



Projection free methods on product domains

Immanuel Bomze¹ · Francesco Rinaldi²  · Damiano Zeffiro²

Received: 1 December 2023 / Accepted: 18 May 2024
© The Author(s) 2024

Abstract

Projection-free block-coordinate methods avoid high computational cost per iteration, and at the same time exploit the particular problem structure of product domains. Frank–Wolfe-like approaches rank among the most popular ones of this type. However, as observed in the literature, there was a gap between the classical Frank–Wolfe theory and the block-coordinate case, with no guarantees of linear convergence rates even for strongly convex objectives in the latter. Moreover, most of previous research concentrated on convex objectives. This study now deals also with the non-convex case and reduces above-mentioned theory gap, in combining a new, fully developed convergence theory with novel active set identification results which ensure that inherent sparsity of solutions can be exploited in an efficient way. Preliminary numerical experiments seem to justify our approach and also show promising results for obtaining global solutions in the non-convex case.

Keywords Projection free optimization · First order optimization · Block coordinate descent

Mathematics Subject Classification 90C06 · 90C26 · 90C30

✉ Francesco Rinaldi
rinaldi@math.unipd.it

Immanuel Bomze
immanuel.bomze@univie.ac.at

Damiano Zeffiro
damiano.zeffiro@studenti.unipd.it

¹ Vienna Center of OR and Research Network Data Science, Universität Wien, Vienna, Austria

² Department of Mathematics, University of Padua, Padua, Italy

1 Introduction

We consider the problem

$$\min_{\mathbf{x} \in \mathcal{C}} f(\mathbf{x}), \quad (1)$$

with objective f having L -Lipschitz regular gradient, and feasible set $\mathcal{C} \subseteq \mathbb{R}^n$ closed and convex. Furthermore, we assume that \mathcal{C} is block separable, that is

$$\mathcal{C} = \mathcal{C}_{(1)} \times \dots \times \mathcal{C}_{(m)} \quad (2)$$

with $\mathcal{C}_{(i)} \subset \mathbb{R}^{n_i}$ closed and convex for $i \in [1:m]$, and of course $\sum_{i=1}^m n_i = n$.

Notice that problem (1) falls in the class of composite optimization problems

$$\min_{\mathbf{x} \in \mathcal{C}} [f(\mathbf{x}) + g(\mathbf{x})] \quad (3)$$

with f smooth and $g(\mathbf{x}) = \sum_{i=1}^m \chi_{\mathcal{C}_{(i)}}(\mathbf{x}^{(i)})$ convex and block separable (see, e.g., [37] for an overview of methods for this class of problems); here $\chi_D : \mathbb{R}^d \rightarrow [0, +\infty]$ denotes the indicator function of a convex set $D \subseteq \mathbb{R}^d$, and for a block vector $\mathbf{x} \in \mathbb{R}^n = \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_m}$ we denote by $\mathbf{x}^{(i)} \in \mathbb{R}^{n_i}$ the component corresponding to the i -th block, so that $\mathbf{x} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)})$.

Problems of this type arise in a wide number of real-world applications like, e.g., traffic assignment [31], structural SVMs [28], trace-norm based tensor completion [32], reduced rank nonparametric regression [22], semi-relaxed optimal transport [23], structured submodular minimization [26], group fused lasso [1], and dictionary learning [18].

Block-coordinate gradient descent (BCGD) strategies (see, e.g., [4]) represent a standard approach to solve problem (1) in the convex case. When dealing with non-convex objectives, those methods can anyway still be used as an efficient tool to perform local searches in probabilistic global optimization frameworks (see, e.g., [33] for further details). The way BCGD approaches work is very easy to understand: those methods build up, at each iteration, a suitable model of the original function for a block of variables and then perform a projection on the feasible set related to that block.

Projection-based strategies (see, e.g., [6, 19, 21, 40] for further details, with significant contributions by Daniela di Serafino) are in practice widely used also in a block-coordinate fashion (see, e.g., [35]). However, they might be costly even when the projection is performed over some structured sets like, e.g., the flow polytope, the nuclear-norm ball, the Birkhoff polytope, or the permutahedron (see, e.g., [20]). This is the reason why, in recent years, projection-free methods (see, e.g., [14, 25, 29]) have been massively used when dealing with those structured constraints.

These methods simply rely on a suitable oracle that minimizes, at each iteration, a linear approximation of the function over the original feasible set, returning a point in

$$\operatorname{argmin}_{\mathbf{x} \in \mathcal{C}} \langle \mathbf{g}, \mathbf{x} \rangle.$$

When \mathcal{C} is defined as in (2), this decomposes in m independent problems thanks to the block separable structure of the feasible set. In turn, the resulting problems on the blocks can then be solved in parallel, a possibility that has widely been explored in the literature, especially in the context of traffic assignment (see, e.g., [31]). In a big data context, performing a full update of the variables might still represent a computational bottleneck that needs to be properly handled in practice. This is the reason why block-coordinate variants of the classic Frank–Wolfe (FW) method have been recently proposed (see, e.g., [28, 36, 41]). This method is proposed in [28] for structured support vector machine training, and randomly selects a block at each iteration to perform an FW update on the block. Several improvements on this algorithm, e.g., adaptive block sampling, use of pairwise and away-step directions, or oracle call caching, are described in [36], which obviously work in a sequential fashion.

However, in case one wants to take advantage of modern multicore architectures or of distributed clusters, parallel and distributed versions of the block-coordinate FW algorithm are also available [41]. It is important to highlight that all the papers mentioned above only consider convex programming problems and use random sampling variants as the main block selection strategy.

Furthermore, as noticed in [36], the standard convergence analysis for FW variants (e.g., pairwise and away step FW) cannot be easily extended to the block-coordinate case. In particular, there has been no extension in this setting of the well known linear convergence rate guarantees for FW variants applied to strongly convex objectives (see [14] and references therein). This is mainly due to the difficulties in handling the bad/short steps (i.e., those steps that do not give a good progress and are taken to guarantee feasibility of the iterate) within a block-coordinate framework. In [36], the authors hence extend the convergence analysis of FW variants to the block coordinate setting under the strong assumption that there are no bad steps, claiming that novel proof techniques are required to carry out the analysis in general and close the gap between FW and BCFW in this context.

Here we focus on the non-convex case and define a new general block-coordinate algorithmic framework that gives flexibility in the use of both block selection strategies and FW-like directions. Such a flexibility is mainly obtained thanks to the way we perform approximate minimizations in the blocks. At each iteration, after selecting one block at least, we indeed use the Short Step Chain (SSC) procedure described in [39], which skips gradient computations in consecutive short steps until proper conditions are satisfied, to get the approximate minimization done in the selected blocks.

Concerning the block selection strategies, we explore three different options. The first one we consider is a parallel or Jacobi-like strategy (see, e.g., [5]), where the SSC procedure is performed for all blocks. This obviously reduces the computational burden with respect to the use of the SSC in the whole variable space (see, e.g., [39]) and eventually enables to use multicore architectures to perform those tasks in parallel. The second one is the random sampling (see, e.g., [28]), where the SSC procedure is performed at each iteration on a randomly selected subset of blocks. Finally we have a variant of the Gauss–Southwell rule (see, e.g., [34]), where we perform SSC in all blocks and then select a block which violates optimality conditions at most.

Such a greedy rule may make more progress in the objective function, since it uses first order information to choose the right block, but is, in principle, more expensive than the other options we mentioned before (notice that the SSC is performed, at each iteration, for all blocks).

Furthermore, we consider the following projection-free strategies: Away-step Frank–Wolfe (AFW), Pairwise Frank–Wolfe (PFW), and Frank–Wolfe method with in face directions (FDFW), see, e.g., [39] and references therein for further details. The AFW and PFW strategies depend on a set of “elementary atoms” A such that $\mathcal{C} = \text{conv}(A)$. Given A , for a base point $\mathbf{x} \in \mathcal{C}$ we can define

$$S_{\mathbf{x}} = \{S \subset A : \mathbf{x} \text{ is a proper convex combination of all the elements in } S\},$$

the family of possible active sets for a given point \mathbf{x} . For $\mathbf{x} \in \mathcal{C}$ and $S \in S_{\mathbf{x}}$, \mathbf{d}^{PFW} is a PFW direction with respect to the active set S and gradient $-\mathbf{g}$ if and only if

$$\mathbf{d}^{\text{PFW}} = \mathbf{s} - \mathbf{q} \text{ with } \mathbf{s} \in \operatorname{argmax}_{\mathbf{s} \in \mathcal{C}} \langle \mathbf{s}, \mathbf{g} \rangle \text{ and } \mathbf{q} \in \operatorname{argmin}_{\mathbf{q} \in S} \langle \mathbf{q}, \mathbf{g} \rangle. \quad (4)$$

Similarly, given $\mathbf{x} \in \mathcal{C}$ and $S \in S_{\mathbf{x}}$, \mathbf{d}^{AFW} is an AFW direction with respect to the active set S and gradient $-\mathbf{g}$ if and only if

$$\mathbf{d}^{\text{AFW}} \in \operatorname{argmax}\{\langle \mathbf{g}, \mathbf{d} \rangle : \mathbf{d} \in \{\mathbf{d}^{\text{FW}}, \mathbf{d}^{\text{AS}}\}\}, \quad (5)$$

where \mathbf{d}^{FW} is a classic Frank–Wolfe direction

$$\mathbf{d}^{\text{FW}} = \mathbf{s} - \mathbf{x} \text{ with } \mathbf{s} \in \operatorname{argmax}_{\mathbf{s} \in \mathcal{C}} \langle \mathbf{s}, \mathbf{g} \rangle, \quad (6)$$

and \mathbf{d}^{AS} is the away direction

$$\mathbf{d}^{\text{AS}} = \mathbf{x} - \mathbf{q} \text{ with } \mathbf{q} \in \operatorname{argmin}_{\mathbf{q} \in S} \langle \mathbf{q}, \mathbf{g} \rangle. \quad (7)$$

The FDFW only requires the current point \mathbf{x} and gradient $-\mathbf{g}$ to select a descent direction (i.e., it does not need to keep track of the active set) and is defined as

$$\mathbf{d}^F = \mathbf{x} - \mathbf{x}_F \text{ with } \mathbf{x}_F \in \operatorname{argmin}\{\langle \mathbf{g}, \mathbf{y} \rangle : \mathbf{y} \in \mathcal{F}(\mathbf{x})\}$$

for $\mathcal{F}(\mathbf{x})$ the minimal face of \mathcal{C} containing \mathbf{x} . The selection criterion is then analogous to the one used by the AFW:

$$\mathbf{d}^{\text{FD}} \in \operatorname{argmax}\{\langle \mathbf{g}, \mathbf{d} \rangle : \mathbf{d} \in \{\mathbf{d}^F, \mathbf{d}^{\text{FW}}\}\}. \quad (8)$$

From a theoretical point of view, this new algorithmic framework enables us to give:

- a local linear convergence rate for any choice of block selection strategy and FW-like direction. This result is obtained under a Kurdyka–Łojasiewicz (KL) property (see, e.g., [3, 7, 8]) and a tailored angle condition (see, e.g., [39]). Thanks to the

way we handle short steps in our framework we are thus able to extend the analysis given for FW variants to the block-coordinate case and then to close the relevant gap in the theory highlighted in [36].

- a local active set identification result (see, e.g., [12, 13, 15, 24]) for a specific structure of the Cartesian product defining the feasible set \mathcal{C} , suitable choices of projection-free strategy (i.e., AFW direction is used), and general smooth non convex objectives. In particular, we prove that our framework identifies in finite time the support of a solution. Such a theoretical feature allows to reduce the dimension of the problem at hand and, consequently, the overall computational cost of the optimization procedure.

This is, to the best of our knowledge, the first time that both a (bad step free) linear convergence rate and an active set identification result are given for block-coordinate FW variants. In particular, we solve the open question from [36] discussed above, proving that the linear convergence rate of FW variants can indeed be extended to the block coordinate setting. Furthermore, our results guarantee, for the first time in the literature of projection free optimization methods, identification of the local active set in a *single iteration* without a tailored active set strategy.

We also report some preliminary numerical results on a specific class of structured problems with a block separable feasible set. Those results show that the proposed framework outperforms the classic block-coordinate FW and, thanks to its flexibility, it can be effectively embedded into a probabilistic global optimization framework thus significantly boosting its performances.

The paper is organized as follows. Section 2 describes the details of our new algorithmic framework. An in-depth analysis of its convergence properties is reported in Sect. 3. An active set identification result is reported in Sect. 4. Preliminary numerical results, focusing on the computational analysis of both the local identification and the convergence properties of our framework, are reported in Sect. 5. Finally, some concluding remarks are included in Sect. 6.

1.1 Notation

For a closed and convex set $C \subset \mathbb{R}^h$ we denote by $\pi(C, \mathbf{x})$ the Euclidean projection of $\mathbf{x} \in \mathbb{R}^h$ onto C , and by $T_C(\mathbf{x})$ the tangent cone to C at $\mathbf{x} \in C$, itself again a closed convex set:

$$T_C(\mathbf{x}) = \text{closure}\{t(\mathbf{v} - \mathbf{x}) : t \geq 0, \mathbf{v} \in C\}.$$

For $\mathbf{g} \in \mathbb{R}^h$ we also use $\pi_{\mathbf{x}}(\mathbf{g})$ as a shorthand for $\|\pi(T_C(\mathbf{x}), \mathbf{g})\|$. We denote by $\hat{\mathbf{y}}$ the vector $\frac{\mathbf{y}}{\|\mathbf{y}\|}$ for $\mathbf{y} \neq \mathbf{0}$, and $\hat{\mathbf{y}} = \mathbf{0}$ otherwise. We finally denote by $\bar{B}_r(\mathbf{x})$ and $B_r(\mathbf{x})$ the closed and open balls of radius r centered at \mathbf{x} .

2 A new block-coordinate projection-free method

The block-coordinate framework we consider here applies the Short Step Chain (SSC) procedure from [39], described below as Algorithm 2, to some of the blocks at every iteration. A detailed scheme is specified as Algorithm 1; recall notation $\mathbf{x} = (\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)})$ with $\mathbf{x}^{(i)} \in \mathcal{C}_{(i)}$, all $i \in [1:m]$.

Algorithm 1 Block coordinate method with Short Step Chain (SSC) procedure

- 1: $\mathbf{x}_0 \in \mathcal{C}$, $k = 0$.
 - 2: If \mathbf{x}_k is stationary, then STOP
 - 3: Choose $\mathcal{M}_k \subset [1:m]$.
 - 4: For all $i \notin \mathcal{M}_k$ set $\mathbf{x}_{k+1}^{(i)} = \mathbf{x}_k^{(i)}$
 - 5: For all $i \in \mathcal{M}_k$ set $\mathbf{x}_{k+1}^{(i)} = \text{SSC}(\mathbf{x}_k^{(i)}, -\nabla f(\mathbf{x}_k)^{(i)})$
 - 6: $k = k + 1$. Go to step 2.
-

In Algorithm 1, we perform two main operations at each iteration. First, in Step 3, we pick a suitable subset of blocks \mathcal{M}_k according to a given block selection strategy. We then update (Steps 4 and 5) the variables related to the selected blocks by means of the SSC procedure, while keeping all the variables in the other blocks unchanged.

We now briefly recall the SSC procedure from [39], designed to recycle the gradient in consecutive bad steps until suitable stopping conditions are met, in Algorithm 2. Here and in the sequel we denote by $\alpha_{\max}(\mathbf{y}_j, \mathbf{d}_j)$ the set of feasible step sizes at \mathbf{y}_j in direction \mathbf{d}_j .

Algorithm 2 Short Step Chain procedure – SSC($\bar{\mathbf{x}}$, \mathbf{g})

- 1: **Initialization.** $\mathbf{y}_0 = \bar{\mathbf{x}}$, $j = 0$
 - 2: Select $\mathbf{d}_j \in \mathcal{A}(\mathbf{y}_j, \mathbf{g})$, $\alpha_{\max}^{(j)} \in \alpha_{\max}(\mathbf{y}_j, \mathbf{d}_j)$,
 - 3: **if** $\mathbf{d}_j = 0$ **then return** \mathbf{y}_j
 - 4: **end if**
 - 5: compute an auxiliary step size β_j
 - 6: let $\alpha_j = \min(\alpha_{\max}^{(j)}, \beta_j)$
 - 7: $\mathbf{y}_{j+1} = \mathbf{y}_j + \alpha_j \mathbf{d}_j$
 - 8: **if** $\alpha_j = \beta_j$ **then return** \mathbf{y}_{j+1}
 - 9: **end if**
 - 10: $j = j + 1$, go to Step 2
-

By \mathcal{A} we indicate a projection-free strategy to generate first-order feasible descent directions for smooth functions on the block where the SSC is applied (e.g., FW, PFW, AFW directions). Since the gradient, $-\mathbf{g}$, is constant during the SSC procedure, it is easy to see that the procedure represents an application of \mathcal{A} to minimize the linearized objective $f_{\mathbf{g}}(\mathbf{z}) = \langle -\mathbf{g}, \mathbf{z} - \bar{\mathbf{x}} \rangle + f(\bar{\mathbf{x}})$, with suitable stepsizes and stopping condition. More specifically, after a stationarity check (see Steps 2–4), the stepsize α_j is the minimum of an auxiliary stepsize $\beta_j > 0$ and the maximal stepsize $\alpha_{\max}^{(j)}$ (which we always assume to be strictly positive). The point \mathbf{y}_{j+1} generated at Step 7 is always

feasible since $\alpha_j \leq \alpha_{\max}^{(j)}$. Notice that if the method \mathcal{A} used in the SSC performs a FW step (see equation (6) for the definition of FW step), then the SSC terminates, with $\alpha_j = \beta_j$ or with \mathbf{y}_{j+1} a global minimizer of $f_{\mathbf{g}}$.

The auxiliary step size β_j (see Step 5 of the SSC procedure) is thus defined as the maximal feasible stepsize (at \mathbf{y}_j) for the trust region

$$\Omega_j = B_{\|\mathbf{g}\|/2L}(\bar{\mathbf{x}} + \frac{\mathbf{g}}{2L}) \cap B_{\langle \mathbf{g}, \hat{\mathbf{d}}_j \rangle / L}(\bar{\mathbf{x}}). \tag{9}$$

This guarantees the sufficient decrease condition

$$f(\mathbf{y}_j) \leq f(\mathbf{x}_k) - \frac{L}{2} \|\mathbf{x}_k - \mathbf{y}_j\|^2 \tag{10}$$

and hence a monotone decrease of f in the SSC. For further details see [39].

2.1 Block selection strategies

As briefly mentioned in the introduction, we will consider three different block selection strategies in our analysis. The first one is a parallel or Jacobi-like strategy (see, e.g., [5]). In this case, we select all the blocks at each iteration. As we already observed, this is computationally cheaper than handling the whole variable space at once. Furthermore, multicore architectures might eventually be considered to perform those tasks in parallel. A definition of the strategy is given below:

Definition 1 (Parallel selection) Set $\mathcal{M}_k = [1 : m]$.

The second strategy is a variant of the GS rule (see, e.g., [34]), where we first perform SSC in all blocks and then select a block that violates optimality conditions at most. The formal definition is reported below.

Definition 2 (Gauss–Southwell (GS) selection) Set $\mathcal{M}_k = \{i(k)\}$, with

$$i(k) \in \operatorname{argmax}_{i \in [1:m]} \langle \mathbf{g}^{(i)}, \operatorname{SSC}(\mathbf{x}_k^{(i)}, -\nabla f(\mathbf{x}_k)^{(i)}) - \mathbf{x}_k^{(i)} \rangle.$$

Finally, we have random sampling (see, e.g., [28]). Here we randomly generate one index at each iteration with uniform probability distribution. The definition we have in this case is the following:

Definition 3 (Random sampling) Set $\mathcal{M}_k = \{i(k)\}$, with $i(k)$ index chosen uniformly at random in $[1 : m]$.

3 Convergence analysis

In this section, we analyze the convergence properties of our algorithmic framework. In particular, we show that under a suitably defined angle condition on the blocks and

a local KL condition on the objective function, we get, for any block selection strategy used, a linear convergence rate. The convergence analysis presented in this section extends the results given in [39] to the block coordinate setting, a demanding task which is by no means straightforward. Hence, some novel arguments are required for this extension, which are now introduced, and then described in detail in the appendix.

Our convergence framework makes use of the angle condition introduced in [38, 39]. Such a condition ensures that the slope of the descent direction selected by the method is optimal up to a constant. We now recall this angle condition. For $\mathbf{x} \in \mathcal{C}$ and $\mathbf{g} \in \mathbb{R}^n$ we first define the directional slope lower bound as

$$\text{DSB}_{\mathcal{A}}(\mathcal{C}, \mathbf{x}, \mathbf{g}) = \inf_{\mathbf{d} \in \mathcal{A}(\mathbf{x}, \mathbf{g})} \frac{\langle \mathbf{g}, \mathbf{d} \rangle}{\pi_{\mathbf{x}}(\mathbf{g}) \|\mathbf{d}\|}, \quad (11)$$

if \mathbf{x} is not stationary for $-\mathbf{g}$, otherwise we set $\text{DSB}_{\mathcal{A}}(\mathcal{C}, \mathbf{x}, \mathbf{g}) = 1$. Given a subset P of \mathcal{C} , we then define the slope lower bound as

$$\text{SB}_{\mathcal{A}}(\mathcal{C}, P) = \inf_{\substack{\mathbf{g} \in \mathbb{R}^n \\ \mathbf{x} \in P}} \text{DSB}_{\mathcal{A}}(\mathcal{C}, \mathbf{x}, \mathbf{g}) = \inf_{\mathbf{g}: \pi_{\mathbf{x}}(\mathbf{g}) \neq 0, \mathbf{x} \in P} \text{DSB}_{\mathcal{A}}(\mathcal{C}, \mathbf{x}, \mathbf{g}). \quad (12)$$

We use $\text{SB}_{\mathcal{A}}(\mathcal{C})$ as a shorthand for $\text{SB}_{\mathcal{A}}(\mathcal{C}, \mathcal{C})$, and say that the angle condition holds for the method \mathcal{A} if

$$\text{SB}_{\mathcal{A}}(\mathcal{C}) = \tau > 0. \quad (13)$$

Remark 1 AFW, PFW and FDFW all satisfy the angle condition, when \mathcal{C} is a polytope. A detailed proof of this result is reported in [39], together with some other examples of methods satisfying the angle condition for convex sets with smooth boundary described in [38].

We now report the local KL condition used to analyze the convergence of our algorithm. The same condition was used previously as well [39, Assumption 2.1].

Assumption 1 Given a stationary point $\mathbf{x}_* \in \mathcal{C}$, there exists $\eta, \delta, \mu > 0$ such that for every $\mathbf{x} \in B_{\delta}(\mathbf{x}_*)$ with $f(\mathbf{x}_*) < f(\mathbf{x}) < f(\mathbf{x}_*) + \eta$ we have

$$\pi_{\mathbf{x}}(-\nabla f(\mathbf{x})) \geq \sqrt{2\mu} [f(\mathbf{x}) - f(\mathbf{x}_*)]^{\frac{1}{2}}. \quad (14)$$

When dealing with convex programming problems, a Hölderian error bound with exponent 2 on the solution set implies condition (14), see [9, Corollary 6]. Therefore, our assumption holds when dealing with μ -strongly convex functions (see, e.g., [27]), and in particular for the setting of the open question from [36] discussed in the introduction. It is however important to highlight that the error bound (14) holds in a variety of both convex and non-convex settings (see [39] for a detailed discussion on this matter). An interesting example for our analysis is the setting where f is (non-convex) quadratic, i.e., $f(\mathbf{x}) = \mathbf{x}^{\top} \mathbf{Q} \mathbf{x} + \mathbf{b}^{\top} \mathbf{x}$, and \mathcal{C} is a polytope.

We now report our main convergence result.

Theorem 1 *Let Assumption 1 hold at \mathbf{x}_* . Let us consider the sequence $\{\mathbf{x}_k\}$ generated by Algorithm 1. Assume that:*

- the angle condition (13) holds in every block for the same $\tau > 0$;
- the SSC procedure always terminates in a finite number of steps.
- $f(\mathbf{x}_*)$ is a minimum in the connected component of $\{\mathbf{x} \in \mathcal{C} : f(\mathbf{x}) \leq f(\mathbf{x}_0)\}$ containing \mathbf{x}_0 .

Then, there exists $\tilde{\delta} > 0$ such that, if $\mathbf{x}_0 \in B_{\tilde{\delta}}(\mathbf{x}_*)$:

- for the parallel block selection strategy, we have

$$f(\mathbf{x}_k) - f(\mathbf{x}_*) \leq (q_P)^k [f(\mathbf{x}_0) - f(\mathbf{x}_*)], \tag{15}$$

and $\mathbf{x}_k \rightarrow \tilde{\mathbf{x}}_*$ with

$$\|\mathbf{x}_k - \tilde{\mathbf{x}}_*\| \leq \frac{\sqrt{2 - 2q_P}}{\sqrt{L}(1 - \sqrt{q_P})} (q_P)^{\frac{k}{2}} [f(\mathbf{x}_0) - f(\tilde{\mathbf{x}}_*)], \tag{16}$$

for

$$q_P = 1 - \frac{\mu\tau^2}{4L(1 + \tau)^2}. \tag{17}$$

- for the GS block selection strategy, we have

$$f(\mathbf{x}_k) - f(\mathbf{x}_*) \leq (q_{GS})^k [f(\mathbf{x}_0) - f(\mathbf{x}_*)], \tag{18}$$

and $\mathbf{x}_k \rightarrow \tilde{\mathbf{x}}_*$ with

$$\|\mathbf{x}_k - \tilde{\mathbf{x}}_*\| \leq \frac{\sqrt{2 - 2q_{GS}}}{\sqrt{L}(1 - \sqrt{q_{GS}})} (q_{GS})^{\frac{k}{2}} [f(\mathbf{x}_0) - f(\tilde{\mathbf{x}}_*)], \tag{19}$$

for

$$q_{GS} = 1 - \frac{\mu\tau^2}{4mL(1 + \tau)^2}, \tag{20}$$

- for the random block selection strategy we have, under the additional condition that

$$\min\{f(\mathbf{x}) : \|\mathbf{x} - \mathbf{x}_*\| = \delta\} > f(\mathbf{x}_*) \tag{21}$$

holds for some $\delta > 0$, that

$$\mathbb{E}[f(\mathbf{x}_k) - f(\mathbf{x}_*)] \leq (q_R)^k [f(\mathbf{x}_0) - f(\mathbf{x}_*)], \tag{22}$$

and $\mathbf{x}_k \rightarrow \tilde{\mathbf{x}}_*$ almost surely with

$$\mathbb{E}[\|\mathbf{x}_k - \tilde{\mathbf{x}}_*\|] \leq \frac{\sqrt{2 - 2q_R}}{\sqrt{L}(1 - \sqrt{q_R})} (q_R)^{\frac{k}{2}} [f(\mathbf{x}_0) - f(\tilde{\mathbf{x}}_*)] \quad (23)$$

for $q_R = q_{GS}$.

This convergence result extends [39, Theorem 4.2] to our block coordinate setting. However, since the SSC is here applied independently to different blocks, we cannot directly apply the results from [39]. Instead, we combine the properties of the SSC applied in single blocks by exploiting the structure of the tangent cone for product domains:

$$T_{\mathcal{C}}(\mathbf{x}) = T_{\mathcal{C}_{(1)}}(\mathbf{x}^{(1)}) \times \dots \times T_{\mathcal{C}_{(m)}}(\mathbf{x}^{(m)}). \quad (24)$$

This requires proving stronger properties for the sequence generated by the SSC than those presented in [39]. The details with references to relevant results from [39] can be found in the appendix. Finite termination of the SSC procedure is instead directly ensured by the results proved in [38, 39], in particular for the AFW, PFW and FDFW applied on polytopes.

Remark 2 If the feasible set \mathcal{C} is a polytope and if we assume that the objective function f satisfies condition (14) on every point generated by the algorithm, with fixed $f(\mathbf{x}_*)$, then Algorithm 1 with AFW (PFW or FDFW) in the SSC converges at the rates given above. Condition (14) holds in case of μ -strongly convex functions, and hence we have that in those cases our algorithm globally converges with the rates given in Theorem 1.

Remark 3 Both the parallel and the GS strategy give the same rate with different constants. In particular, the constant ruling the GS case depends on the number of blocks used (the larger the number of blocks, the worse the rate) and is larger than the one we have for the parallel case.

Remark 4 The random block selection strategy has the same rate as the GS strategy, but it is given in expectation. In particular, the constant ruling the rate is the same as the GS one, hence depends on the number of blocks used. Note that a further technical assumption (21) on \mathbf{x}_* is needed in this case.

4 Active set identification

We now report an active set identification result for our framework, assuming that the sets in the Cartesian product have a specific structure:

$$\mathcal{C} = \Delta^{n_1} \times \dots \times \Delta^{n_m}, \quad (25)$$

so that the set $\mathcal{C}_{(i)}$ is the $(n_i - 1)$ -dimensional standard simplex

$$\Delta^{n_i} = \{\mathbf{x} \in \mathbb{R}_+^{n_i} : \mathbf{x}^\top \mathbf{e}^{(i)} = 1\}, \quad i \in [1:m],$$

for $\mathbf{e} \in \mathbb{R}^n$ the vector with components all equal to 1. We only focus on Algorithm 1 with AFW in the SSC and assume that strict complementarity holds at a stationary point \mathbf{x}_* , i.e., either $(x_*)_i = 0$ or $\frac{\partial f}{\partial x_i}(\mathbf{x}_*) = \langle \mathbf{x}_*, \nabla f(\mathbf{x}_*) \rangle$ holds but not both simultaneously, for all i . As usual, the support $\text{supp } \mathbf{x} = \{i \in [1:n] : x_i > 0\}$. We now report our main identification result. A detailed proof is included in the appendix.

Theorem 2 *Under the above assumptions on \mathcal{C} , let $\mathcal{A}^{(i)}$ be the AFW for $i \in [1:m]$, and let strict complementarity conditions hold at $\mathbf{x}_* \in \mathcal{C}$.*

- *If $\{\mathbf{x}_k\}$ is generated by Algorithm 1 with parallel selection, then there exists a neighborhood U of \mathbf{x}_* such that if $\mathbf{x}_k \in U$ then $\text{supp}(\mathbf{x}_{k+1}) = \text{supp}(\mathbf{x}_*)$.*
- *If $\{\mathbf{x}_k\}$ is generated by Algorithm 1 with randomized or GS selection, then there exists a neighborhood U of \mathbf{x}_* such that if $\mathbf{x}_k \in U$ then $\text{supp}(\mathbf{x}_{k+1}^{i(k)}) = \text{supp}(\mathbf{x}_*^{i(k)})$.*

When the sequence generated by our algorithm converges to the point \mathbf{x}_* , it is then easy to see that the support of the iterate matches the final support of \mathbf{x}_* for k large enough.

Corollary 1 *Under the above assumptions on \mathcal{C} , let $\mathcal{A}^{(i)}$ be the AFW for $i \in [1:m]$, and let strict complementarity conditions hold at $\mathbf{x}_* \in \mathcal{C}$. If $\mathbf{x}_k \rightarrow \mathbf{x}_*$ (almost surely), then for parallel and GS selection (for random sampling) we have $\text{supp}(\mathbf{x}_k) = \text{supp}(\mathbf{x}_*)$ for k large enough.*

This result has relevant practical implications, especially when handling sparse optimization problems. Since the algorithm iterates have a constant support when k is large, we can simply focus on the few support components and forget about the others in this case. We hence can exploit this by embedding sophisticated tools (like, e.g., caching strategies, second-order methods) in the algorithm, thus obtaining a significant speed up in the end.

5 Numerical results

We report here some preliminary numerical results for a non-convex quadratic optimization problem referred to as Multi-StQP [16] on a product of (here identical) simplices, that is

$$\min \left\{ \mathbf{x}^\top \mathbf{Q} \mathbf{x} : \mathbf{x} \in (\Delta^l)^m \right\}. \tag{26}$$

The matrix \mathbf{Q} was generated in such a way that the solutions of problem (26) had components sparse but different from vertices. This is in fact the setting where FW variants have proved to be more effective [15, 39]. In order to obtain the desired property, we consider a perturbation of a stochastic StQP [11]. Given $\{\bar{\mathbf{Q}}_i\}_{i \in [1:m]}$ representing m possible StQPs, with $\bar{\mathbf{Q}}_i \in \mathbb{R}^{l \times l}$ for $i \in [1:m]$, the corresponding stochastic StQP with sample space $[1:m]$ is given by

$$\max \left\{ \sum_{i=1}^m p_i \mathbf{y}_i^\top \bar{\mathbf{Q}}_i \mathbf{y}_i : \mathbf{y}_i \in \Delta^l \text{ for all } i \in [1:m] \right\}. \tag{27}$$

with p_i probability of the StQP i . Equivalently, (27) is an instance of problem (26) with $Q = \bar{Q}$, for

$$\bar{Q} = \begin{bmatrix} -p_1\bar{Q}_1 & 0 & \cdots & 0 \\ 0 & -p_2\bar{Q}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & -p_m\bar{Q}_m \end{bmatrix}. \quad (28)$$

In our tests, we added to the stochastic StQP a perturbation coupling the blocks. More precisely, the matrix Q was set equal to $\bar{Q} + \varepsilon\tilde{Q}$, for \tilde{Q} a random matrix with standard Gaussian independent entries. The coefficient ε was set equal to $\frac{1}{2m^2}$. We set $\bar{Q}_i = \bar{A}_i + \alpha I_i$, for $\alpha = 0.5$ and \bar{A}_i the adjacency matrix of an Erdős-Rényi random graph, where each couple of vertices has probability p of being connected by an edge, independently from the other couples. Hence, for $i \in [1:m]$ the problem

$$\min \left\{ -\mathbf{y}^\top \bar{Q}_i \mathbf{y} : \mathbf{y} \in \Delta^l \right\} \quad (29)$$

is a regularized maximum-clique formulation, where each maximal clique corresponds to a unique strict local maximizer with support equal to its vertices, and conversely (see [10] and references therein). The probability p is set as follows

$$p = \binom{l}{s}^{\frac{2}{s(s-1)}}, \quad (30)$$

for s the nearest integer to $0.4l$, so that the expected number of cliques with size $\approx 0.4l$ is 1, (see, e.g., [2]). Notice that the perturbation term \tilde{Q} ensures that problem (26) cannot be solved by optimizing each block separately.

We remark here that different ways to build large StQPs starting from smaller instances and preserving the structure of their solutions have been discussed in [17]. However, while the resulting problems decouple on the feasible set of the larger problem, they still decouple on the product of the feasible sets of the smaller instances, and for our purposes are equivalent to the block diagonal structure.

We tested four methods in total: AFW + SSC with parallel, GS and randomized updates (PAFW + SSC, GSAFW + SSC, BCAFW + SSC respectively), and FW with randomized updates (BCFW, coinciding with the block coordinate FW introduced in [28]). Our tests focused on the local identification and on the convergence properties of our methods.

The code was written in Python using the numpy package, and the tests were performed on an Intel Core i9-12900KS CPU 3.40GHz, 32GB RAM. The codes relevant to the numerical tests are available at the following link: <https://github.com/DamianoZeffiro/Projection-free-product-domain>.

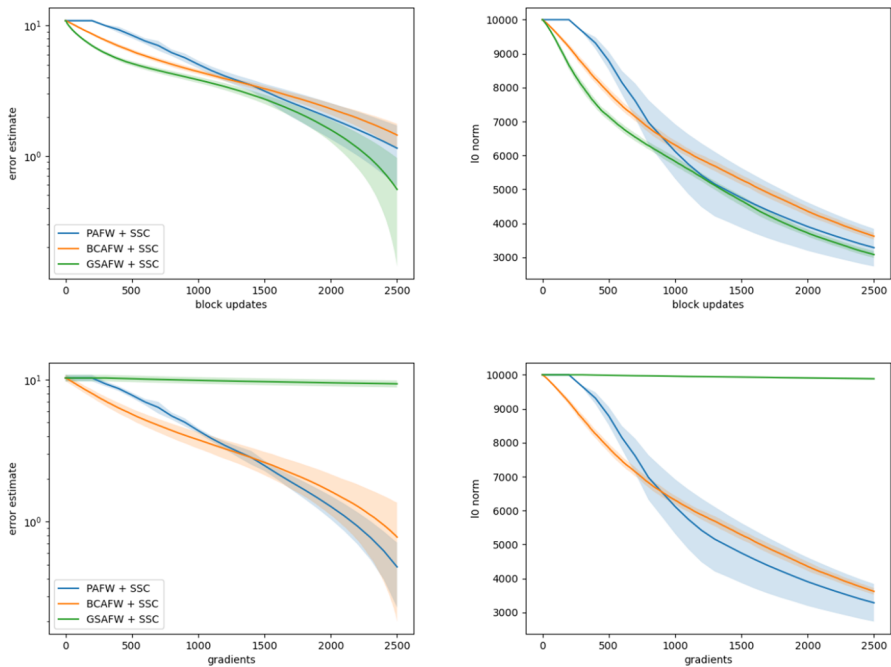


Fig. 1 Comparison using multistart between GSAFW + SSC, PAFW + SSC and BCAFW + SSC. $l = m = 100$

5.1 Multistart

We first considered a multistart approach, where the results are averaged across 20 runs, choosing 4 starting points for each of 5 random initializations of the objective.

We measure both optimality gap (error estimate) and sparsity (number of nonzero components, ℓ_0 norm) of the iterates, reporting average and standard deviation in the plots. The estimated global optimum used in the optimality gap is obtained by subtracting 10^{-5} from the best local solution found by the algorithms. We mostly consider the performance with respect to block gradient computations, with one gradient counted each time the SSC is performed in one of the blocks, as in previous works (see, e.g., [28]). In some tests involving the GSAFW + SSC, we consider instead block updates, with one block update counted each time the algorithms modifies the current iterate in one of the blocks. It is important to highlight that, since at each block update the gradient is constant and only one linear minimization is required at the beginning of the SSC, the number of gradient computations for our algorithms also coincides with the number of linear minimizations on the blocks for the FW variants we consider.

We first compare PAFW + SSC, BCAFW + SSC and GSAFW + SSC (Fig. 1). As expected, while GSAFW + SSC shows good performance with respect to block updates, it has a very poor performance with respect to block gradient computations, since at every iteration m gradients must be computed to update a single block. We then compared PAFW + SSC, BCAFW + SSC and BCFW. The results (Fig. 2) clearly

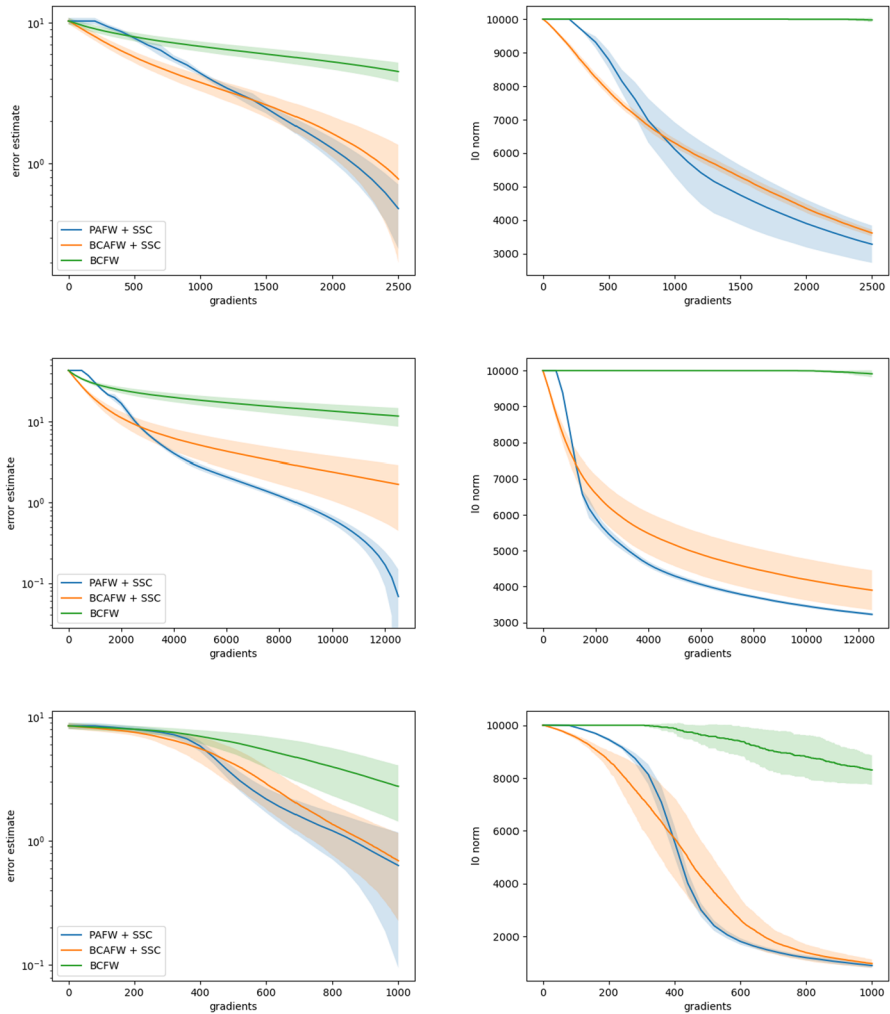


Fig. 2 Comparison using multistart between BCFW, PAFW + SSC and BCAFW + SSC. $l = m = 100$ in the first row, $l = 40$ and $m = 250$ in the second row, $l = 250$ and $m = 40$ in the third row

show that PAFW + SSC and BCAFW + SSC outperform BCFW. All these findings are consistent with the theoretical results described in “An active set identification criterion” section.

5.2 Monotonic basin hopping

We then consider the monotonic basin hopping approach (see, e.g., [30, 33]) described in Algorithm 3. The method computes a local optimizer $\mathbf{x}_{*,i}$ close to the current iterate $\bar{\mathbf{x}}_i$ (Step 2). There \mathcal{M} is a local optimization algorithm, and given as input \mathcal{M} and $\bar{\mathbf{x}}_i$, the subroutine LO returns the result of applying \mathcal{M} starting from $\bar{\mathbf{x}}_i$, with a suitable




Fig. 3 Comparison using Monotonic Basin Hopping with BCFW, PAFW + SSC and BCFW + SSC. From left to right: $l = m = 100$; $l = 40$ and $m = 250$; $l = 250$ and $m = 40$

stopping criterion which in our case is given by a limit on the number of gradient computations, set to $10m$. The sequence of best points found in the first i iterations $\{\bar{\mathbf{x}}_{*,i}\}$ is updated in Step 3, and in Step 5, $\bar{\mathbf{x}}_{i+1}$ is chosen in a neighborhood of $\bar{\mathbf{x}}_{*,i}$. The neighborhood $B(\mathbf{x}, \gamma)$ for $\mathbf{x} \in \mathcal{C}$ and $\gamma \in (0, 1]$ is given by

$$B(\mathbf{x}, \gamma) = \{\mathbf{x} + \gamma(\mathbf{y} - \mathbf{x}) : \mathbf{y} \in \mathcal{C}\}. \tag{31}$$

In the tests, we chose \mathbf{y} uniformly at random in \mathcal{C} and set $\bar{\mathbf{x}}_{i+1} = \bar{\mathbf{x}}_i + \gamma(\mathbf{y} - \bar{\mathbf{x}}_i)$,

Algorithm 3 Monotonic Basin Hopping Strategy

- 1: $\bar{\mathbf{x}}_0 \in \mathcal{C}, i_{\max} \in \mathbb{N}, \gamma \in [0, 1], i = 0$. Set $\bar{\mathbf{x}}_{*, -1} = \bar{\mathbf{x}}_0$
 - 2: Compute a local optimizer $\mathbf{x}_{*, i} = \text{LO}(\mathcal{M}, \bar{\mathbf{x}}_i)$
 - 3: **If** $f(\mathbf{x}_{*, i}) < f(\bar{\mathbf{x}}_{*, i-1})$ **then** set $\bar{\mathbf{x}}_{*, i} = \mathbf{x}_{*, i}$, **else** set $\bar{\mathbf{x}}_{*, i} = \bar{\mathbf{x}}_{*, i-1}$.
 - 4: **If** $i = i_{\max}$, **then** STOP.
 - 5: Randomly chose $\bar{\mathbf{x}}_{i+1}$ in $B(\bar{\mathbf{x}}_{*, i}, \gamma)$.
 - 6: Set $i = i + 1$. Go to step 2.
-

with $\gamma = 0.25$. The methods we consider as subroutines in Step 2 are PAFW + SSC, BCFW + SSC and BCFW. We set $i_{\max} = 9$, and perform 10 runs of Algorithm 3, randomly initializing the starting point. We plot once again (Fig. 3) average and standard deviation for $\{f(\bar{\mathbf{x}}_{*, i}) - \tilde{f}^*\}$ with \tilde{f}^* estimating the global optimum (obtained by subtracting 10^{-1} from the best solution found by the methods).

The results again show that PAFW + SSC and BCFW + SSC find better solutions than BCFW, with BCFW + SSC outperforming PAFW + SSC in most instances if $l \leq m$.

6 Conclusions

For a quite general optimization problem on product domains, we offer a seemingly new convergence theory, which ensures both convergence of objective values and (local) linear convergence of the iterates under widely accepted conditions, for block-coordinate FW variants. Convergence is global for μ -strongly convex objectives, but we mainly focus on the non-convex case. In case of randomized selection of the blocks, all results are in expectation, and need a further technical assumption. As

usual, constants and rates are specified in terms of the Lipschitz constant L for the gradient map, the constant μ used in the local Kurdyka-Łojasiewicz-condition, and the parameter τ in the so-called angle condition.

The results are complemented by an active set identification result for a specific structure of the product domain and suitable choices of a projection-free strategy (FW-approach with away steps for the search direction): it is proved that our framework identifies the support of a solution in a finite number of iterations.

To the best of our knowledge, this is the first time that both a linear convergence rate and an active set identification result are given for (bad step-free) block-coordinate FW variants, in an effort to narrow the research gap observed in [36].

In our preliminary experiments, numerical evidence clearly points out the advantages of our strategy to exploit structural knowledge. On randomly generated non-convex Multi-StQPs where easy instances were carefully avoided, our approach (AFW with parallel or randomized updates, both combined with the Short Step Chain strategy SSC) is dominating the block-coordinate FW method with randomized updates.

We tested resilience of our reported observations by employing two experimental setups, pure multistart and monotonic basin hopping. The same effects seem to prevail.

Instance construction was motivated by a stochastic variant of the StQP, varying both domain dimension l and number m of possible scenarios. In case $l \leq m$ there seems to be a slight edge towards the combination of AFW with randomized updates and SSC, compared to the parallel variant. This effect does not seem to happen with large l in comparison to m , but would not change superiority over traditional block-coordinate FW methods.

7. Appendix

7.1 Proofs

In the rest of this section, we always assume that the SSC terminates in a finite number of steps and that the angle condition holds.

The first lemma is related to the single block setting, and strengthens some of the properties proved for the SSC in [39, Proposition 4.1].

Lemma 1 *For a fixed $i \in [1:m]$, let $\{\mathbf{w}_k\} = \{\mathbf{x}_k^{(i)}\}$, and let $\mathbf{w}_{k+1} = \text{SSC}(\mathbf{w}_k, \mathbf{g})$. Then there exists $\tilde{\mathbf{w}}_k \in \{\mathbf{y}_j\}_{j=0}^T$ such that*

$$\|\mathbf{w}_{k+1} - \mathbf{w}_k\| \geq \frac{\tau}{L} \|\pi(T_{C_{(i)}}(\tilde{\mathbf{w}}_k), \mathbf{g})\| \quad (32)$$

and

$$\|\tilde{\mathbf{w}}_k - \mathbf{w}_k\| \leq \|\mathbf{w}_{k+1} - \mathbf{w}_k\|, \quad (33)$$

$$\langle \mathbf{g}, \tilde{\mathbf{w}}_k - \mathbf{w}_k \rangle \leq \langle \mathbf{g}, \mathbf{w}_{k+1} - \mathbf{w}_k \rangle. \quad (34)$$

Furthermore, we have

$$L\|y - w_k\|^2 \leq \langle g, y - w_k \rangle \tag{35}$$

for $y \in \{w_{k+1}, \tilde{w}_k\}$.

Proof Let $\bar{B} = \bar{B}_{\frac{\|g\|}{2L}}(w_k + \frac{g}{2L})$ and let T be such that $w_{k+1} = y_T$. By [39, (4.4)] we have that (35) holds for every $z \in \bar{B}$ (in place of y), and therefore as desired for every

$$y \in \{w_{k+1}, \tilde{w}_k\} \subset \{y_j : j \in [0:T]\} \subset \bar{B}.$$

Let now $\tilde{p}_j = \|\pi(T_{C_{(i)}}(y_j), g)\|$. Notice that, if $\tilde{w}_k = y_l$, then

$$\frac{\tau}{L} \|\pi(T_{C_{(i)}}(\tilde{w}_k), g)\| = \frac{\tau}{L} \tilde{p}_l \leq \frac{1}{L} \langle g, \hat{d}_l \rangle, \tag{36}$$

where the inequality follows from $\frac{\langle g, \hat{d}_l \rangle}{\tilde{p}_l} \geq \text{DSB}_{\mathcal{A}}(C_{(i)}, y_l, g) \geq \text{SB}_{\mathcal{A}}(C_{(i)}) = \tau$. Thus for proving (32), in the rest of the proof it will be enough to prove

$$\langle g, \hat{d}_l \rangle \leq L\|w_{k+1} - w_k\|. \tag{37}$$

Furthermore, since by definition of the SSC, the scalar product $\langle g, y_j \rangle$ is increasing in j , we have

$$\langle g, \tilde{w}_k - w_k \rangle = \langle g, y_l - w_k \rangle \leq \langle g, y_T - w_k \rangle = \langle g, w_{k+1} - w_k \rangle.$$

We distinguish four cases, according to how the SSC terminates. In the first two, we show we can choose the last step, $\tilde{w} = y_T$; in the third, the penultimate choice $\tilde{w} = y_{T-1}$ satisfies all conditions, and in the fourth case, an intermediate step is an appropriate choice. We abbreviate $B_j = \bar{B}_{\langle g, \hat{d}_j \rangle / L}(w_k)$.

Case 1: $T = 0$ or $d_T = o$. Since there are no descent directions, $w_{k+1} = y_T$ must be stationary for the gradient $-g$. Equivalently, $\tilde{p}_T = \|\pi(T_{C_{(i)}}(w_{k+1}), g)\| = 0$. Finally, it is clear that if $T = 0$ then $d_0 = o$, since y_0 must be stationary for $-g$. Thus taking $\tilde{w}_k = y_T$ the desired properties follow.

Before examining the remaining cases we remark that if the SSC terminates in Phase II, then $\alpha_{T-1} = \beta_{T-1}$ must be maximal w.r.t. the conditions $y_T \in B_{T-1}$ or $y_T \in \bar{B}$. If $\alpha_{T-1} = 0$ then $y_{T-1} = y_T$, and in this case we cannot have $y_{T-1} \in \partial \bar{B}$, otherwise the SSC would terminate in Phase II of the previous cycle. Therefore necessarily $y_T = y_{T-1} \in \text{int}(B_{T-1})^c$ (Case 2). If $\beta_{T-1} = \alpha_{T-1} > 0$ we must have $y_{T-1} \in C_{T-1} = B_{T-1} \cap \bar{B}$, and $y_T \in \partial B_{T-1}$ (Case 3) or $y_T \in \partial \bar{B}$ (Case 4) respectively.

Case 2: $y_{T-1} = y_T \in \text{int}(B_{T-1})^c$, which means $\|y_{T-1} - w_k\| \geq \frac{\langle g, \hat{d}_{T-1} \rangle}{L}$. We can rewrite the condition as

$$\langle g, \hat{d}_{T-1} \rangle \leq L\|y_{T-1} - w_k\| = L\|y_T - w_k\|, \tag{38}$$

which is exactly (37). Then $\tilde{\mathbf{w}}_k = \mathbf{w}_{k+1} = \mathbf{y}_T$ satisfies the desired conditions.

Case 3: $\mathbf{y}_T = \mathbf{y}_{T-1} + \beta_{T-1} \mathbf{d}_{T-1}$ and $\mathbf{y}_T \in \partial B_{T-1}$. Then from $\mathbf{y}_{T-1} \in B_{T-1}$ it follows

$$L \|\mathbf{y}_{T-1} - \mathbf{w}_k\| \leq \langle \mathbf{g}, \hat{\mathbf{d}}_{T-1} \rangle, \tag{39}$$

and $\mathbf{y}_T \in \partial B_{T-1}$ implies

$$\langle \mathbf{g}, \hat{\mathbf{d}}_{T-1} \rangle = L \|\mathbf{y}_T - \mathbf{w}_k\|, \tag{40}$$

which is (37) for $l = T - 1$. Combining (39) with (40) we also obtain

$$L \|\mathbf{y}_{T-1} - \mathbf{w}_k\| \leq L \|\mathbf{y}_T - \mathbf{w}_k\|, \tag{41}$$

so that in particular we can take $\tilde{\mathbf{w}}_k = \mathbf{y}_{T-1}$.

Case 4: $\mathbf{y}_T = \mathbf{y}_{T-1} + \beta_{T-1} \mathbf{d}_{T-1}$ and $\mathbf{y}_T \in \partial \bar{B}$.

The condition $\mathbf{w}_{k+1} = \mathbf{y}_T \in \partial \bar{B}$ can be rewritten as

$$L \|\mathbf{w}_{k+1} - \mathbf{w}_k\|^2 = \langle \mathbf{g}, \mathbf{w}_{k+1} - \mathbf{w}_k \rangle. \tag{42}$$

Indeed, for any $\{\mathbf{a}, \mathbf{b}\} \subset \mathbb{R}^n$, the equation $\|\mathbf{b} - \mathbf{a}\| = \|\mathbf{b}\|$ implies $\|\mathbf{a}\|^2 = 2\langle \mathbf{a}, \mathbf{b} \rangle$, wherefrom (42) follows, putting $\mathbf{a} = \mathbf{w}_{k+1} - \mathbf{w}_k$ and $\mathbf{b} = \frac{\mathbf{g}}{2L}$. For every $j \in [0 : T]$ we have

$$\mathbf{w}_{k+1} = \mathbf{y}_j + \sum_{i=j}^{T-1} \alpha_i \mathbf{d}_i. \tag{43}$$

We now want to prove that for every $j \in [0 : T]$

$$\|\mathbf{w}_{k+1} - \mathbf{w}_k\| \geq \|\mathbf{y}_j - \mathbf{w}_k\|. \tag{44}$$

Indeed, we have

$$\begin{aligned} L \|\mathbf{w}_{k+1} - \mathbf{w}_k\|^2 &= \langle \mathbf{g}, \mathbf{w}_{k+1} - \mathbf{w}_k \rangle = \langle \mathbf{g}, \mathbf{y}_j - \mathbf{w}_k \rangle + \sum_{i=j}^{T-1} \alpha_i \langle \mathbf{g}, \mathbf{d}_i \rangle \\ &\geq \langle \mathbf{g}, \mathbf{y}_j - \mathbf{w}_k \rangle \geq L \|\mathbf{y}_j - \mathbf{w}_k\|^2, \end{aligned}$$

where we used (42) in the first equality, (43) in the second, $\langle \mathbf{g}, \mathbf{d}_j \rangle \geq 0$ for every j in the first inequality and $\mathbf{y}_j \in \bar{B}$ in the second inequality (using an argument similar to that for (42)), which proves (44). We also have

$$\langle \mathbf{g}, \sum_{j=0}^{T-1} \alpha_j \mathbf{d}_j \rangle = \sum_{j=0}^{T-1} \alpha_j \|\mathbf{d}_j\| \frac{\langle \mathbf{g}, \mathbf{d}_j \rangle}{\|\mathbf{d}_j\|} \geq \min \left\{ \frac{\langle \mathbf{g}, \mathbf{d}_i \rangle}{\|\mathbf{d}_i\|} : i \in [0 : T - 1] \right\} \sum_{j=0}^{T-1} \alpha_j \|\mathbf{d}_j\|,$$

which implies

$$\begin{aligned} \frac{\langle \mathbf{g}, \mathbf{w}_{k+1} - \mathbf{w}_k \rangle}{\|\mathbf{w}_{k+1} - \mathbf{w}_k\|} &= \frac{\langle \mathbf{g}, \sum_{j=0}^{T-1} \alpha_j \mathbf{d}_j \rangle}{\|\sum_{j=0}^{T-1} \alpha_j \mathbf{d}_j\|} \geq \frac{\langle \mathbf{g}, \sum_{j=0}^{T-1} \alpha_j \mathbf{d}_j \rangle}{\sum_{j=0}^{T-1} \alpha_j \|\mathbf{d}_j\|} \\ &\geq \min \left\{ \frac{\langle \mathbf{g}, \mathbf{d}_i \rangle}{\|\mathbf{d}_i\|} : i \in [0:T-1] \right\}, \end{aligned} \tag{45}$$

using the triangle inequality for the left inequality. Thus for $\tilde{T} \in \operatorname{argmin} \left\{ \frac{\langle \mathbf{g}, \mathbf{d}_i \rangle}{\|\mathbf{d}_i\|} : i \in [0:T-1] \right\}$ we have

$$\langle \mathbf{g}, \hat{\mathbf{d}}_{\tilde{T}} \rangle \leq \frac{\langle \mathbf{g}, \mathbf{w}_{k+1} - \mathbf{w}_k \rangle}{\|\mathbf{w}_{k+1} - \mathbf{w}_k\|} = L \|\mathbf{w}_{k+1} - \mathbf{w}_k\|, \tag{46}$$

where we used (45) in the first inequality and (42) in the second (equality).

In particular $\tilde{\mathbf{w}}_k = \mathbf{y}_{\tilde{T}}$ satisfies the desired properties, where $\|\tilde{\mathbf{w}}_k - \mathbf{w}_k\| \leq \|\mathbf{w}_{k+1} - \mathbf{w}_k\|$ by (44) and (37) holds by (46). \square

We denote by $\overline{\text{SSC}}(\mathbf{w}_k, \mathbf{g})$ a point $\tilde{\mathbf{w}}_k$ with the properties stated in the above lemma. It is also useful to define U_0 as the connected component of $\{\mathbf{x} \in \mathcal{C} : f(\mathbf{x}) \leq f(\mathbf{x}_0)\}$ containing \mathbf{x}_0 . The next result shows how in our block coordinate setting the assumption of Theorem 1 on U_0 allows us to retrieve a lower bound on the objective for points generated by the SSC. This lower bound is analogous to the lower bound required in [39, Theorem 4.2].

Lemma 2 *Let $\{\mathbf{x}_k\}$ be a sequence generated by Algorithm 1. Let $\bar{\mathbf{x}}_k = [\overline{\text{SSC}}(\mathbf{x}_k^{(i)}, -\nabla f(\mathbf{x}_k^{(i)}))]_{i=1}^m$. Assume also that $f(\mathbf{x}_*)$ is a minimum in U_0 . Then, $\{f(\mathbf{x}_k)\}$ is decreasing, and for every k , $\{\mathbf{x}_k, \bar{\mathbf{x}}_k\} \subset U_0$, with $f(\mathbf{y}) \in [f(\mathbf{x}_*), f(\mathbf{x}_0)]$ for $\mathbf{y} \in \{\mathbf{x}_k, \bar{\mathbf{x}}_k\}$.*

Proof Let U_k be the minimal connected component of $\{\mathbf{x} \in \mathcal{C} : f(\mathbf{x}) \leq f(\mathbf{x}_k)\}$ containing \mathbf{x}_k , let $\mathbf{g} = -\nabla f(\mathbf{x}_k)$ and let $\bar{B}_k^{\mathcal{C}} = \mathcal{C} \cap \prod_i \bar{B}_{\frac{\|\mathbf{g}^{(i)}\|}{2L}}(\mathbf{x}_k^{(i)} + \frac{\mathbf{g}^{(i)}}{2L})$. For $\mathbf{y} \in \bar{B}_k^{\mathcal{C}}$, we have $\mathbf{x}_k \in \bar{B}_k^{\mathcal{C}}$ and

$$\begin{aligned} f(\mathbf{y}) &\leq f(\mathbf{x}_k) + \langle \nabla f(\mathbf{x}_k), \mathbf{y} - \mathbf{x}_k \rangle + \frac{L}{2} \|\mathbf{y} - \mathbf{x}_k\|^2 \\ &= f(\mathbf{x}_k) + \sum_{i=1}^m \langle \nabla f(\mathbf{x}_k)^{(i)}, \mathbf{y}^{(i)} - \mathbf{x}_k^{(i)} \rangle + \frac{L}{2} \|\mathbf{y}^{(i)} - \mathbf{x}_k^{(i)}\|^2 \\ &\leq f(\mathbf{x}_k) - \frac{L}{2} \sum_{i=1}^m \|\mathbf{y}^{(i)} - \mathbf{x}_k^{(i)}\|^2 = f(\mathbf{x}_k) - \frac{L}{2} \|\mathbf{x}_k - \mathbf{y}\|^2 \end{aligned} \tag{47}$$

where we used the standard Descent Lemma in the first inequality and the the second follows by definition of $\bar{B}_k^{\mathcal{C}}$. From (47) it follows that $\{f(\mathbf{x}_k)\}$ is decreasing, and that $\bar{B}_k^{\mathcal{C}} \subset \{\mathbf{x} \in \mathcal{C} : f(\mathbf{x}) \leq f(\mathbf{x}_k)\}$. Furthermore, since $\bar{B}_k^{\mathcal{C}}$ is connected and contains \mathbf{x}_k ,

the stronger inclusion $\bar{B}_k^C \subset U_k$ is also true. Thus $\{\mathbf{x}_{k+1}, \bar{\mathbf{x}}_k\} \subset \bar{B}_k^C \subset U_k$, so that in particular $U_{k+1} \subset U_k$ since $f(\mathbf{x}_{k+1}) \leq f(\bar{\mathbf{x}}_k)$, and by induction we can conclude $\{\mathbf{x}_{k+1}, \bar{\mathbf{x}}_k\} \subset U_0$. Finally, $f(\mathbf{y}) \in [f(\mathbf{x}_*), f(\bar{\mathbf{x}}_k)]$ for $\mathbf{y} \in \{\mathbf{x}_{k+1}, \bar{\mathbf{x}}_k\}$, where the lower bound follows from the assumption that $f(\mathbf{x}_*)$ is a minimum in U_0 , and the upper bound follows from (47). \square

In the following lemma, the properties of the SSC proved in Lemma 1 for single blocks are combined to obtain analogous properties on the whole product of blocks, and the KL condition is then used to lower bound suitable improvement measures with an optimality gap for the objective. We would like to highlight that, unlike the single block case, this optimality gap is measured with respect to an auxiliary point which is not necessarily among those generated by the algorithm. Proof of the linear convergence rate hence requires proper handling in this case.

Lemma 3 *Let $\{\mathbf{x}_k\}$ be a sequence generated by Algorithm 1, and assume that the angle condition holds for the method $\mathcal{A}^{(i)}$ with the same τ , for all $i \in [1 : m]$. Let $\bar{\mathbf{x}}_k = [\bar{\text{SSC}}(\mathbf{x}_k^{(i)}, -\nabla f(\mathbf{x}_k)^{(i)})]_{i=1}^m$ and $\tilde{\mathbf{x}}_{k+1} = [\text{SSC}(\mathbf{x}_k^{(i)}, -\nabla f(\mathbf{x}_k)^{(i)})]_{i=1}^m$. If (14) holds at $\bar{\mathbf{x}}_k$, we then have, abbreviating $\mathbf{g} = -\nabla f(\mathbf{x}_k)$:*

$$\|\tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k\|^2 \geq \frac{\tau^2}{2(1 + \tau^2)L^2} \|\pi(T_C(\bar{\mathbf{x}}_k), -\nabla f(\bar{\mathbf{x}}_k))\|^2 \geq \frac{\tau^2 \mu}{L^2(1 + \tau^2)} [f(\bar{\mathbf{x}}_k) - f^*], \quad (48)$$

$$\frac{1}{2} \langle \mathbf{g}, \tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k \rangle \geq \frac{1}{3} [f(\mathbf{x}_k) - f(\bar{\mathbf{x}}_k)]. \quad (49)$$

Proof Let $\bar{\mathbf{g}} = -\nabla f(\bar{\mathbf{x}}_k)$, $\bar{q}_{(i)} = \|\pi(T_{C_{(i)}}(\bar{\mathbf{x}}_k^{(i)}), \bar{\mathbf{g}}^{(i)})\|$, and $q_{(i)} = \|\pi(T_{C_{(i)}}(\bar{\mathbf{x}}_k^{(i)}), \mathbf{g}^{(i)})\|$. Observe that by the Lipschitz continuity of the gradient, we have the inequality

$$\bar{q}_{(i)} \leq q_{(i)} + L \|\bar{\mathbf{x}}_k^{(i)} - \mathbf{x}_k^{(i)}\| \quad (50)$$

and thus

$$\bar{q}_{(i)}^2 \leq 2q_{(i)}^2 + 2L^2 \|\bar{\mathbf{x}}_k^{(i)} - \mathbf{x}_k^{(i)}\|^2 \leq \frac{2L^2(1 + \tau^2)}{\tau^2} \|\bar{\mathbf{x}}_{k+1}^{(i)} - \mathbf{x}_k^{(i)}\|^2, \quad (51)$$

where we applied Jensen's inequality to (50) in the first inequality, and (32) together with (33) in the second inequality.

Thus we can write

$$\begin{aligned} \|\tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k\|^2 &= \sum_{i=1}^m \|\tilde{\mathbf{x}}_{k+1}^{(i)} - \mathbf{x}_k^{(i)}\|^2 \geq \frac{\tau^2}{2L^2(1 + \tau^2)} \sum_{i=1}^m \bar{q}_{(i)}^2 \\ &= \frac{\tau^2}{2L^2(1 + \tau^2)} \|\pi(T_C(\bar{\mathbf{x}}_k), -\nabla f(\bar{\mathbf{x}}_k))\|^2 \geq \frac{\tau^2 \mu}{L^2(1 + \tau^2)} [f(\bar{\mathbf{x}}_k) - f^*], \end{aligned} \quad (52)$$

where we used (51) in the first inequality and the KL property in the second. This proves (48).

Using the standard Descent Lemma, we can give the upper bound

$$\begin{aligned}
 f(\mathbf{x}_k) - f(\bar{\mathbf{x}}_k) &\leq \langle \mathbf{g}, \bar{\mathbf{x}}_k - \mathbf{x}_k \rangle + \frac{L}{2} \|\bar{\mathbf{x}}_k - \mathbf{x}_k\|^2 \leq \frac{3}{2} \langle \mathbf{g}, \bar{\mathbf{x}}_k - \mathbf{x}_k \rangle \\
 &= \frac{3}{2} \sum_{i=1}^m \langle \mathbf{g}^{(i)}, \bar{\mathbf{x}}_k^{(i)} - \mathbf{x}_k^{(i)} \rangle,
 \end{aligned}
 \tag{53}$$

where we used (35) in the second inequality. We can finally prove (49):

$$\frac{1}{2} \langle \mathbf{g}, \tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k \rangle \geq \frac{1}{2} \langle \mathbf{g}, \bar{\mathbf{x}}_k - \mathbf{x}_k \rangle = \frac{1}{2} \sum_{i=1}^m \langle \mathbf{g}^{(i)}, \bar{\mathbf{x}}_k^{(i)} - \mathbf{x}_k^{(i)} \rangle \geq \frac{1}{3} [f(\mathbf{x}_k) - f(\bar{\mathbf{x}}_k)],
 \tag{54}$$

where we used (34) in the first inequality and (53) in the second one. □

The next result, which directly follows from the previous lemma, explicitly lower bounds the improvement on the objective with the optimality gap introduced above.

Lemma 4 *Let $\{\mathbf{x}_k\}$ be a sequence generated by Algorithm 1, and assume that the angle condition holds for the method $\mathcal{A}^{(i)}$ with the same τ , for all $i \in [1 : m]$. Let $\bar{\mathbf{x}}_k = (\overline{\text{SSC}}(\mathbf{x}_k^{(i)}, -\nabla f(\mathbf{x}_k)^{(i)}))_{i=1}^m$. Then, if the KL property (14) holds at $\bar{\mathbf{x}}_k$, for parallel updates*

$$f(\mathbf{x}_k) - f(\mathbf{x}_{k+1}) \geq \frac{\tau^2 \mu}{2L(1 + \tau^2)} (f(\bar{\mathbf{x}}_k) - f^*),
 \tag{55}$$

for GS updates

$$f(\mathbf{x}_k) - f(\mathbf{x}_{k+1}) \geq \frac{\tau^2 \mu}{2L(1 + \tau^2)} \frac{1}{m} (f(\bar{\mathbf{x}}_k) - f^*),
 \tag{56}$$

and for random updates

$$\mathbb{E}[f(\mathbf{x}_k) - f(\mathbf{x}_{k+1})] \geq \frac{\tau^2 \mu}{2L(1 + \tau^2)} \frac{1}{m} \mathbb{E}[f(\bar{\mathbf{x}}_k) - f^*].
 \tag{57}$$

Proof We first prove the inequality for parallel updates. We have

$$\begin{aligned}
 f(\mathbf{x}_k) - f(\mathbf{x}_{k+1}) &\geq \frac{L}{2} \|\mathbf{x}_k - \mathbf{x}_{k+1}\|^2 \geq \frac{L}{2} \frac{\tau^2}{2L^2(1 + \tau^2)} \|\pi(T_{\mathcal{C}}(\bar{\mathbf{x}}_k), -\nabla f(\bar{\mathbf{x}}_k))\|^2 \\
 &= \frac{\tau^2}{4L(1 + \tau^2)} \|\pi(T_{\mathcal{C}}(\bar{\mathbf{x}}_k), -\nabla f(\bar{\mathbf{x}}_k))\|^2 \geq \frac{\tau^2 \mu}{2L(1 + \tau^2)} [f(\bar{\mathbf{x}}_k) - f^*],
 \end{aligned}
 \tag{58}$$

where the first inequality follows from (47), the second inequality by (48) where with the notation introduced in Lemma 3 we have by definition $\mathbf{x}_{k+1} = \tilde{\mathbf{x}}_{k+1}$. For GS updates, we have

$$\begin{aligned} f(\mathbf{x}_k) - f(\mathbf{x}_{k+1}) &\geq \langle \mathbf{g}, \mathbf{x}_{k+1} - \mathbf{x}_k \rangle - \frac{L}{2} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \geq \frac{1}{2} \langle \mathbf{g}, \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \\ &= \frac{1}{2} \max_{i \in [1:m]} \langle \mathbf{g}^{(i)}, \tilde{\mathbf{x}}_{k+1}^{(i)} - \mathbf{x}_k \rangle \geq \frac{1}{2m} \langle \mathbf{g}, \tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k \rangle \geq \frac{L}{2m} \|\tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k\|^2 \quad (59) \\ &\geq \frac{\tau^2 \mu}{2mL(1 + \tau^2)} [f(\bar{\mathbf{x}}_k) - f^*], \end{aligned}$$

where in the first inequality we used the standard Descent Lemma, (35) in the second inequality; the equality follows by definition of GS updates, in the fourth inequality we applied again (35), and (48) in the last one.

Finally, for random updates we have, denoting as $i(k) = j$ the event that the index chosen at the step k is j :

$$\begin{aligned} \mathbb{E}[f(\mathbf{x}_k) - f(\mathbf{x}_{k+1})] &\geq \frac{L}{2} \mathbb{E}[\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2] = \frac{L}{2} \sum_{j=1}^m \mathbb{P}(i(k) = j) \mathbb{E}[\|\tilde{\mathbf{x}}_{k+1}^{(j)} - \mathbf{x}_k^{(j)}\|^2] \quad (60) \\ &= \frac{L}{2m} \sum_{j=1}^m \mathbb{E}[\|\tilde{\mathbf{x}}_{k+1}^{(j)} - \mathbf{x}_k^{(j)}\|^2] = \frac{L}{2m} \mathbb{E}[\|\tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k\|^2] \geq \frac{\tau^2 \mu}{2L(1 + \tau^2)} \mathbb{E}[f(\bar{\mathbf{x}}_k) - f^*], \end{aligned}$$

where the first inequality follows from (47), we used $\mathbb{P}(\{i(k) = j\}) = \frac{1}{m}$ in the second equality and (48) in the last inequality. \square

In the next two lemmas, we relate the improvement measured with respect to the auxiliary point to the true improvement of the objective, and thus manage to extend the linear convergence rate in [39, Lemma 4.3] to the block coordinate setting.

Lemma 5 *Let $\{\mathbf{x}_k\}$ be a sequence generated by Algorithm 1, and assume that the angle condition holds for the method $\mathcal{A}^{(i)}$ with the same τ , for all $i \in [1 : m]$. Let $\bar{\mathbf{x}}_k = (\text{SSC}(\mathbf{x}_k^{(i)}, -\nabla f(\mathbf{x}_k)^{(i)))_{i=1}^m$. Then for parallel updates*

$$f(\mathbf{x}_k) - f(\mathbf{x}_{k+1}) \geq \frac{1}{3} [f(\mathbf{x}_k) - f(\bar{\mathbf{x}}_k)], \quad (61)$$

for GS updates

$$f(\mathbf{x}_k) - f(\mathbf{x}_{k+1}) \geq \frac{1}{3m} [f(\mathbf{x}_k) - f(\bar{\mathbf{x}}_k)], \quad (62)$$

and for random updates

$$\mathbb{E}[f(\mathbf{x}_k) - f(\mathbf{x}_{k+1})] \geq \frac{1}{3m} \mathbb{E}[f(\mathbf{x}_k) - f(\bar{\mathbf{x}}_k)] \quad (63)$$

Proof For parallel updates, we have

$$\begin{aligned} f(\mathbf{x}_k) - f(\mathbf{x}_{k+1}) &\geq \langle \mathbf{g}, \mathbf{x}_{k+1} - \mathbf{x}_k \rangle - \frac{L}{2} \|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \geq \frac{1}{2} \langle \mathbf{g}, \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \\ &= \frac{1}{2} \langle \mathbf{g}, \tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k \rangle \geq \frac{1}{3} [f(\mathbf{x}_k) - f(\tilde{\mathbf{x}}_k)], \end{aligned} \tag{64}$$

where we have used the standard descent Lemma in the first inequality, (35) in the second inequality, and (49) in the last inequality.

The proof follows analogously for GS updates, after noticing

$$\langle \mathbf{g}, \mathbf{x}_{k+1} - \mathbf{x}_k \rangle \geq \frac{1}{m} \langle \mathbf{g}, \tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k \rangle, \tag{65}$$

as showed in (59), and for random updates, using

$$\mathbb{E}[\langle \mathbf{g}, \mathbf{x}_{k+1} - \mathbf{x}_k \rangle] = \frac{1}{m} \mathbb{E}[\langle \mathbf{g}, \tilde{\mathbf{x}}_{k+1} - \mathbf{x}_k \rangle], \tag{66}$$

respectively. □

Lemma 6 *Let $\{\mathbf{x}_k\}$ be a sequence generated by Algorithm 1, and assume that the angle condition holds for the method $\mathcal{A}^{(i)}$ with the same τ , for all $i \in [1 : m]$. Then, if the KL property (14) holds at \mathbf{x}_k , for parallel updates*

$$f(\mathbf{x}_{k+1}) - f^* \leq q_P (f(\mathbf{x}_k) - f^*), \tag{67}$$

for GS updates

$$f(\mathbf{x}_{k+1}) - f^* \leq q_{GS} (f(\mathbf{x}_k) - f^*), \tag{68}$$

and for random updates

$$\mathbb{E}[f(\mathbf{x}_{k+1}) - f^*] \leq q_R \mathbb{E}[f(\mathbf{x}_k) - f^*]. \tag{69}$$

Proof First observe that since $\tau \in [0, 1]$ and $\mu \leq L$ we have

$$\frac{\tau^2 \mu}{2L(1 + \tau^2)} \leq \frac{\mu}{3L} \leq \frac{1}{3}. \tag{70}$$

Then, combining (55) and (61), we can write

$$\begin{aligned} 2[f(\mathbf{x}_k) - f(\mathbf{x}_{k+1})] &\geq \frac{1}{3} [f(\mathbf{x}_k) - f(\tilde{\mathbf{x}}_k)] + \frac{\tau^2 \mu}{2L(1 + \tau^2)} [f(\tilde{\mathbf{x}}_k) - f^*] \\ &\geq \frac{\tau^2 \mu}{2L(1 + \tau^2)} [f(\mathbf{x}_k) - f(\tilde{\mathbf{x}}_k)] + \frac{\tau^2 \mu}{2L(1 + \tau^2)} [f(\tilde{\mathbf{x}}_k) - f^*] \\ &= \frac{\tau^2 \mu}{2L(1 + \tau^2)} [f(\mathbf{x}_k) - f^*], \end{aligned} \tag{71}$$

and rearranging an inequality of the form $2(a - b) \geq \gamma(a - f^*)$ into the form $b - f^* \leq (1 - \frac{\gamma}{2})(a - f^*)$, using (17),

$$f(\mathbf{x}_{k+1}) - f^* \leq q_P [f(\mathbf{x}_k) - f^*]. \tag{72}$$

The thesis follows for GS and random updates analogously, invoking (20). \square

Given the previous results for the block coordinate setting, the remaining part of the proof is a straightforward adaptation of arguments used in the proof of [39, Theorem 4.2].

Proof of Theorem 1 We need to prove that the KL property (14) holds in $\{\bar{\mathbf{x}}_k\}$. The bounds on $f(\mathbf{x}_k) - f(\mathbf{x}_*)$ then follow immediately by induction from Lemma 6, and in turn the bounds on $\|\mathbf{x}_k - \mathbf{x}_*\|$ follow as in the proof of [39, Lemma 4.3].

For random updates, we can take $\tilde{\delta} < \delta$ small enough so that $f(\mathbf{x}_0) < f(\mathbf{x}_*) + \eta$. Then by construction the KL property (14) holds in U_0 , and since $\{\bar{\mathbf{x}}_k\}$ is contained in U_0 by Lemma 2, (14) holds in particular in $\{\bar{\mathbf{x}}_k\}$.

For parallel updates, thanks to Lemma 2 we have that $\{f(\mathbf{x}_k)\}$ is decreasing and $f(\bar{\mathbf{x}}_k), f(\mathbf{x}_k) \geq f(\mathbf{x}_*)$. It can then be proved with an argument analogous to the proof of [39, Theorem 4.2] that for $\tilde{\delta}$ small enough, (14) holds in $\{\bar{\mathbf{x}}_k\}$. We include the argument here for completeness. Let $f_k = f(\mathbf{x}_k) - f(\mathbf{x}_*)$, and let $\tilde{\delta} < \delta/2$ defined as in the proof of [39, Theorem 4.2] so that

$$\tilde{\delta} < \frac{\delta}{2} < \delta - \frac{\sqrt{2f_0(1-q)}}{L(1-\sqrt{q})} - \sqrt{\frac{2}{L}}\sqrt{f_0}, \tag{73}$$

with $q = q_P$ here. We now want to prove $\bigcup_{\{0:k-1\}}\{\mathbf{x}_i, \bar{\mathbf{x}}_i\} \cup \{\mathbf{x}_k\} \subset B_\delta(\mathbf{x}_*)$ for every $k \in \mathbb{N}$, by induction on k . Notice that $\mathbf{x}_0 \in B_\delta(\mathbf{x}_*)$ by construction. To start with the inductive step,

$$\sum_{i=0}^{k-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\| \leq \sqrt{\frac{2}{L}} \sum_{i=0}^{k-1} \sqrt{f_i - f_{i+1}} \leq \frac{\sqrt{2f_0(1-q)}}{\sqrt{L}(1-\sqrt{q})} \tag{74}$$

where we used (47) in the first inequality, and the second can be derived from [39, Lemma 8.1] as in the proof of [39, Theorem 4.2]. But then

$$\begin{aligned} \|\mathbf{x}_{k+1} - \mathbf{x}_*\| &\leq \|\mathbf{x}_0 - \mathbf{x}_*\| + \left(\sum_{i=0}^{k-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\| \right) + \|\mathbf{x}_k - \mathbf{x}_{k+1}\| \\ &\leq \tilde{\delta} + \frac{\sqrt{2f_0(1-q)}}{L(1-\sqrt{q})} + \sqrt{\frac{2}{L}}\sqrt{f_k - f_{k+1}} \\ &< \tilde{\delta} + \frac{\sqrt{2f_0(1-q)}}{L(1-\sqrt{q})} + \sqrt{\frac{2}{L}}\sqrt{f_k} < \delta, \end{aligned} \tag{75}$$

where we used (74) together with (47) in the second inequality, $f_{k+1} = f(\mathbf{x}_{k+1}) - f(\mathbf{x}_*) \geq 0$ in the third inequality, and (73) together with $f_0 \geq f_k$ in the last inequality.

We now have

$$\begin{aligned} \|\tilde{\mathbf{x}}_k - \mathbf{x}_*\| &\leq \|\mathbf{x}_0 - \mathbf{x}_*\| + \left(\sum_{i=0}^{k-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\| \right) + \|\mathbf{x}_k - \tilde{\mathbf{x}}_k\| \\ &\leq \|\mathbf{x}_0 - \mathbf{x}_*\| + \left(\sum_{i=0}^{k-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\| \right) + \|\mathbf{x}_k - \mathbf{x}_{k+1}\| < \delta, \end{aligned} \tag{76}$$

where we used $\|\tilde{\mathbf{x}}_k - \mathbf{x}_k\| \leq \|\mathbf{x}_{k+1} - \mathbf{x}_k\|$ in the second inequality and the last inequality follows as in (75). Thus $\tilde{\mathbf{x}}_k \in B_\delta(\mathbf{x}_*)$ as well, and the induction is complete. For GS updates the proof that $\{\tilde{\mathbf{x}}_k\} \subset B_\delta(\mathbf{x}_*)$ is analogous. \square

7.2 An active set identification criterion

We prove in this section Theorem 2, proposing a general active set identification criterion for Algorithm 1 in the special case where the feasible set \mathcal{C} is the product of simplices. With the notation introduced in Sect. 4, let $\mathcal{C}_* = \{\mathbf{x} \in \mathcal{C} : \text{supp}(\mathbf{x}) = \text{supp}(\mathbf{x}_*)\}$ and $S_* = \{\mathbf{x} \in \mathbb{R}^n : \text{supp}(\mathbf{x}) = \text{supp}(\mathbf{x}_*)\}$ be the subset of points in \mathcal{C} and the subspace of directions with the same support of \mathbf{x}_* respectively. For ease of notation, we will use the Bachmann/Landau symbols $\mathbf{y} = o(1)$ for a sequence $\mathbf{y} = \mathbf{y}_\nu \in \mathbb{R}^d$ if $\mathbf{y}_\nu \rightarrow \mathbf{0}$ as $\nu \rightarrow \infty$, and $\mathbf{y} = \Theta(1)$ if $\|\mathbf{y}_\nu\|$ is both bounded away from zero and bounded above, for all sufficiently large ν .

Definition 4 We say that the method \bar{A} has active set related directions in \mathbf{x}_* if it can do a bounded number of consecutive maximal steps, and if for some neighborhood V of \mathbf{x}_* , $\mathbf{x} \rightarrow \mathbf{x}_*$, $\mathbf{g} \rightarrow -\nabla f(\mathbf{x}_*)$ and $\mathbf{d} \in \bar{A}(\mathbf{x}, \mathbf{g})$:

- if $\mathbf{x} \in \mathcal{C}_*$ then $\mathbf{d} \in S_*$ with $\alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) = \Theta(1)$,
- if $\mathbf{x} \in \mathcal{C} \setminus \mathcal{C}_*$ then $\langle \mathbf{g}, \hat{\mathbf{d}} \rangle = \Theta(1)$.

Lemma 7 Under the assumptions of Definition 4:

- if $\mathbf{x} \in \mathcal{C} \setminus \mathcal{C}_*$ we have $\alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) = o(1)$,
- if $\mathbf{x} \in \mathcal{C}_*$, then $\langle \mathbf{g}, \hat{\mathbf{d}} \rangle = o(1)$.

Proof Notice that

$$\begin{aligned} 0 \leq \alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) \langle \mathbf{g}, \hat{\mathbf{d}} \rangle &= \langle \mathbf{g}, (\mathbf{x} + \alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) \hat{\mathbf{d}}) - \mathbf{x} \rangle \\ &= \langle -\nabla f(\mathbf{x}), (\mathbf{x} + \alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) \hat{\mathbf{d}}) - \mathbf{x} \rangle + o(1) \leq \langle -\nabla f(\mathbf{x}), \mathbf{x}_* - \mathbf{x} \rangle + o(1) = o(1). \end{aligned} \tag{77}$$

Thus

$$\alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) \leq \frac{o(1)}{\langle \mathbf{g}, \hat{\mathbf{d}} \rangle} = o(1), \tag{78}$$

where in the equality we used $\langle \mathbf{g}, \hat{\mathbf{d}} \rangle = \Theta(1)$ by assumption. This proves the first part of the claim. As for the second part, we have

$$\langle \mathbf{g}, \hat{\mathbf{d}} \rangle = \langle -\nabla f(\mathbf{x}_*), \hat{\mathbf{d}} \rangle + o(1) = o(1), \tag{79}$$

where we used $\langle -\nabla f(\mathbf{x}_*), \hat{\mathbf{d}} \rangle = 0$ in the second equality, guaranteed by stationarity conditions since $\mathbf{d} \in S_*$. \square

Proposition 1 *Let Algorithm 1 be applied to a method with active set related directions in \mathbf{x}_* as in Definition 4. Then there is a neighborhood U of \mathbf{x}_* such that if $\mathbf{x}_k \in U$ then $\text{supp}(\mathbf{x}_{k+1}) = \text{supp}(\mathbf{x}_*)$.*

Proof Let $\{\mathbf{y}_i : i \in [0 : j]\}$ be the set of points generated by $\text{SSC}(\mathbf{x}_k, -\nabla f(\mathbf{x}_k))$, \bar{T} the upper bound on the number of consecutive maximal steps, so that $T \leq \bar{T} + 1$, and let \bar{B} and B_j as in the proof of Lemma 1. We assume without loss of generality that $\|\mathbf{d}_j\| = 1$ for $j \in [0 : T]$.

We will show that for \mathbf{x}_k sufficiently close to \mathbf{x}_* certain inequalities, namely (81), (82) and (85) are satisfied, allowing us to deduce the identification property. Let

$$T^* = \max\{j \in [0 : T] : \{\mathbf{y}_i : i \in [0 : j]\} \subset \mathcal{C} \setminus \mathcal{C}_*\}$$

whenever $\mathbf{y}_0 \notin \mathcal{C}_*$, and $T^* = -1$ otherwise. We first claim $\bar{T} \geq T^* + 1$. This is clear by the definition of T^* if $T^* = -1$. Otherwise, for $j \in [0 : T^*]$ let $\tilde{\mathbf{y}}_{j+1} = \mathbf{y}_j + \alpha_{\max}^{(j)} \mathbf{d}_j$. We now show $\tilde{\mathbf{y}}_{j+1} \in \mathcal{C}_j$. First, we check $\tilde{\mathbf{y}}_{j+1} \in \text{int}(\bar{B})$. On one hand we have

$$\begin{aligned} L \|\tilde{\mathbf{y}}_{j+1} - \mathbf{y}_0\|^2 &= L \|\alpha_{\max}^{(j)} \mathbf{d}_j + \sum_{i=0}^{j-1} \alpha_i \mathbf{d}_i\|^2 \leq L \left(\sum_{i=0}^{j-1} \alpha_{\max}^{(i)} \|\mathbf{d}_i\| \right)^2 \\ &= L \left(\sum_{i=0}^{j-1} \alpha_{\max}^{(i)} \right)^2 = O(\max_{i \in [0:j]} (\alpha_{\max}^{(i)})^2), \end{aligned} \quad (80)$$

where we used $\|\mathbf{d}_i\| = 1$ by assumption in the second equality. On the other hand

$$\langle \mathbf{g}, \tilde{\mathbf{y}}_{j+1} - \mathbf{y}_0 \rangle = \alpha_{\max}^{(j)} \langle \mathbf{g}, \mathbf{d}_j \rangle + \sum_{i=0}^{j-1} \alpha_i \langle \mathbf{g}, \mathbf{d}_i \rangle = O(\max_{i \in [0:j]} \alpha_{\max}^{(i)}). \quad (81)$$

Since $\mathbf{y}_{(i)} \in \mathcal{C} \setminus \mathcal{C}_*$, by the active set related property $\alpha_{\max}^{(i)} = o(1)$ for $\mathbf{x}_k \rightarrow \mathbf{x}_*$. Let now M_1 and M_2 be the implicit constants in (80) and (81). For $\mathbf{y}_0 = \mathbf{x}_k$ close enough to \mathbf{x}_* we obtain

$$L \|\tilde{\mathbf{y}}_{j+1} - \mathbf{y}_0\|^2 \leq M_1 \max_{i \in [0:j]} (\alpha_{\max}^{(i)})^2 < M_2 \max_{i \in [0:j]} \alpha_{\max}^{(i)} \leq \langle \mathbf{g}, \tilde{\mathbf{y}}_{j+1} - \mathbf{y}_0 \rangle, \quad (82)$$

where we used (80) in the first inequality, $\max_{i \in [0:j]} \alpha_{\max}^{(i)} = o(1)$ in the second inequality and (81) in the last inequality. From (82), $\tilde{\mathbf{y}}_{j+1} \in \text{int}\bar{B}$ follows easily as desired.

We now need to check $\tilde{\mathbf{y}}_{j+1} \in B_j$. Reasoning as above, on the one hand we have $\frac{\|\mathbf{g}\|}{2L} = \Theta(1)$ for $\mathbf{x}_k \rightarrow \mathbf{x}$ (setting aside the trivial case where $-\nabla f(\mathbf{x}_*) = 0$), and on

the other hand $\|\tilde{\mathbf{y}}_{j+1} - \mathbf{y}_0\| = o(1)$ by (80), so that

$$\|\tilde{\mathbf{y}}_{j+1} - \mathbf{y}_0\| < \frac{\|\mathbf{g}\|}{2L} \tag{83}$$

for \mathbf{y}_0 close enough to \mathbf{x}_* and $\tilde{\mathbf{y}}_{j+1} \in B_j$ as desired. Then $\tilde{\mathbf{y}}_{j+1} \in \text{int}(\mathcal{C}_j)$ for $j \in [0: T^*]$, or equivalently $\beta_j > \alpha_{\max}^{(j)}$ the SSC does always maximal steps in the first $T^* + 1$ iterations. In particular, it generates the point $\mathbf{y}_{T^*+1} \in \mathcal{C}_* \setminus \{\mathbf{y}_{T^*}\}$. The claim is thus proved.

If \mathbf{y}_{T^*+1} is stationary for g , the SSC terminates at step 4 with output $\mathbf{y}_{T^*+1} \in \mathcal{C}_*$ and the thesis is proved. Otherwise, we claim that the SSC terminates with output $\mathbf{y}_{T^*+2} \in \mathcal{C}_*$ and $\beta_{T^*+1} < \alpha_{\max}^{(T^*+1)}$. First, observe that by assumption we must have $\mathbf{d}_{T^*+1} \in S_*$, and therefore $\mathbf{y}_{T^*+2} = \mathbf{y}_{T^*+1} + \alpha_{T^*+1} \mathbf{d}_{T^*+1} \in \mathcal{C}_*$. Second, we have $\alpha_{\max}^{(T^*+1)} = \Theta(1)$, and at the same time

$$\beta_{T^*+1} \leq \text{diam}(\mathcal{C}_{T^*+1}) \leq \text{diam}(B_{T^*+1}) \leq 2\langle \mathbf{g}_{T^*+1}, \mathbf{d}_{T^*+1} \rangle = o(1). \tag{84}$$

Thus for \mathbf{y}_0 close enough to \mathbf{x}_* we must have

$$\beta_{T^*+1} < \alpha_{\max}^{(T^*+1)}, \tag{85}$$

and the claim is proved. Since the SSC terminates either with \mathbf{y}_{T^*+1} or \mathbf{y}_{T^*+2} , and both of these points are in \mathcal{C}_* , the thesis follows. □ □

Lemma 8 For $\mathbf{x} \rightarrow \mathbf{x}_*$, $\mathbf{g} \rightarrow -\nabla f(\mathbf{x}_*)$, if $\mathbf{x} \in \mathcal{C}_*$, $\mathbf{d} \in S_*$ and $\alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}})$ coincides with the maximal feasible stepsize, then $\alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) = \Theta(1)$.

Proof We have

$$\alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) = \min_{i:\hat{d}_i < 0} \frac{x_i}{|\hat{d}_i|} \geq \min_{i:\hat{d}_i < 0} x_i \geq \min_{i \in \text{supp}(\mathbf{x}_*)} x_i = \Theta(1), \tag{86}$$

where we used $|\hat{d}_i| \leq \|\hat{\mathbf{d}}\| \leq 1$ in the first inequality, $\text{supp}(\mathbf{d}) \subseteq \text{supp}(\mathbf{x}_*)$ in the second inequality, and $x_i \rightarrow x_{*,i} > 0$ in the third one. □

For $\mathbf{x} \in \mathcal{C}$, we define the expression

$$\lambda^{(i)}(\mathbf{x}, \mathbf{g}) = \langle \mathbf{g}^{(i)}, \mathbf{x}^{(i)} \rangle \mathbf{e}^{(i)} - \mathbf{g}^{(i)}, \quad i \in [1:m],$$

and the Lagrangian multiplier vector

$$\lambda^{(i)}(\mathbf{x}) = \lambda^{(i)}(\mathbf{x}, -\nabla f(\mathbf{x})) = \nabla f(\mathbf{x})^{(i)} - \langle \nabla f(\mathbf{x})^{(i)}, \mathbf{x}^{(i)} \rangle \mathbf{e}^{(i)}, \quad i \in [1:m]. \tag{87}$$

We notice that strict complementarity holds at a stationary point $\mathbf{x}_* \in \mathcal{C}$ for $\nabla f(\mathbf{x}_*)$ if and only if it holds for every $i \in [1:m]$ at $\mathbf{x}_*^{(i)} \in \mathcal{C}^{(i)}$ and $\nabla f(\mathbf{x}_*)^{(i)}$.

Lemma 9 Assume that strict complementarity holds at \mathbf{x}_* . Then the AFW applied to the simplex has active set related directions in \mathbf{x}_* as in Definition 4.

Proof For $\mathbf{x} \rightarrow \mathbf{x}_*$ and $\mathbf{g} \rightarrow -\nabla f(\mathbf{x}_*)$ we have $\lambda(\mathbf{x}, \mathbf{g}) \rightarrow \lambda(\mathbf{x}_*)$, and therefore in particular $\lambda_i(\mathbf{x}, \mathbf{g}) \rightarrow 0$ for $i \in \text{supp}(\mathbf{x}_*)$ while $\lambda_i(\mathbf{x}, \mathbf{g}) \rightarrow \lambda_i(\mathbf{x}_*) > 0$ for $i \in [1:n] \setminus \text{supp}(\mathbf{x}_*)$. Therefore, for $\mathbf{x} \in \mathcal{C} \setminus \mathcal{C}_*$ close enough to \mathbf{x}_* we must have $\max\{\lambda_i(\mathbf{x}, \mathbf{g}) : i \in \text{supp}(\mathbf{x})\} > \max\{-\lambda_i(\mathbf{x}, \mathbf{g}) : i \in [1:n]\}$, so that by [13, Lemma 3.2(a)] we have that the descent direction selected by the AFW satisfies $\mathbf{d} = \mathbf{x} - \mathbf{e}_{\hat{i}}$ for some $\hat{i} \in \text{argmax}\{\lambda_i(\mathbf{x}, \mathbf{g}) : i \in \text{supp}(\mathbf{x})\} \subset [1:n] \setminus \text{supp}(\mathbf{x}_*)$. Therefore

$$\langle \mathbf{g}, \hat{\mathbf{d}} \rangle = \langle \mathbf{g}, \frac{\mathbf{x} - \mathbf{e}_{\hat{i}}}{\|\mathbf{x} - \mathbf{e}_{\hat{i}}\|} \rangle = \frac{\lambda_{\hat{i}}(\mathbf{x}, \mathbf{g})}{\|\mathbf{x} - \mathbf{e}_{\hat{i}}\|} = \Theta(1) \quad (88)$$

for $\mathbf{x} \rightarrow \mathbf{x}_*$ and $\mathbf{g} \rightarrow -\nabla f(\mathbf{x}_*)$.

As for the case $\mathbf{x} \in \mathcal{C}_*$, then if \mathbf{x}, \mathbf{g} are close enough to \mathbf{x}_* we must have $\lambda_i(\mathbf{x}, \mathbf{g}) > 0$ for every i in $[1:n] \setminus \text{supp}(\mathbf{x}_*)$. Therefore by [13, Lemma 3.2(b)] if \mathbf{y} is obtained from \mathbf{x} with a FW update we must have $\mathbf{y}_i = 0$ for $i \in [1:n] \setminus \text{supp}(\mathbf{x}_*)$, which is equivalent to say that the update direction must be in S_* . The property $\alpha_{\max}(\mathbf{x}, \hat{\mathbf{d}}) = \Theta(1)$ follows by Lemma 8. \square

Proof of Theorem 2 Follows by applying the property proved in Lemma 9 to each block selected by the method.

Acknowledgements The authors are grateful for the diligence of two excellent referees who suggested significant improvements which we were happy to follow. Thanks are also due to the editorial team for enabling an efficient and constructive evaluation process.

Funding Open access funding provided by Università degli Studi di Padova within the CRUI-CARE Agreement. The work of Francesco Rinaldi has been partially funded by the European Union - Next Generation EU under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.1 - Call PRIN 2022 No. 104 of February 2, 2022 of Italian Ministry of University and Research; Project 2022BMBW2A (subject area: PE - Physical Sciences and Engineering) "Large-scale optimization for sustainable and resilient energy systems."

Declarations

Conflict of interest Nothing to declare by all of the authors.

Ethical approval The entire research work and writing of this article was performed under strict compliance with widely accepted ethical standards by all authors.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alaíz, C.M., Barbero, A., Dorronsoro, J.R.: Group fused lasso. In: International Conference on Artificial Neural Networks, pp. 66–73. Springer (2013)
2. Alon, N., Spencer, J.H.: The Probabilistic Method. Wiley, New York (2016)
3. Attouch, H., Bolte, J., Redont, P., Soubeyran, A.: Proximal alternating minimization and projection methods for nonconvex problems: an approach based on the kurdyka-lojasiewicz inequality. *Math. Oper. Res.* **35**(2), 438–457 (2010)
4. Beck, A.: First-Order Methods in Optimization. SIAM, Philadelphia (2017)
5. Bertsekas, D., Tsitsiklis, J.: Parallel and Distributed Computation: Numerical Methods. Athena Scientific, Nashua (2015)
6. Birgin, E.G., Martínez, J.M., Raydan, M.: Nonmonotone spectral projected gradient methods on convex sets. *SIAM J. Optim.* **10**(4), 1196–1211 (2000)
7. Bolte, J., Daniilidis, A., Lewis, A., Shiota, M.: Clarke subgradients of stratifiable functions. *SIAM J. Optim.* **18**(2), 556–572 (2007)
8. Bolte, J., Daniilidis, A., Ley, O., Mazet, L.: Characterizations of Łojasiewicz inequalities: subgradient flows, talweg, convexity. *Trans. Am. Math. Soc.* **362**(6), 3319–3363 (2010)
9. Bolte, J., Nguyen, T.P., Peyrouquet, J., Suter, B.W.: From error bounds to the complexity of first-order descent methods for convex functions. *Math. Program.* **165**(2), 471–507 (2017)
10. Bomze, I.M., Budinich, M., Pardalos, P.M., Pelillo, M.: The maximum clique problem. In: Handbook of Combinatorial Optimization, pp. 1–74. Springer, Cham (1999)
11. Bomze, I.M., Gabl, M., Maggioni, F., Pflug, G.: Two-stage stochastic standard quadratic optimization. *Eur. J. Oper. Res.* **299**(1), 21–34 (2022)
12. Bomze, I.M., Rinaldi, F., Rota Bulò, S.: First-order methods for the impatient: support identification in finite time with convergent Frank–Wolfe variants. *SIAM J. Optim.* **29**(3), 2211–2226 (2019)
13. Bomze, I.M., Rinaldi, F., Zeffiro, D.: Active set complexity of the away-step Frank–Wolfe algorithm. *SIAM J. Optim.* **30**(3), 2470–2500 (2020)
14. Bomze, I.M., Rinaldi, F., Zeffiro, D.: Frank–Wolfe and friends: a journey into projection-free first-order optimization methods. *4OR* **19**(3), 313–345 (2021)
15. Bomze, I.M., Rinaldi, F., Zeffiro, D.: Fast cluster detection in networks by first order optimization. *SIAM J. Math. Data Sci.* **4**(1), 285–305 (2022)
16. Bomze, I.M., Schachinger, W.: Multi-standard quadratic optimization: interior point methods and cone programming reformulation. *Comput. Optim. Appl.* **45**(2), 237–256 (2010)
17. Bomze, I.M., Schachinger, W., Ullrich, R.: The complexity of simple models: a study of worst and typical hard cases for the standard quadratic optimization problem. *Math. Oper. Res.* **43**(2), 347–692 (2017)
18. Boumal, N.: An Introduction to Optimization on Smooth Manifolds, vol. 3. Cambridge University Press, Cambridge (2020)
19. Calamai, P.H., Moré, J.J.: Projected gradient methods for linearly constrained problems. *Math. Program.* **39**(1), 93–116 (1987)
20. Combettes, C.W., Pokutta, S.: Complexity of linear minimization and projection on some sets. *Oper. Res. Lett.* **49**(4), 565–571 (2021)
21. Di Serafino, D., Toraldo, G., Viola, M., Barlow, J.: A two-phase gradient method for quadratic programming problems with a single linear constraint and bounds on the variables. *SIAM J. Optim.* **28**(4), 2809–2838 (2018)
22. Foygel, R., Horrell, M., Drton, M., Lafferty, J.: Nonparametric reduced rank regression. *Adv. Neural Inf. Process. Syst.* **25** (2012)
23. Fukunaga, T., Kasai, H.: Fast block-coordinate Frank–Wolfe algorithm for semi-relaxed optimal transport. arXiv preprint [arXiv:2103.05857](https://arxiv.org/abs/2103.05857) (2021)
24. Garber, D.: Revisiting Frank–Wolfe for polytopes: strict complementary and sparsity. arXiv preprint [arXiv:2006.00558](https://arxiv.org/abs/2006.00558) (2020)
25. Jaggi, M.: Revisiting Frank–Wolfe: projection-free sparse convex optimization. *ICML* **1**, 427–435 (2013)
26. Jegelka, S., Bach, F., Sra, S.: Reflection methods for user-friendly submodular optimization. *Adv. Neural Inf. Process. Syst.* **26** (2013)

27. Karimi, H., Nutini, J., Schmidt, M.: Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, pp. 795–811. Springer (2016)
28. Lacoste-Julien, S., Jaggi, M., Schmidt, M., Pletscher, P.: Block-coordinate Frank-Wolfe optimization for structural SVMs. In: S. Dasgupta, D. McAllester (eds.) Proceedings of the 30th International Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 28(1), pp. 53–61. PMLR, Atlanta, Georgia, USA (2013). <http://proceedings.mlr.press/v28/lacoste-julien13.html>
29. Lan, G.: First-order and Stochastic Optimization Methods for Machine Learning. Springer (2020)
30. Leary, R.H.: Global optimization on funneling landscapes. *J. Global Optim.* **18**(4), 367–383 (2000)
31. LeBlanc, L.J., Morlok, E.K., Pierskalla, W.P.: An efficient approach to solving the road network equilibrium traffic assignment problem. *Transp. Res.* **9**(5), 309–318 (1975)
32. Liu, J., Musialski, P., Wonka, P., Ye, J.: Tensor completion for estimating missing values in visual data. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(1), 208–220 (2012)
33. Locatelli, M., Schoen, F.: Global Optimization: Theory, Algorithms, and Applications. SIAM, Philadelphia (2013)
34. Luo, Z.Q., Tseng, P.: On the convergence of the coordinate descent method for convex differentiable minimization. *J. Optim. Theory Appl.* **72**(1), 7–35 (1992)
35. Nesterov, Y.: Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM J. Optim.* **22**(2), 341–362 (2012)
36. Osokin, A., Alayrac, J.B., Lukasevitz, I., Dokania, P., Lacoste-Julien, S.: Minding the gaps for block Frank-Wolfe optimization of structured svms. In: International Conference on Machine Learning, pp. 593–602. PMLR (2016)
37. Richtárik, P., Takáč, M.: Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Math. Program.* **144**(1), 1–38 (2014)
38. Rinaldi, F., Zeffiro, D.: A unifying framework for the analysis of projection-free first-order methods under a sufficient slope condition. arXiv preprint [arXiv:2008.09781](https://arxiv.org/abs/2008.09781) (2020)
39. Rinaldi, F., Zeffiro, D.: Avoiding bad steps in Frank Wolfe variants. *Comput. Optim. Appl.* **84**, 225–264 (2023)
40. di Serafino, D., Hager, W.W., Toraldo, G., Viola, M.: On the stationarity for nonlinear optimization problems with polyhedral constraints. *Mathematical Programming* pp. 1–28 (2023)
41. Wang, Y.X., Sadhanala, V., Dai, W., Neiswanger, W., Sra, S., Xing, E.: Parallel and distributed block-coordinate Frank-Wolfe algorithms. In: International Conference on Machine Learning, pp. 1548–1557. PMLR (2016)

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