bar{\beta} := \min_{S_1, S_2} \left| \sum_{i \in S_1} \beta_i - \sum_{j \in S_2} \beta_j \right|.$$

Then

$$\left| \sum_{i=1}^q \beta_i \xi_i \right| \geq \left(\frac{c}{c+k} - \frac{C_2}{k^2} \right) \bar{\beta} \geq \bar{\beta} \frac{\max\{C_2, c\}}{k}.$$

Picking $\epsilon = \max\{C_2, c\}$ concludes the proof. \square

E AVERAGED FRANK WOLFE PROOFS

E.1 Accumulation terms

Lemma E.1. *For an averaging term $\bar{s}(t)$ satisfying*

$$\dot{\bar{s}}(t) = \beta(t)(s(t) - \bar{s}(t)), \quad \bar{s}(0) = s(0) = 0$$

where $\beta(t) = \frac{c^p}{(c+t)^p}$, then

$$\bar{s}(t) = \begin{cases} e^{-\alpha(t)} \int_0^t \frac{c^p e^{\alpha(\tau)}}{(c+\tau)^p} s(\tau) d\tau, & p \neq 1 \\ \frac{c}{(c+t)^c} \int_0^t (c+\tau)^{c-1} s(\tau) d\tau & p = 1 \end{cases}$$

where $\alpha(t) = \frac{c^p(c+t)^{1-p}}{1-p}$. If $s(t) = 1$ for all t , then we have an accumulation term

$$\bar{s}(t) = \begin{cases} 1 - \frac{e^{\alpha(0)}}{e^{\alpha(t)}}, & p \neq 1 \\ 1 - \left(\frac{c}{c+t}\right)^c, & p = 1 \end{cases}$$

Proof. This can be done through simple verification.

- If $p \neq 1$,

$$\alpha'(t) = \frac{c^p}{(c+t)^p} = \beta(t),$$

and via chain rule,

$$\bar{s}'(t) = \underbrace{e^{-\alpha(t)} \frac{c^p \exp(\alpha(t))}{(c+t)^p}}_{=\beta(t)} s(t) - \underbrace{\alpha'(t) \exp(-\alpha(t)) \int_0^t \frac{c^p \exp(\alpha(\tau))}{(c+\tau)^p} s(\tau) d\tau}_{\bar{s}(t)}.$$

The accumulation term can be verified if

$$e^{-\alpha(t)} \int_0^t \frac{c^p e^{\alpha(\tau)}}{(c+\tau)^p} d\tau = 1 - \frac{e^{\alpha(0)}}{e^{\alpha(t)}}$$

which is true since

$$e^{-\alpha(t)} \int_0^t \frac{c^p e^{\alpha(\tau)}}{(c+\tau)^p} d\tau = e^{-\alpha(t)} \int_0^t \left(\frac{d}{d\tau} e^{\alpha(\tau)}\right) d\tau.$$

- If $p = 1$

$$\bar{s}'(t) = \frac{c}{(c+t)} s(t) - \frac{c^2}{(c+t)^{c+1}} \int_0^t (c+\tau)^{c-1} s(\tau) d\tau = \frac{c}{(c+t)} (s(t) - \bar{s}(t)).$$

For the accumulation term,

$$\frac{c}{(c+t)^c} \int_0^t (c+\tau)^{c-1} d\tau = \frac{c}{(c+t)^c} \int_0^t \left(\frac{\partial}{\partial \tau} \frac{(c+\tau)^c}{c}\right) d\tau = 1 - \left(\frac{c}{c+t}\right)^c.$$

□

For convenience, we define

$$\beta_{t,\tau} := \begin{cases} \frac{c^p e^{\alpha(\tau) - \alpha(t)}}{(c+\tau)^p}, & p \neq 1 \\ \frac{c(c+\tau)^{c-1}}{(c+t)^b}, & p = 1, \end{cases} \quad \bar{\beta}_t := \begin{cases} 1 - \frac{\exp(\alpha(0))}{\exp(\alpha(t))}, & p \neq 1 \\ 1 - \left(\frac{c}{c+t}\right)^b, & p = 1 \end{cases}$$

Lemma E.2. For the averaging sequence \bar{s}_k defined recursively as

$$\bar{s}_{k+1} = \bar{s}_k + \beta_k(\mathbf{s}_k - \bar{s}_k), \quad \bar{s}_0 = 0.$$

Then

$$\bar{s}_k = \sum_{i=1}^k \beta_{k,i} \mathbf{s}_i, \quad \beta_{k,i} = \frac{c^p}{(c+i)^p} \prod_{j=0}^{k-i-1} \left(1 - \frac{c^p}{(c+k-j)^p}\right) \stackrel{p=1}{=} \frac{c}{c+i} \prod_{j=0}^c \frac{i+j+1}{c+k-j}$$

and moreover, $\sum_{i=1}^k \beta_{k,i} = 1$.

Proof.

$$\begin{aligned}
 \bar{s}_{k+1} &= \frac{c^p}{(c+k)^p} s_k + \left(1 - \frac{c^p}{(c+k)^p}\right) \bar{s}_k \\
 &= \frac{c^p}{(c+k)^p} s_k + \frac{c^p}{(c+k-1)^p} \left(1 - \frac{c^p}{(c+k)^p}\right) s_{k-1} + \left(1 - \frac{c^p}{(c+k)^p}\right) \left(1 - \frac{c^p}{(c+k-1)^p}\right) \bar{s}_{k-1} \\
 &= \sum_{i=0}^k \frac{c^p}{(c+k-i)^p} \prod_{j=0}^{i-1} \left(1 - \frac{c^p}{(c+k-j)^p}\right) s_{k-i} \\
 &\stackrel{l=k-i}{=} \sum_{l=1}^k \underbrace{\frac{c^p}{(c+l)^p} \prod_{j=0}^{k-l-1} \left(1 - \frac{c^p}{(c+k-j)^p}\right)}_{\beta_{k,l}} s_l.
 \end{aligned}$$

If $p = 1$, then

$$\beta_{k,i} = \frac{c}{c+i} \prod_{l=0}^{k-i-1} \frac{k-l}{c+k-l} = \frac{c}{c+i} \frac{k(k-1)(k-2)\cdots(i+1)}{(c+k)(c+k-1)\cdots(c+i+1)} = \frac{c}{c+i} \prod_{j=0}^i \frac{i+j+1}{c+k-j}$$

For all p , to show the sum is 1, we do so recursively. At $k = 1$, $\beta_{1,1} = \frac{c^p}{(c+1)^p}$. Now, if $\sum_{i=0}^{k-1} \beta_{k-1,i} = 1$, then for $i \leq k-1$

$$\beta_{k,i} = \left(1 - \frac{c^p}{(c+k)^p}\right) \beta_{k-1,i}, \quad i \leq k-1$$

and for $i = k$, $\beta_{k,k} = \frac{c^p}{(c+k)^p}$. Then

$$\sum_{i=1}^k \beta_{k,i} = \beta_{k,k} + \left(1 - \frac{c^p}{(c+k)^p}\right) \sum_{l=1}^{k-1} \beta_{k-1,i} = \frac{c^p}{(c+k)^p} + \left(1 - \frac{c^p}{(c+k)^p}\right) = 1.$$

□

E.2 Averaging

In the vanilla Frank-Wolfe method, we have two players s and x , and as $x \rightarrow x^*$, s may oscillate around the solution facet however it would like, so that its average is x^* but $\|s - x^*\|_2$ remains bounded away from 0. However, we now show that if we replace s with \bar{s} , whose velocity slows down, then it must be that $\|s - x^*\|_2$ decays.

Lemma E.3 (Continuous averaging). *Consider some vector trajectory $v(t) \in \mathbb{R}^n$, and suppose*

- $\|v(t)\|_2 \leq D$ for arbitrarily large t
- $\|v'(t)\|_2 = \beta(t)D$
- $\frac{1}{2} \left\| \int_t^\infty \gamma(\tau) v(\tau) d\tau \right\|_2^2 = O(1/t^q)$ for $q > 0$.

Then $\|v(t)\|_2^2 \leq O(t^{q/2+p-1})$.

Proof. We start with the orbiting property.

$$\frac{d}{dt} \left(\frac{1}{2} \left\| \int_t^\infty \gamma(\tau) v(\tau) d\tau \right\|_2^2 \right) = - \int_t^\infty \gamma(t) \gamma(\tau) v(\tau)^T v(t) d\tau \leq 0.$$

Since this is happening asymptotically, then the negative derivative of the LHS must be upper bounded by the negative derivative of the RHS. That is, if a function is decreasing asymptotically at a certain rate, then its negative derivative should be decaying asymptotically at the negative derivative of this rate. So,

$$\int_t^\infty \gamma(\tau) v(\tau)^T v(t) d\tau \leq O(1/t^q).$$

This indicates either that $\|v(t)\|_2$ is getting smaller (converging) or $v(t)$ and its average are becoming more and more uncorrelated (orbiting).

Doing the same trick again with the negative derivative,

$$-\frac{d}{dt} \int_t^\infty \gamma(\tau) v(\tau)^T v(t) d\tau = \gamma(t) \|v(t)\|_2^2 - \int_t^\infty \gamma(\tau) v(\tau)^T v'(t) d\tau$$

By similar logic, this guy should also be decaying at a rate $O(1/t^{q+1})$, so

$$\gamma(t) \|v(t)\|_2^2 \leq \frac{C_2}{t^{q+1}} + \underbrace{\int_t^\infty \gamma(\tau) v(\tau)^T v'(t) d\tau}_{\leq O(1/t^{q/2})} \leq \frac{C_2}{t^{q+1}} + \frac{C_3}{t^{q/2+p}}$$

Therefore

$$\|v(t)\|_2^2 \leq \frac{C_2}{t^q} + \frac{C_3}{t^{q/2+p-1}} = O\left(\frac{1}{t^{q/2+p-1}}\right).$$

□

Corollary E.4. *Suppose f is μ -strongly convex. Then*

$$\|\bar{s}(t) - x(t)\|_2^2 \leq Ct^{-(q/2+p-1)}$$

for some constant $C > 0$.

Proof. Taking $v(t) = \bar{s}(t) - x(t)$, it is clear that if $\beta(t) \geq \gamma(t)$ then the first two conditions are satisfied. In the third condition, note that

$$\int_t^\infty \gamma(\tau) (\bar{s}(\tau) - x(\tau)) d\tau = \int_t^\infty \dot{x}(\tau) d\tau = x^* - x(t)$$

and therefore

$$\frac{1}{2} \left\| \int_t^\infty \gamma(\tau) v(\tau) d\tau \right\|_2^2 = \frac{1}{2} \|x^* - x(t)\|_2^2 \leq \mu(f(x) - f^*)$$

by strong convexity. □

Lemma E.5 (Discrete averaging). *Consider some vector trajectory $\mathbf{v}_k \in \mathbb{R}^n$. Then the following properties cannot all be true.*

- $\|\mathbf{v}_k\|_2 \leq D$ for arbitrarily large k
- $\|\mathbf{v}_{k+1} - \mathbf{v}_k\|_2 \leq \beta_k D$
- $\frac{1}{2} \left\| \sum_{i=k}^\infty \gamma_i \mathbf{v}_i \right\|_2^2 \leq \frac{C_1}{k^q}$ for $q > 0$.

Then $\|\mathbf{v}_k\|_2^2 \leq O(k^{q/2+p-1})$.

Proof. The idea is to recreate the same proof steps as in the previous lemma. Note that the claim is not that these inequalities happen at each step, but that they must hold asymptotically in order for the asymptotic decay rates to hold. So

$$\frac{1}{2} \left\| \sum_{i=k}^\infty \gamma_i \mathbf{v}_i \right\|_2^2 - \frac{1}{2} \left\| \sum_{i=k}^\infty \gamma_{i+1} \mathbf{v}_{i+1} \right\|_2^2 = \frac{\gamma_k^2}{2} \|\mathbf{v}_k\|_2^2 + \gamma_k \mathbf{v}_k^T \left(\sum_{i=k}^\infty \gamma_{i+1} \mathbf{v}_{i+1} \right) \leq \frac{C_1}{(k+1)^q} - \frac{C_1}{k^q} = \frac{C_2}{k^q}$$

and therefore

$$\frac{\gamma_k}{2} \|\mathbf{v}_k\|_2^2 + \mathbf{v}_k^T \left(\sum_{i=k}^\infty \gamma_{i+1} \mathbf{v}_{i+1} \right) \leq \frac{C_1}{(k+1)^{q+1}} - \frac{C_1}{k^{q+1}} = \frac{C_2}{k^{q+1}}.$$

Next,

$$\begin{aligned}
 & \mathbf{v}_k^T \left(\sum_{i=k}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right) - \mathbf{v}_{k+1}^T \left(\sum_{i=k+1}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right) + \mathbf{v}_k^T \left(\sum_{i=k+1}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right) - \mathbf{v}_k^T \left(\sum_{i=k+1}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right) \\
 = & \gamma_{k+1} \mathbf{v}_k^T \mathbf{v}_{k+1} + (\mathbf{v}_k - \mathbf{v}_{k+1})^T \left(\sum_{i=k+1}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right) \\
 & \frac{\gamma_k}{2} \|\mathbf{v}_k\|_2^2 - \frac{\gamma_{k+1}}{2} \|\mathbf{v}_{k+1}\|_2^2 + \mathbf{v}_k^T \left(\sum_{i=k}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right) - \mathbf{v}_{k+1}^T \left(\sum_{i=k+1}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right) = \\
 & \frac{\gamma_k}{2} \|\mathbf{v}_k\|_2^2 - \frac{\gamma_{k+1}}{2} \|\mathbf{v}_{k+1}\|_2^2 + \underbrace{\gamma_{k+1} \mathbf{v}_k^T \mathbf{v}_{k+1} + (\mathbf{v}_k - \mathbf{v}_{k+1})^T \left(\sum_{i=k+1}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right)}_{-\frac{\gamma_{k+1}}{2} \|\mathbf{v}_{k+1} - \mathbf{v}_k\|_2^2 + \frac{\gamma_{k+1}}{2} \|\mathbf{v}_k\|_2^2} \leq \frac{C_3}{k^{q+1}}
 \end{aligned}$$

Therefore

$$\frac{\gamma_k + \gamma_{k+1}}{2} \|\mathbf{v}_k\|_2^2 \leq \frac{C_3}{k^{q+1}} + \underbrace{(\mathbf{v}_{k+1} - \mathbf{v}_k)^T \left(\sum_{i=k+1}^{\infty} \gamma_{i+1} \mathbf{v}_{i+1} \right)}_{O(\beta_k/k^{q/2})} + \underbrace{\frac{\gamma_{k+1}}{2} \|\mathbf{v}_{k+1} - \mathbf{v}_k\|_2^2}_{O(\gamma_k \beta_k^2)}$$

Finally,

$$\|\mathbf{v}_k\|_2^2 \leq \frac{C_3}{k^q} + \frac{C_4}{k^{q/2+p-1}} + \frac{C_5}{k^{2p}} = O(1/k^{\min\{q/2+p-1, 2p\}}).$$

□

Corollary E.6. Suppose f is μ -strongly convex. Then if $f(x) - f^* = O(k^{-q})$

$$\|\bar{\mathbf{s}}_k - \mathbf{x}_k\|_2^2 \leq C \max\{k^{-(q/2+p-1)}, k^{-2p}\}$$

for some constant $C > 0$.

Proof. Taking $\mathbf{v}_k = \bar{\mathbf{s}}_k - \mathbf{x}_k$, it is clear that if $\beta(t) \geq \gamma(t)$ then the first two conditions are satisfied. In the third condition, note that

$$\sum_{i=k}^{\infty} \gamma_i (\bar{\mathbf{s}}_i - \mathbf{x}_i) = \sum_{i=k}^{\infty} \mathbf{x}_{i+1} - \mathbf{x}_i = \mathbf{x}^* - \mathbf{x}_k$$

and therefore

$$\frac{1}{2} \left\| \sum_{i=k}^{\infty} \gamma_i (\bar{\mathbf{s}}_i - \mathbf{x}_i) \right\|_2^2 = \frac{1}{2} \|\mathbf{x}^* - \mathbf{x}_k\|_2^2 \leq \mu (f(\mathbf{x}_k) - f^*)$$

by strong convexity. □

E.3 Global rates

Lemma E.7 (Continuous energy function decay). Suppose $c \geq q - 1$, and

$$g(t) \leq - \int_0^t \frac{\exp(\alpha(\tau)) (c + \tau)^c}{\exp(\alpha(t)) (c + t)^c} \frac{C_1}{(b + \tau)^r} d\tau$$

Then

$$g(t) \leq - \frac{\left(1 - \frac{\alpha(1)}{\exp(\alpha(1))}\right) C_1}{\alpha(t)(1-p)} (c+t)^{1-r}$$

Proof.

$$\begin{aligned}
 \frac{g(t)}{C_1} \exp(\alpha(t))(c+t)^c &\leq - \int_0^t \exp(\alpha(\tau))(c+\tau)^{c-r} d\tau \\
 &= - \int_0^t \sum_{k=1}^{\infty} \frac{\alpha(\tau)^k}{k!} (c+\tau)^{c-r} d\tau \\
 &= - \int_0^t \sum_{k=1}^{\infty} \frac{c^{kp}}{k!(1-p)^k} (c+\tau)^{k-pk+c-r} d\tau \\
 &\stackrel{\text{Fubini}}{=} - \sum_{k=1}^{\infty} \frac{c^{kp}}{(1-p)^k k!} \int_0^t (c+\tau)^{k-pk+c-r} d\tau \\
 &= - \sum_{k=1}^{\infty} \frac{c^{kp}}{(1-p)^k k!} \frac{(c+t)^{k-pk+c-r+1} - c^{k-pk+c-r+1}}{k-pk+c-r+1} \\
 &= - \sum_{k=1}^{\infty} \frac{1}{(k+1)!} \left(\frac{c^p}{(c+t)^{p-1}(1-p)} \right)^k (c+t)^{1+c-r} \underbrace{\frac{1}{(1-p) + (c-r+1)/k} \frac{k+1}{k}}_{\geq C_2} \\
 &\leq -C_2 \sum_{k=1}^{\infty} \frac{(c+t)^{1+c-r} \alpha(t)^{k+1}}{(k+1)! \alpha(t)} \\
 &= - \frac{C_2 (c+t)^{1+c-r}}{\alpha(t)} (\exp(\alpha(t)) - \alpha(1))
 \end{aligned}$$

Then

$$\begin{aligned}
 g(t) &\leq - \frac{C_1 C_2}{\alpha(t)} (c+t)^{1-r} \left(1 - \frac{\alpha(1)}{\exp(\alpha(t))} \right) \\
 &\leq - \frac{C_1 C_2}{\alpha(t)} (c+t)^{1-r} \left(1 - \frac{\alpha(1)}{\exp(\alpha(1))} \right) \\
 &\leq - \frac{C_1 C_3}{\alpha(t)} (c+t)^{1-r}
 \end{aligned}$$

where $C_3 = C_2 \left(1 - \frac{\alpha(1)}{\exp(\alpha(1))} \right)$ and $C_2 = \frac{1}{1-p}$ satisfies the condition. \square

Theorem E.8 (Continuous global rate). *Suppose $0 < p < 1$. Then the averaged FW flow decays as $O(1/t^{1-p})$.*

Proof. Consider the error function

$$g(t) := \nabla f(x(t))^T (\bar{s}(t) - x(t)), \quad g(0) = 0.$$

$$\begin{aligned}
 \dot{g}(t) &= \frac{\partial}{\partial t} \nabla f(x)^T (\bar{s} - x) \\
 &= \left(\frac{\partial}{\partial t} \nabla f(x)^T \right) (\bar{s} - x) + \nabla f(x)^T \left(\frac{\partial}{\partial t} (\bar{s} - x) \right) \\
 &= \underbrace{\left(\frac{\partial}{\partial t} \nabla f(x)^T \right)}_{=\dot{x}^T \nabla^2 f(x)} (\bar{s} - x) + \nabla f(x)^T (\beta(t)(s(t) - \bar{s}(t)) - \gamma(t)(\bar{s}(t) - x(t))) \\
 &\leq \underbrace{\gamma (\bar{s} - x)^T \nabla^2 f(x) (\bar{s} - x)}_{\leq 4LD^2 \gamma(t)} + \underbrace{\beta(t) \nabla f(x)^T (s(t) - x(t))}_{-\beta(t) \mathbf{gap}(x)} - (\beta(t) + \gamma(t)) \underbrace{\nabla f(x)^T (\bar{s}(t) - x(t))}_{=g(t)}
 \end{aligned}$$

$$g(t) \leq \int_0^t \frac{\exp(\alpha(\tau))}{\exp(\alpha(t))} \frac{(c+\tau)^c}{(c+t)^c} \underbrace{\left(\frac{4LD^2c}{c+\tau} - \frac{b^p}{(b+\tau)^p} \mathbf{gap}(x(\tau)) \right)}_{A(\tau)} d\tau$$

$$\dot{h}(t) \leq \gamma(t)g(t) \leq \gamma(t) \int_0^t \underbrace{\frac{\exp(\alpha(\tau))}{\exp(\alpha(t))} \frac{(c+\tau)^c}{(c+t)^c}}_{\mu(\tau)} A(\tau) d\tau$$

In order for $h(t)$ to decrease, it must be that $\dot{h}(t) \leq 0$. However, since $\mu(\tau) \geq 0$ for all $\tau \geq 0$, it must be that $A(\tau) \leq 0$, e.g.

$$\frac{c^p}{(c+\tau)^p} \mathbf{gap}(x(\tau)) \geq \frac{4LD^2c}{c+\tau}.$$

which would imply $h(t) = O(1/(c+t)^{1-p})$. Let us therefore test the candidate solution

$$h(t) = \frac{C_3}{(c+t)^{1-p}}.$$

Additionally, from Lemma E.7, if

$$A(\tau) \leq -\frac{C_1}{(c+t)} \Rightarrow g(t) \leq -\frac{C_1}{\alpha(t)(1-p)}$$

and therefore

$$\dot{h}(t) \leq \gamma(t)g(t) \leq -\frac{c}{c+t} \frac{C_1}{c^p} \cdot (c+t)^{p-1}$$

$$\int_0^t (h(\tau)) d\tau \leq \frac{C_1 c}{(1-p)c^p} (c+t)^{p-1}$$

which satisfies our candidate solution for $C_3 = \frac{C_1 c}{(p-1)c^p}$. □

This term $(\bar{s} - x)^T \nabla^2 f(x) (\bar{s} - x) \leq 4LD^2 \gamma(t)$ is an important one to consider when talking about local vs global distance. The largest values of the Hessian will probably not correspond to the indices that are “active”, and thus this bound is very loose near optimality.

Lemma E.9 (Discrete energy decay). *Suppose $0 < p < 1$. Consider the error function*

$$\mathbf{g}_k := \nabla f(\mathbf{x}_k)^T (\bar{\mathbf{s}}_k - \mathbf{x}_k).$$

Then

$$\mathbf{g}_k \leq -\sum_{i=0}^{k-1} \beta_{i,i} \left(\frac{i+1+c}{k+c} \right)^c \mathbf{gap}(\mathbf{x}_i) - \beta_{k,k} \mathbf{gap}(\mathbf{x}_k) + \frac{4D^2 L C_1}{(k+c)^p} + \left(\frac{c}{k+c} \right)^c \mathbf{g}_0.$$

where $C_1 = c^p \left(1 + \frac{1}{(c+1)(1-p)-1} \right)$.

Importantly, C_1 is finite only if $p < 1$. When $p = 1$, the right hand side is at best bounded by a constant, and does not decay, which makes it impossible to show method convergence.

Proof. Define $\mathbf{z}_k = \nabla f(\mathbf{x}_k)$, $\mathbf{g}_k = \mathbf{z}_k^T (\bar{\mathbf{s}}_k - \mathbf{x}_k)$. Then

$$\begin{aligned}
 \mathbf{g}_k &= \underbrace{\beta_{k,k} \mathbf{z}_k^T (\mathbf{s}_k - \mathbf{x}_k)}_{-\beta_{k,k} \mathbf{gap}(\mathbf{x}_k)} + \underbrace{\sum_{i=0}^{k-1} \beta_{k,i} \mathbf{z}_k^T (\mathbf{s}_i - \mathbf{x}_k)}_A \\
 A &= \sum_{i=0}^{k-1} \underbrace{\frac{\beta_{k,i}}{\beta_{k-1,i}}}_{=(1-\frac{c}{c+i})^p} \beta_{k-1,i} \mathbf{z}_k^T (\mathbf{s}_i - \mathbf{x}_k) \leq (1 - \frac{c^p}{(c+k)^p}) \underbrace{\mathbf{z}_k^T (\bar{\mathbf{s}}_{k-1} - \mathbf{x}_k)}_{=B} \\
 B &= \mathbf{z}_k^T (\bar{\mathbf{s}}_{k-1} - \underbrace{(\mathbf{x}_{k-1} + \gamma_{k-1}(\bar{\mathbf{s}}_{k-1} - \mathbf{x}_{k-1}))}_{\mathbf{x}_k}) \\
 &= (1 - \gamma_{k-1}) \mathbf{z}_k^T (\bar{\mathbf{s}}_{k-1} - \mathbf{x}_{k-1}) \\
 &= (1 - \gamma_{k-1}) (\mathbf{z}_k - \mathbf{z}_{k-1})^T \underbrace{(\bar{\mathbf{s}}_{k-1} - \mathbf{x}_{k-1})}_{(\mathbf{x}_k - \mathbf{x}_{k-1}) \gamma_{k-1}^{-1}} + (1 - \gamma_{k-1}) \underbrace{\mathbf{z}_{k-1}^T (\bar{\mathbf{s}}_{k-1} - \mathbf{x}_{k-1})}_{\mathbf{g}_{k-1}} \\
 &= \frac{(1 - \gamma_{k-1})}{\gamma_{k-1}} \underbrace{(\mathbf{z}_k - \mathbf{z}_{k-1})^T (\mathbf{x}_k - \mathbf{x}_{k-1})}_{\leq L \|\mathbf{x}_k - \mathbf{x}_{k-1}\|_2^2} + (1 - \gamma_{k-1}) \mathbf{g}_{k-1} \\
 &\leq (1 - \gamma_{k-1}) \gamma_{k-1} L \underbrace{\|\bar{\mathbf{s}}_{k-1} - \mathbf{x}_{k-1}\|_2^2}_{4D^2} + (1 - \gamma_{k-1}) \mathbf{g}_{k-1}
 \end{aligned}$$

Overall,

$$\begin{aligned}
 \mathbf{g}_k &\leq -\beta_{k,k} \mathbf{gap}(\mathbf{x}_k) + 4D^2 L \gamma_{k-1} (1 - \gamma_{k-1}) (1 - \beta_k) + \underbrace{(1 - \beta_k)(1 - \gamma_{k-1})}_{\mu_k} \mathbf{g}_{k-1} \\
 &= -\beta_{k,k} \mathbf{gap}(\mathbf{x}_k) + 4D^2 L \gamma_{k-1} \mu_k + \mu_k \mathbf{g}_{k-1} \\
 &= -\beta_{k,k} \mathbf{gap}(\mathbf{x}_k) + 4D^2 L \gamma_{k-1} \mu_k - \beta_{k-1,k-1} \mu_k \mathbf{gap}(\mathbf{x}_{k-1}) + 4D^2 L \gamma_{k-2} \mu_{k-1} \mu_k + \mu_k \mu_{k-1} \mathbf{g}_{k-2} \\
 &= -\sum_{i=0}^{k-1} \beta_{i,i} \mathbf{gap}(\mathbf{x}_i) \prod_{j=i+1}^k \mu_j - \beta_{k,k} \mathbf{gap}(\mathbf{x}_k) + 4D^2 L \sum_{i=0}^k \gamma_{k-i} \prod_{j=i}^k \mu_j + \prod_{j=1}^k \mu_j \mathbf{g}_0
 \end{aligned}$$

Now we compute $\prod_{j=i}^k \mu_j$

$$\prod_{j=i}^k (1 - \gamma_{j-1}) = \prod_{j=i}^k \frac{j-1}{c+j-1} = \prod_{j=0}^c \frac{i-1+j}{k+j} \leq \left(\frac{i+c}{k+c}\right)^c$$

Using $1 - \frac{c^p}{(c+k)^p} \leq \exp(-(\frac{c}{c+k})^p)$,

$$\log\left(\prod_{j=i}^k (1 - \frac{c^p}{(c+j)^p})\right) \leq -\sum_{j=i}^k \left(\frac{c}{c+j}\right)^p \leq -\int_i^k \left(\frac{c}{c+j}\right)^p dj = \frac{c^p}{p+1} ((c+i)^{p+1} - (c+k)^{p+1})$$

and therefore

$$\prod_{j=i}^k (1 - \frac{c^p}{(c+k)^p}) \leq \frac{\exp(\frac{c^p}{p+1} (c+i)^{p+1})}{\exp(\frac{c^p}{p+1} (c+k)^{p+1})}$$

which means

$$\prod_{j=i}^k \mu_j \leq \left(\frac{i+c}{k+c}\right)^c \exp\left(\frac{c^p}{p+1} ((c+i)^{p+1} - (c+k)^{p+1})\right).$$

Now we bound the constant term coefficient.

$$\begin{aligned}
 \sum_{i=0}^k \gamma_{i-1} \prod_{j=i}^k \mu_j &\leq \sum_{i=0}^k \underbrace{\frac{c}{(c+i-1)} \left(\frac{i+c}{k+c}\right)^c}_{\max \text{ at } i=0} \underbrace{\exp\left(\frac{c^p}{p+1}((c+i)^{p+1} - (c+k)^{p+1})\right)}_{\leq \frac{1}{(k-i)^{(c+1)(1-p)}}} \\
 &\stackrel{C-S}{\leq} \frac{c}{c-1} \frac{c^c}{(k+c)^c} \sum_{i=0}^{k-1} \frac{1}{(k-i)^{(c+1)(1-p)}} \\
 &\leq \frac{c}{c-1} \frac{c^c}{(k+c)^c} \frac{1}{(c+1)(1-p)-1} \frac{1}{(1-k)^{(c+1)(1-p)-1}} \\
 &\leq \frac{C_1}{(k+c)^p} \frac{1}{(1-k)^{(c+1)(1-p)-1}}
 \end{aligned}$$

where (*) if c is chosen such that $(c+1)(1-p) > 1$ and $C_1 > 0$ big enough. Note that necessarily, $p < 1$, and the size of C_1 depends on how close p is to 1.

Also, to simplify terms,

$$\prod_{j=i}^k \mu_j \leq \left(\frac{i+c}{k+c}\right)^c \underbrace{\exp\left(\frac{c^p}{p+1}((c+i)^{p+1} - (c+k)^{p+1})\right)}_{\leq 1}.$$

Now, we can say

$$\mathbf{g}_k \leq - \sum_{i=0}^{k-1} \beta_{i,i} \left(\frac{i+1+c}{k+c}\right)^c \mathbf{gap}(\mathbf{x}_i) - \beta_{k,k} \mathbf{gap}(\mathbf{x}_k) + \frac{4D^2LC_1}{(k+c)^p} + \left(\frac{c}{k+c}\right)^c \mathbf{g}_0.$$

□

Theorem E.10 (Global rate, $p < 1$). *Suppose $0 < p < 1$ and $c \geq \frac{c-1}{c^p}$. Then $h(\mathbf{x}_k) =: \mathbf{h}_k = O\left(\frac{1}{(k+c)^p}\right)$.*

Proof. Start with

$$\mathbf{g}_k \leq - \sum_{i=0}^{k-1} \beta_{i,i} \left(\frac{i+1+c}{k+c}\right)^c \mathbf{gap}(\mathbf{x}_i) - \beta_{k,k} \mathbf{gap}(\mathbf{x}_k) + \frac{4D^2LC_1}{(k+c)^p} + \left(\frac{c}{k+c}\right)^c \mathbf{g}_0.$$

$$\begin{aligned}
 \mathbf{h}_{k+1} - \mathbf{h}_k &\leq \gamma_k \mathbf{g}_k + 2\gamma_k LD^2 \\
 &\leq - \frac{c}{c+k} \frac{1}{(k+c)^c} \sum_{i=0}^{k-1} \beta_{i,i} (i+c+1)^c \mathbf{h}_i - \frac{c}{c+k} \beta_{k,k} \mathbf{h}_k \\
 &\quad + 2D^2L\gamma_k \left(\frac{2C_1}{(k+c)^p} + \gamma_k\right) + \frac{c}{c+k} \frac{c^c}{(k+c)^c} \mathbf{g}_0 \\
 &\leq - \frac{c^{p+1}}{c+k} \frac{1}{(k+c)^c} \sum_{i=0}^{k-1} (i+c)^{c-p} \mathbf{h}_i - \frac{c}{c+k} \beta_{k,k} \mathbf{h}_k \\
 &\quad + 2D^2L\gamma_k \left(\frac{2C_1}{(k+c)^p} + \gamma_k\right) + \frac{c}{c+k} \frac{c^c}{(k+c)^c} \mathbf{g}_0
 \end{aligned}$$

Suppose $\mathbf{h}_k \leq \frac{C_2}{(k+c)^p}$. Then

$$\begin{aligned} \mathbf{h}_{k+1} - \mathbf{h}_k &\leq \underbrace{-\frac{c^{p+1}}{c+k} \frac{C_2}{(k+c)^c} \sum_{i=0}^{k-1} (c+i)^{c-2p} - \frac{cC_2}{c+k} \frac{c^p}{(c+k)^{2p}}}_{-AC_2} \\ &\quad + \underbrace{2D^2L \frac{c}{c+k} \left(\frac{2C_1}{(k+c)^p} + \frac{c}{c+k} \right) + \frac{c}{c+k} \frac{c^c}{(k+c)^c} \mathbf{g}_0}_B \\ B &= \frac{c}{c+k} \left(2D^2L \left(\frac{2C_1}{(k+c)^p} + \frac{c}{c+k} \right) + \frac{c^c}{(k+c)^c} \mathbf{g}_0 \right) \\ &\leq \frac{2cD^2L(2C_1+c) + c^{c+1} \mathbf{g}_0}{(c+k)^p(c+k)} \\ &=: \frac{C_3}{(c+k)^{p+1}} \end{aligned}$$

where $c > 1$. Then,

$$\begin{aligned} (k+c)^c A &= \frac{c^{p+1}}{c+k} \sum_{i=0}^{k-1} (c+i)^{c-2p} + \frac{c}{c+k} \frac{c^p}{(c+k)^{2p-1}} \\ &\geq \frac{c^{p+1}}{c+k} \frac{(c+k-1)^{c-2p+1} - (c+1)^{c-2p+1}}{c-2p+1} + \frac{c}{c+k} \frac{c^p}{(c+k)^{2p-1}} \\ &= \frac{c^{p+1}}{c-2p+1} \frac{(c+k-1)^{c-2p+1}}{c+k} + O(1/k) \\ &\stackrel{c \geq 2p+1}{\geq} \frac{c^{p+1}}{c-2p+1} + O(1/k) \\ &\stackrel{k \text{ big enough}}{\geq} \frac{c^{p+1}}{c-1} \end{aligned}$$

Therefore,

$$\mathbf{h}_{k+1} \leq \frac{C_2(1 - \frac{c^{p+1}}{c-1})}{(k+c)^p} + \frac{C_3}{(c+k)^{p+1}}.$$

Define $\epsilon = \frac{2c^{p+1}}{c^{p+1}+1-c}$ and pick $C_2 > \frac{C_3}{c\epsilon}$. By assumption, $\epsilon > 0$. Consider $k > K$ such that for all k ,

$$\frac{(k+c+1)^p}{(k+c)^p} \leq 1 + \frac{\epsilon}{2}, \quad \frac{C_3}{c+k} \leq \frac{\epsilon}{2} \frac{(c+k)^p}{(c+k+1)^p}.$$

Then

$$\mathbf{h}_{k+1} \leq \frac{C_2(1 - \frac{\epsilon}{2})}{(k+c+1)^p} + \frac{\frac{\epsilon}{2}}{(k+c+1)^p} \leq \frac{C_2}{(k+c+1)^p}.$$

We have now proved the inductive step. Picking $C_2 \geq \mathbf{h}_0$ gives the starting condition, completing the proof. \square

E.4 Local rates

Lemma E.11 (Local convergence). *Define, for all t ,*

$$\tilde{s}(t) = \operatorname{argmin}_{\tilde{s} \in \mathbf{conv}(S(\mathbf{x}^*))} \|s(t) - \tilde{s}\|_2, \quad \hat{s}(t) = \bar{\beta}_t^{-1} \int_0^t \beta_{t,\tau} \tilde{s}(\tau) d\tau. \quad (11)$$

e.g., $\hat{s}(t)$ is the closest point in the convex hull of the support of \mathbf{x}^* to the point $s(t) = \text{LMO}_{\mathcal{D}}(x(t))$.

Then,

$$\|\bar{s}(t) - \hat{s}(t)\|_2 \leq \frac{c_1}{(c+t)^c}.$$

The proof of this theorem actually does not really depend on how well the FW method works inherently, but rather is a consequence of the averaging. Intuitively, the proof states that after the manifold has been identified, all new support components must also be in the convex hull of the support set of the *optimal* solution; thus, in fact $s_2(t) - \hat{s}(t) = 0$. However, because the accumulation term in the flow actually is not a true average until $t \rightarrow +\infty$, there is a pesky normalization term which must be accounted for. Note, importantly, that this normalization term does not appear in the method, where the accumulation weights always equal 1 (pure averaging).

Proof. First, note that

$$\bar{s}(t) - \hat{s}(t) = \bar{s}(t) - \bar{\beta}_t^{-1} \int_0^t \beta_{t,\tau} \tilde{s}(\tau) d\tau = \int_0^t \beta_{t,\tau} (s(\tau) - \bar{\beta}_t^{-1} \tilde{s}(\tau)) d\tau.$$

Using triangle inequality,

$$\|\bar{s}(t) - \hat{s}(t)\|_2 \leq \underbrace{\left\| \int_0^{\bar{t}} \beta_{t,\tau} (s(\tau) - \bar{\beta}_t^{-1} \tilde{s}(\tau)) d\tau \right\|_2}_{\epsilon_1} + \underbrace{\left\| \int_{\bar{t}}^t \beta_{t,\tau} (s(\tau) - \bar{\beta}_t^{-1} \tilde{s}(\tau)) d\tau \right\|_2}_{\epsilon_2}.$$

Expanding the first term, via Cauchy Schwartz for integrals, we can write, elementwise,

$$\int_0^{\bar{t}} \beta_{t,\tau} (s(\tau)_i - \bar{\beta}_t^{-1} \tilde{s}(\tau)_i) d\tau \leq \int_0^{\bar{t}} \beta_{t,\tau} d\tau \int_0^{\bar{t}} |s(\tau)_i - \bar{\beta}_t^{-1} \tilde{s}(\tau)_i| d\tau$$

and thus,

$$\left\| \int_0^t \beta_{t,\tau} (s(\tau) - \bar{\beta}_t^{-1} \tilde{s}(\tau)) d\tau \right\| \leq \int_0^{\bar{t}} \beta_{t,\tau} d\tau \underbrace{\left\| \int_0^t s(\tau)_i - \bar{\beta}_t^{-1} \tilde{s}(\tau)_i d\tau \right\|_2}_{\leq 2D(1+\bar{\beta}_t^{-1})\bar{t}}.$$

and moreover,

$$\int_0^{\bar{t}} \beta_{t,\tau} d\tau = \int_0^{\bar{t}} \frac{c(c+\tau)^{c-1}}{(c+t)^c} d\tau = \frac{\hat{c}_0}{(c+t)^c}$$

since $\int_0^{\bar{t}} b(b+\tau)^{b-1} d\tau$ does not depend on t . Thus the first error term

$$\epsilon_1 \leq \frac{2\hat{c}_0 D \bar{t} (1 + \bar{\beta}_t^{-1})}{(c+t)^c} \leq \frac{c_0}{(c+t)^c}$$

where

$$\hat{c}_0 := 2D\bar{t} \int_0^{\bar{t}} c(c+\tau)^{c-1} d\tau.$$

In the second error term, *because the manifold has now been identified*, $s(\tau) = \tilde{s}(\tau)$, and so

$$\int_{\bar{t}}^t \beta_{t,\tau} (s(\tau) - \bar{\beta}_t^{-1} \tilde{s}(\tau)) d\tau = \int_{\bar{t}}^t \beta_{t,\tau} (1 - \bar{\beta}_t^{-1}) s(\tau) d\tau$$

and using the same Cauchy-Schwartz argument,

$$\int_{\bar{t}}^t \beta_{t,\tau} (1 - \bar{\beta}_t^{-1}) s(\tau) d\tau \leq D \int_{\bar{t}}^t \beta_{t,\tau} (1 - \bar{\beta}_t^{-1}) d\tau.$$

The term

$$1 - \bar{\beta}_t^{-1} = \left| 1 - \frac{1}{1 - \left(\frac{c}{c+t}\right)^c} \right| = \frac{c^c}{(c+t)^c - c^c} \leq \frac{2c^2}{(c+t)^c}$$

and thus

$$\epsilon_t \leq \int_{\bar{t}}^t \beta_{t,\tau} (1 - \bar{\beta}_t^{-1}) d\tau \leq \frac{2c^3}{(c+t)^{2c}} \int_{\bar{t}}^t (c+\tau)^{c-1} = \frac{2c^2}{(c+t)^{2c}} ((c+t)^c - (c+\bar{t})^c) \leq \frac{2c^2}{(c+t)^c}.$$

Thus,

$$\|\bar{s}(t) - \hat{s}(t)\|_2 \leq \frac{\hat{c}_0 + 2c^2}{(c+t)^c} = O\left(\frac{1}{(c+t)^c}\right).$$

□

Corollary E.12 (Local flow rate). *Suppose that for all $x \in \mathcal{D}$, $\|\nabla f(x)\|_2 \leq G$ for some G large enough. Consider $\gamma(t) = \beta(t) = \frac{c}{c+t}$. Then the ODE*

$$\dot{h}(x(t)) = \gamma(t) \nabla f(x)^T (\bar{s} - x)$$

has solutions $h(t) = O\left(\frac{\log(t)}{(c+t)^c}\right)$ when $t \geq \bar{t}$.

Proof. First, we rewrite the ODE in a more familiar way, with an extra error term

$$\dot{h}(x(t)) = \gamma(t) \nabla f(x)^T (\bar{s} - \hat{s}) + \gamma(t) \nabla f(x)^T (\hat{s} - x)$$

where \hat{s} is as defined in (11). By construction, \hat{s} is a convex combination of $\tilde{s} \in \mathcal{S}(\mathbf{x}^*)$. Moreover, after $t \geq \bar{t}$, $\mathcal{S}(\bar{x}(t)) = \mathcal{S}(\mathbf{x}^*)$, and thus

$$\nabla f(x)^T (\hat{s}(t) - x) = \nabla f(x)^T (s(t) - x) = -\mathbf{gap}(t) \leq -h(t).$$

Then, using Cauchy-Schwartz, and piecing it together,

$$h(t) = \nabla f(x)^T (\hat{s}(t) - x) \leq G\gamma(t) \|\bar{s} - \hat{s}\|_2 - \gamma(t)h(t) \leq \frac{G\gamma(t)c_1}{(c+t)^c} - \gamma(t)h(t).$$

Let us therefore consider the system

$$\dot{h}(x(t)) = \frac{2GD\gamma(t)}{(c+t)^c} - \gamma(t)h(x(t)).$$

The solution to this ODE is

$$h(t) = \frac{h(0)c^c + 2GDc \log(c+t) - 2GDc \log(c)}{(c+t)^c} = O\left(\frac{\log(t)}{(c+t)^c}\right).$$

□

Lemma E.13 (Local averaging error). *Define, for all k ,*

$$\tilde{\mathbf{s}}_k = \underset{\tilde{\mathbf{s}} \in \text{conv}(\mathcal{S}(\mathbf{x}^*))}{\text{argmin}} \|\mathbf{s}_k - \tilde{\mathbf{s}}\|_2, \quad \hat{\mathbf{s}}_k = \bar{\beta}_k^{-1} \sum_{i=1}^k \beta_{k,i} \tilde{\mathbf{s}}_i.$$

e.g., $\tilde{\mathbf{s}}(k)$ is the closest point in the convex hull of the support of \mathbf{x}^ to the point $\mathbf{s}_k = \text{LMO}_{\mathcal{D}}(\mathbf{x}_k)$.*

Then,

$$\|\bar{\mathbf{s}}_k - \hat{\mathbf{s}}_k\|_2 \leq \frac{c_2}{k^c}.$$

Proof. First, note that

$$\bar{\mathbf{s}}_k - \hat{\mathbf{s}}_k = \bar{\mathbf{s}}_k - \sum_{i=1}^k \beta_{k,i} \tilde{\mathbf{s}}_i = \sum_{i=1}^k \beta_{k,i} (\mathbf{s}_i - \tilde{\mathbf{s}}_i).$$

Using triangle inequality,

$$\|\bar{\mathbf{s}}_k - \hat{\mathbf{s}}_k\|_2 \leq \underbrace{\left\| \sum_{i=1}^{\bar{k}} \beta_{k,i} (\mathbf{s}_i - \tilde{\mathbf{s}}_i) \right\|_2}_{\epsilon} + \underbrace{\left\| \sum_{i=\bar{k}}^k \beta_{k,i} (\mathbf{s}_i - \tilde{\mathbf{s}}_i) \right\|_2}_{0}.$$

where the second error term is 0 since the manifold has been identified, so $\tilde{\mathbf{s}}_i = \mathbf{s}_i$ for all $i \geq \bar{k}$.

Expanding the first term, using a Holder norm (1 and ∞ norm) argument,

$$\begin{aligned} \left\| \sum_{i=1}^{\bar{k}} \beta_{k,i} (\mathbf{s}_i - \tilde{\mathbf{s}}_i) \right\|_2 &\leq 2D \sum_{i=1}^{\bar{k}} \beta_{k,i} \\ &= 2D \sum_{i=1}^{\bar{k}} \frac{c}{(c+i)} \prod_{j=0}^{k-i-1} \left(1 - \frac{c}{(c+k-j)}\right) \\ &= 2D \sum_{i=1}^{\bar{k}} \frac{c}{(c+i)} \prod_{j=0}^c \frac{i-1+j}{c+k-j} \\ &\leq 2D \left(\frac{\bar{k}-1+c}{k}\right)^c \sum_{i=1}^{\bar{k}} \frac{c}{(c+i)} = O(1/k^c). \end{aligned}$$

□

Corollary E.14 (Local convergence rate bounds). *Suppose that for all $\mathbf{x} \in \mathcal{D}$, $\|\nabla f(\mathbf{x})\|_2 \leq G$ for some G large enough. Define also r the decay constant of $\|\bar{\mathbf{s}}_k - \mathbf{x}_k\|_2^2 (= O(1/k^r))$. Consider $\gamma_k = \beta_k = \frac{c}{c+k}$. Then the difference equation*

$$\mathbf{h}(\mathbf{x}_{k+1}) - \mathbf{h}(\mathbf{x}_k) \leq \gamma_k \nabla f(\mathbf{x})^T (\bar{\mathbf{s}}_k - \mathbf{x}_k) + \frac{C}{k^r}$$

is satisfied with candidate solution $\mathbf{h}(\mathbf{x}_k) = C_4 \max\{\frac{\log(k)}{(c+k)^c}, \frac{1}{k^{r+1}}\}$ when $k \geq \bar{k}$.

Proof. First, we rewrite the ODE in a more familiar way, with an extra error term

$$\mathbf{h}(\mathbf{x}_{k+1}) - \mathbf{h}(\mathbf{x}_k) = \underbrace{\gamma_k \nabla f(\mathbf{x}_k)^T (\bar{\mathbf{s}}_k - \hat{\mathbf{s}}_k)}_{\leq \gamma_k G \|\bar{\mathbf{s}}_k - \hat{\mathbf{s}}_k\|_2} + \gamma_k \nabla f(\mathbf{x}_k)^T (\hat{\mathbf{s}}_k - \mathbf{x}_k) + \frac{C}{k^{r+2}}$$

where $\hat{\mathbf{s}}_k$ is as defined in (11). By construction, $\hat{\mathbf{s}}_k$ is a convex combination of $\tilde{\mathbf{s}}_i \in \mathcal{S}(\mathbf{x}^*)$. Moreover, after $k \geq \bar{k}$, $\mathcal{S}(\bar{\mathbf{x}}_k) = \mathcal{S}(\mathbf{x}^*)$, and thus

$$\nabla f(\mathbf{x}_k)^T (\hat{\mathbf{s}}_k - \mathbf{x}_k) = \nabla f(\mathbf{x}_k)^T (\mathbf{s}_k - \mathbf{x}_k) = -\mathbf{gap}(\mathbf{x}_k) \leq -\mathbf{h}(\mathbf{x}_k).$$

Then, piecing it together,

$$\mathbf{h}(\mathbf{x}_{k+1}) - \mathbf{h}(\mathbf{x}_k) \leq \gamma_k G \|\bar{\mathbf{s}} - \hat{\mathbf{s}}\|_2 - \gamma_k \mathbf{h}(\mathbf{x}_k) + \frac{C}{k^{r+2}} \stackrel{\text{Lemma E.13}}{\leq} \underbrace{\gamma_k \frac{GC_2}{k^c}}_{\leq C_3/k^{c+1}} - \gamma_k \mathbf{h}(\mathbf{x}_k) + \frac{C}{k^{r+2}}$$

Recursively, we can now show that for $C_4 \geq C + C_3$, if

$$\mathbf{h}(\mathbf{x}_k) \leq C_4 \max\left\{\frac{\log(k)}{(c+k)^c}, \frac{1}{k^{r+1}}\right\}$$

then,

$$\begin{aligned} \mathbf{h}(\mathbf{x}_{k+1}) &\leq \frac{C_3}{k^{c+1}} + \frac{C}{k^{r+1}} + (1 - \gamma_k) \mathbf{h}(\mathbf{x}_k) \\ &\leq \frac{C_3}{k^{c+1}} + \frac{C}{k^{r+1}} + \frac{k}{c+k} C_4 \max\left\{\frac{\log(k)}{(c+k)^c}, \frac{1}{k^{r+1}}\right\}. \end{aligned}$$

If $c \leq r + 1$ then

$$\mathbf{h}(\mathbf{x}_{k+1}) \leq \frac{C + C_3}{k^c} + \frac{k}{c+k} \frac{C_4 \log(k)}{(c+k)^c} \leq \frac{k}{c+k} \frac{C_4 \log(k)}{(c+k)^c} \leq \frac{C_4 \log(k+1)}{(c+k+1)^c}$$

for k large enough. Otherwise,

$$\mathbf{h}(\mathbf{x}_{k+1}) \leq \frac{C + C_3}{k^{r+2}} + \frac{k}{c+k} \frac{C_4}{k^{r+1}} \leq \frac{C_4}{(k+1)^{r+1}}.$$

for k large enough. □

Theorem E.15 (Local convergence rate). *Picking $c \geq 3p/2 + 1$, the proposed method AVGFW has an overall convergence $\mathbf{h}(\mathbf{x}_k) = O(k^{-3p/2})$.*

Proof. Putting together Theorem E.10, Lemma E.13, and Corollary E.14, we can resolve the constants

$$q = p, \quad r = \min\{q/2 + p - 1, 2p\} = \frac{3p}{2} - 1$$

and resolves an overall convergence bound of $\mathbf{h}(\mathbf{x}_k) = O(k^{-3p/2})$. □