

SuperMann: A Superlinearly Convergent Algorithm for Finding Fixed Points of Nonexpansive Operators

Themelis, Andreas

Department of Electrical Engineering (ESAT-STADIUS), KU Leuven

Patrinos, Panagiotis

Department of Electrical Engineering (ESAT-STADIUS), KU Leuven

<https://hdl.handle.net/2324/4399977>

出版情報 : IEEE transactions on automatic control. 64 (12), pp.4875-4890, 2019-03-27. IEEE
バージョン :

権利関係 : © 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

SuperMann: a superlinearly convergent algorithm for finding fixed points of nonexpansive operators

Andreas Themelis and Panagiotis Patrinos

Abstract—Operator splitting techniques have recently gained popularity in convex optimization problems arising in various control fields. Being fixed-point iterations of nonexpansive operators, such methods suffer many well known downsides, which include high sensitivity to ill conditioning and parameter selection, and consequent low accuracy and robustness. As universal solution we propose *SuperMann*, a Newton-type algorithm for finding fixed points of nonexpansive operators. It generalizes the classical Krasnosel’skiĭ-Mann scheme, enjoys its favorable global convergence properties and requires exactly the same oracle. It is based on a novel separating hyperplane projection tailored for nonexpansive mappings which makes it possible to include steps along any direction. In particular, when the directions satisfy a Dennis-Moré condition we show that *SuperMann* converges superlinearly under mild assumptions, which, surprisingly, do not entail nonsingularity of the Jacobian at the solution but merely metric subregularity. As a result, *SuperMann* enhances and robustifies all operator splitting schemes for structured convex optimization, overcoming their well known sensitivity to ill conditioning.

I. INTRODUCTION

Operator splitting techniques (also known as proximal algorithms), introduced in the 50’s for solving PDEs and optimal control problems, have been successfully used to reduce complex problems into a series of simpler subproblems. The most well known operator splitting methods are the alternating direction method of multipliers (ADMM), forward-backward splitting (FBS) also known as proximal-gradient method in composite convex minimization, Douglas-Rachford splitting (DRS) and the alternating minimization method (AMM) [1]. Operator splitting techniques pose several advantages over traditional optimization methods such as sequential quadratic programming and interior point methods: (1) they can easily handle nonsmooth terms and abstract linear operators, (2) each iteration requires only simple arithmetic operations, (3) the algorithms scale gracefully as the dimension of the problem increases, and (4) they naturally lead to parallel and distributed implementation. Therefore, operator splitting methods cope well with limited amount of hardware resources making them particularly attractive for (embedded) control [2], signal processing [3], and distributed optimization [4], [5].

The key idea behind these techniques when applied to convex optimization is to reformulate the optimality conditions of the problem at hand into a problem of finding a fixed point

of a nonexpansive operator and then apply relaxed fixed-point iterations. Although sometimes a fast convergence rate can be observed, the norm of the fixed-point residual decreases, at best, with Q -linear rate, and due to an inherent sensitivity to ill conditioning oftentimes the Q -factor is close to one. Moreover, all operator splitting methods are basically “open-loop”, since the tuning parameters, such as stepsizes and preconditioning, must be set before their execution. In fact, such methods are very sensitive to the choice of parameters and sometimes there is not even a concrete way of selecting them, as it is the case of ADMM. All these are serious obstacles when it comes to using such types of algorithms for real-time applications such as embedded MPC, or to reliably solve cone programs.

As an attempt to solve the issue, people have considered the employment of variable metrics to reshape the geometry of the problem and enhance convergence rate [6]. However, unless such metrics have a very specific structure, even for simple problems the cost of operating in the new geometry outweighs the benefits.

Another interesting approach that is gaining more and more popularity tries to exploit possible sparsity patterns by means of chordal decomposition techniques [7]. These methods can improve scalability and reduce memory usage, but unless the problem comes with an inherent sparse structure they yield no tangible benefit.

Alternatively, the task of searching fixed points of an operator T can be translated to that of finding zeros of the corresponding residual $R = \text{id} - T$. Many methods with fast asymptotic convergence rates such as Newton-type exist that can be employed for efficiently solving nonlinear equations, see, e.g., [8, §7] and [9]. However, such methods converge only when close enough to the solution, and in order to globalize the convergence there comes the need of a merit function to perform a linesearch along candidate directions of descent. The typical choice of the square residual $\|Rx\|^2$ unfortunately is of no use, as in meaningful applications R is nonsmooth.

A. Proposed methodology

In response to these issues, in this paper we propose a universal scheme that globalizes Newton-type methods for finding fixed points of any nonexpansive operator on real Hilbert spaces. Admittedly with an intended pun, since it exhibits *superlinear* convergence rates and generalizes the Krasnosel’skiĭ-Mann iterations we name our algorithm *SuperMann*. The method is based on a novel hyperplane projection step tailored for nonexpansive mappings.

Furthermore, we consider a modified Broyden’s scheme which was first introduced in [10] and show how it fits into

Department of Electrical Engineering (ESAT-STADIUS) – KU Leuven, Kasteelpark Arenberg 10, 3001 Leuven, Belgium. {andreas.themelis, panos.patrinos}@esat.kuleuven.be.

The work was supported by: FWO research projects G086518N and G086318N; Fonds de la Recherche Scientifique – FNRS and the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS Project no 30468160; KU Leuven – Internal Funding C11-18-00411.

our framework enabling superlinear asymptotic convergence rates. One of the most appealing properties of *SuperMann* is that thanks to its quasi-Fejérian behavior, achieving superlinear convergence does not necessitate nonsingularity of the Jacobian at the solution, which is the usual requirement of quasi-Newton schemes, but merely metric subregularity. This property considerably widens the range of problems which can be solved efficiently, in that, for instance, the solutions need not be isolated for superlinear convergence to take place.

B. Contributions

Our contributions can be summarized as follows:

- (1) In Section IV we design a universal algorithmic framework (Algorithm 1) for finding fixed points of nonexpansive operators, which generalizes the classical Krasnosel'skiĭ-Mann scheme and possesses its same global and local convergence properties.
- (2) In Section V we introduce a novel separating hyperplane projection tailored for nonexpansive mappings; based on this, in Definition V.3 we then propose a generalized KM iteration (GKM).
- (3) We define a linesearch based on the novel projection, suited for any nonexpansive operator and update direction (Theorem V.4).
- (4) In Section VI we combine these ideas and derive the *SuperMann scheme* (Alg. 2), an algorithm that
 - globalizes the convergence of Newton-type methods for finding fixed points of nonexpansive operators (Theorem VI.1);
 - reduces to the local method $x_{k+1} = x_k + d_k$ when the directions d_k are *superlinear*, as it is the case for a modified Broyden's scheme (Theorems VI.4 and VI.8);
 - has superlinear convergence guarantees even without the usual requirement of nonsingularity of the Jacobian at the limit point, but simply under metric subregularity; in particular, the solution need not be unique!

C. Paper organization

The paper is organized as follows. Section II serves as an informal introduction to highlight the known limitations of fixed-point iterations and to motivate our interest in Newton-type methods with some toy examples. The formal presentation begins in Section III with the introduction of some basic notation and known facts. In Section IV we define the problem at hand and propose a general abstract algorithmic framework for solving it. In Section V we provide a generalization of the classical KM iterations that is key for the global convergence and performance of *SuperMann*, an algorithm which is presented and analyzed in Section VI. Finally, in Section VII we show how the theoretical findings are backed up by promising numerical simulations, where *SuperMann* dramatically improves classical splitting schemes. For the sake of readability some of the proofs are deferred to the Appendix.

II. MOTIVATING EXAMPLES

Given a nonexpansive operator $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$, consider the problem of finding a fixed point, *i.e.*, a point $x_\star \in \mathbb{R}^n$ such that $x_\star = Tx_\star$. The independent works of Krasnosel'skiĭ and Mann [11], [12] provided a very elegant solution which is simply based on recursive iterations $x^+ = (1-\alpha)x + \alpha Tx$ with $\alpha \in (0, \bar{\alpha})$ for some $\bar{\alpha} \geq 1$. The method, known as Krasnosel'skiĭ-Mann scheme or KM scheme for short, has been studied intensively ever since, also because it generalizes a plethora of optimization algorithms. It is well known that the scheme is globally convergent with square-summable and monotonically decreasing residual $R = \text{id} - T$ (in norm), and also locally Q -linearly convergent if R is *metrically subregular* at the limit point x_\star . Metric subregularity basically amounts to requiring the distance from the set of solutions to be upper bounded by a multiple of the norm of R for all points sufficiently close to x_\star ; it is quite mild a requirement — for instance, it does not entail x_\star to be an isolated solution — and as such linear convergence is quite frequent in practice. However, the major drawback of the KM scheme is its high sensitivity to ill conditioning of the problem, and cases where convergence is prohibitively slow in practice despite the theoretical (sub)linear rate are also abundant. Illustrative examples can be easily constructed for the problem of finding a point in the intersection of two closed convex sets C_1 and C_2 with $C_1 \cap C_2 \neq \emptyset$. The problem can be solved by means of fixed-point iterations of the (nonexpansive) *alternating projections* operator $T = \Pi_{C_2} \circ \Pi_{C_1}$.

In Figure 1a we consider the case of two polyhedral cones, namely $C_1 = \{x \in \mathbb{R}^2 \mid 0.1x_1 \leq x_2 \leq 0.2x_1\}$ and $C_2 = \{x \in \mathbb{R}^2 \mid 0.3x_1 \leq x_2 \leq 0.35x_1\}$. Alternating projections is then linearly convergent (to the unique intersection point 0) due to the fact that $R = \text{id} - T$ is piecewise affine and hence globally metrically subregular. However, the convergence is extremely slow due to the pathological small angle between the two cones, as it is apparent in Figure 1a.

As an attempt to overcome this frequent phenomenon, [13] proposes a foretracking linesearch heuristic which is particularly effective when subsequent fixed-point iterations proceed along almost parallel directions. Iteration-wise, in such instances the linesearch does yield a considerable improvement upon the plain KM scheme; however, each foretrack prescribes extra evaluations of T and unless T has a specific structure the computational overhead might outweigh the advantages. Moreover, its asymptotic convergence rates do not improve upon the plain KM scheme. Figure 1b illustrates this fact relative to $C_1 = \{x \in \mathbb{R}^2 \mid x_1^2 + x_2^2 \leq 1\}$ and $C_2 = \{x \in \mathbb{R}^2 \mid x_1 = 1\}$. Despite a good performance on early iterations, the linesearch cannot improve the asymptotic sublinear rate of the plain KM scheme due to the fact that the residual is not metrically subregular at the (unique) solution $x_\star = (0, 1)$. In particular, it is evident that medium-to-high accuracy cannot be achieved in a reasonable number of iterations with either methods.

In response to this limitation there comes the need to include some “first-order-like information”. Specifically, the problem of finding a fixed point of T can be rephrased in terms of solving the system of nonlinear (monotone) equations $Rx = 0$, which could *possibly* be solved efficiently with Newton-type

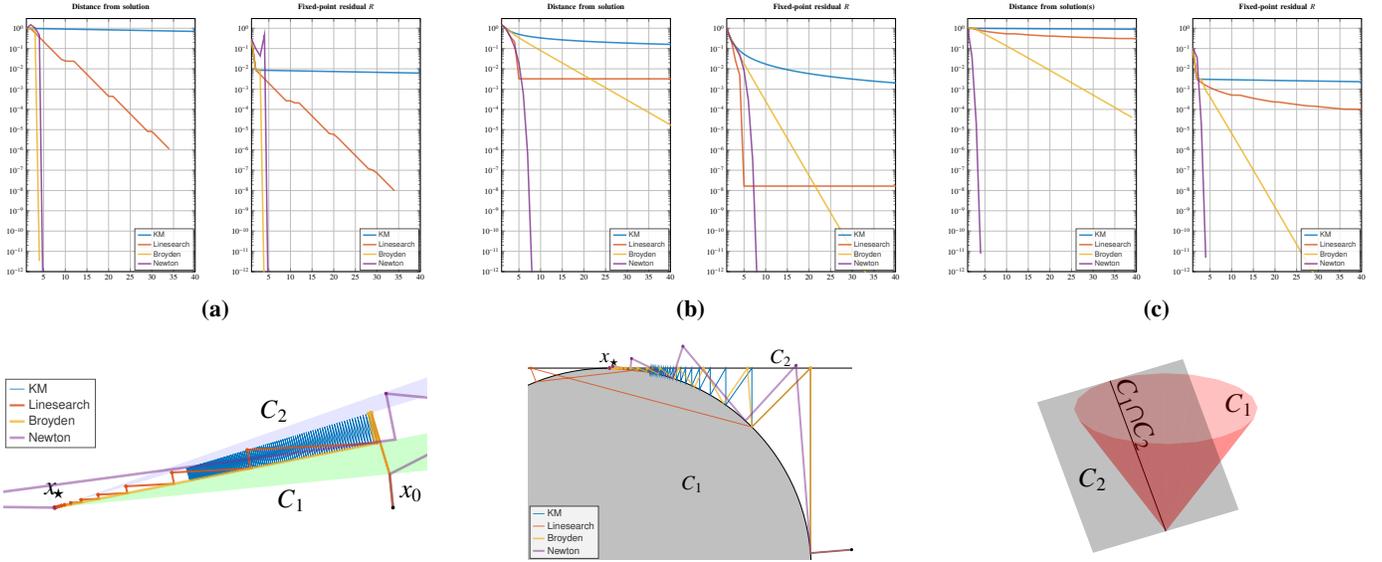


Figure 1: (a) Alternating projections on polyhedral cones. $R = \text{id} - \Pi_{C_2} \circ \Pi_{C_1}$ is globally metrically subregular, however the Q -linear convergence of the KM scheme is very slow. (b) Alternating projections on ball and tangent line. With or without linesearch the KM scheme is not linearly convergent due to the fact that the residual R is not metrically subregular at x_* . (c) Alternating projections on second-order cone and tangent plane. In contrast with the slow sublinear rate of KM both with and without linesearch, and despite the non isolatedness of any solution, Broyden’s scheme exhibits an appealing linear convergence rate.

methods. In the toy simulations of this section, the purple lines correspond to the semismooth Newton iterations

$$x^+ = x - G^{-1}Rx \quad \text{for some } G \in \partial R,$$

where ∂R is the Clarke generalized Jacobian of R [8, Def. 7.1.1]. Interestingly, in the proposed simulations this method exhibits fast convergence even when the limit point is a non isolated solution, as in the case of the second-order cone $C_1 = \{x \in \mathbb{R}^3 \mid x_3 \geq 0.1 \sqrt{x_1^2 + x_2^2}\}$ and the tangent plane $C_2 = \{x \in \mathbb{R}^3 \mid x_3 = 0.1x_2\}$ considered in Figure 1c.

However, computing the generalized Jacobian might be too demanding and require extra information not available in close form. For this reason we focus on quasi-Newton methods

$$x^+ = x - HRx,$$

where the linear operator H is progressively updated with only evaluations of R and direct linear algebra in such a way that the vector HRx is asymptotically a good approximation of a Newton direction $G^{-1}Rx$. The yellow lines in the simulations of this section correspond to H being selected with Broyden’s quasi-Newton method.

The crucial issue is convergence itself. Though in these trivial simulations it is not the case, it is well known that Newton-type methods in general converge only when close to a solution, and may even diverge otherwise. In fact, globalizing the convergence of Newton-type methods is a key challenge in optimization, as the dedicated recent book [9] confirms.

In this paper we provide the SuperMann scheme, a globalization strategy for Newton-type methods (or any local scheme in general) that applies to any (nonsmooth) monotone equation deriving from fixed-point iterations of nonexpansive operators. Our method covers almost all splitting schemes in con-

vex optimization, such as forward-backward splitting (FBS, also known as proximal gradient method), Douglas-Rachford splitting (DRS) and the alternating direction method of multipliers (ADMM), to name a few. We also provide sufficient conditions at the limit point under which the method reduces to the local scheme and converges superlinearly.

III. NOTATION AND KNOWN RESULTS

With $\text{bdry } A$ we denote the boundary of the set A , and given a sequence $(x_k)_{k \in \mathbb{N}}$ we write $(x_k)_{k \in \mathbb{N}} \subset A$ to indicate that $x_k \in A$ for all $k \in \mathbb{N}$. For $p > 0$ we let

$$\ell^p := \{(x_k)_{k \in \mathbb{N}} \subset \mathbb{R} \mid \sum_{k \in \mathbb{N}} |x_k|^p < \infty\}$$

denote the set of real-valued sequences with summable p -th power, and with ℓ_+^p the subset of the positive-valued ones.

The positive part of $x \in \mathbb{R}$ is $[x]_+ := \max\{x, 0\}$.

A. Hilbert spaces and bounded linear operators

Throughout the paper, \mathcal{H} is a real separable Hilbert space endowed with an inner product $\langle \cdot, \cdot \rangle$ and with induced norm $\|\cdot\|$. The Euclidean norm and scalar product are denoted as $\|\cdot\|_2$ and $\langle \cdot, \cdot \rangle_2$, respectively. For $\bar{x} \in \mathcal{H}$ and $r > 0$, the open ball centered at \bar{x} with radius r is indicated as $\mathbf{B}(\bar{x}; r) := \{x \in \mathcal{H} \mid \|x - \bar{x}\| < r\}$. For a closed and nonempty convex set $C \subseteq \mathcal{H}$ we let Π_C denote the projection operator on C .

Given $(x_k)_{k \in \mathbb{N}} \subset \mathcal{H}$ and $x \in \mathcal{H}$ we write $x_k \rightarrow x$ and $x_k \rightharpoonup x$ to denote, respectively, strong and weak convergence of $(x_k)_{k \in \mathbb{N}}$ to x . The set of weak sequential cluster points of $(x_k)_{k \in \mathbb{N}}$ is indicated as $\mathcal{W}(x_k)_{k \in \mathbb{N}}$.

The set of bounded linear operators $\mathcal{H} \rightarrow \mathcal{H}$ is denoted as $\mathcal{B}(\mathcal{H})$. The adjoint operator of $L \in \mathcal{B}(\mathcal{H})$ is indicated as L^* , i.e., the unique operator in $\mathcal{B}(\mathcal{H})$ such that $\langle Lx, y \rangle = \langle x, L^*y \rangle$ for all $x, y \in \mathcal{H}$.

B. Nonexpansive operators and Fejér sequences

We now briefly recap some known definitions and results of nonexpansive operator theory that will be used in the paper.

Definition III.1. An operator $T : \mathcal{H} \rightarrow \mathcal{H}$ is said to be

- (i) nonexpansive (NE) if $\|Tx - Ty\| \leq \|x - y\|$ for all $x, y \in \mathcal{H}$;
- (ii) averaged if it is α -averaged for some $\alpha \in (0, 1)$, i.e., if there exists a nonexpansive operator $S : \mathcal{H} \rightarrow \mathcal{H}$ such that $T = (1 - \alpha)\text{id} + \alpha S$;
- (iii) firmly nonexpansive (FNE) if it is $1/2$ -averaged.

Clearly, for any NE operator T the residual $R = \text{id} - T$ is monotone, in the sense that $\langle Rx - Ry, x - y \rangle \geq 0$ for all $x, y \in \mathcal{H}$; if T is additionally FNE, then not only is R monotone, but it is FNE as well. For notational convenience we extend the definition of α -averagedness to the case $\alpha = 1$ which reduces to plain nonexpansiveness.

Given an operator $T : \mathcal{H} \rightarrow \mathcal{H}$ we let

$$\mathbf{zer} T := \{z \in \mathcal{H} \mid Tz = 0\} \quad \text{and} \quad \mathbf{fix} T := \{z \in \mathcal{H} \mid Tz = z\}$$

denote the set of its *zeros* and *fixed points*, respectively. For $\lambda \in \mathbb{R}$ we define the λ -averaging of T as

$$T_\lambda := (1 - \lambda)\text{id} + \lambda T.$$

Notice that

$$\text{id} - T_\lambda = \lambda(\text{id} - T) \quad \text{for all } \lambda \in \mathbb{R}, \quad (1)$$

and therefore $\mathbf{fix} T_\lambda = \mathbf{fix} T$ for all $\lambda \neq 0$. Moreover, if T is α -averaged and $\lambda \in (0, 1/\alpha]$, then

$$T_\lambda \text{ is } \alpha\lambda\text{-averaged} \quad (2)$$

[14, Cor. 4.28] and in particular $T_{1/2\alpha}$ is FNE.

Definition III.2. Relative to a nonempty set $S \subseteq \mathcal{H}$, a sequence $(x_k)_{k \in \mathbb{N}} \subset \mathcal{H}$ is

- (i) Fejér (-monotone) if $\|x_{k+1} - s\| \leq \|x_k - s\|$ for all $k \in \mathbb{N}$ and $s \in S$;
- (ii) quasi-Fejér (monotone) if for all $s \in S$ there exists a sequence $(\varepsilon_k(s))_{k \in \mathbb{N}} \in \ell_+^1$ such that

$$\|x_{k+1} - s\|^2 \leq \|x_k - s\|^2 + \varepsilon_k(s) \quad \forall k \in \mathbb{N}.$$

This definition of quasi-Fejér monotonicity is taken from [15] where it is referred to as *of type III*, and generalizes the classical definition [16].

Theorem III.3. Let $T : \mathcal{H} \rightarrow \mathcal{H}$ be an NE operator with $\mathbf{fix} T \neq \emptyset$, and suppose that $(x_k)_{k \in \mathbb{N}} \subset \mathcal{H}$ is quasi-Fejér with respect to $\mathbf{fix} T$. If $(x_k - Tx_k)_{k \in \mathbb{N}} \rightarrow 0$, then there exists $x_\star \in \mathbf{fix} T$ such that $x_k \rightarrow x_\star$.

Proof. From [15, Prop. 3.7(i)] we have $\mathcal{W}(x_k)_{k \in \mathbb{N}} \neq \emptyset$; in turn, from [14, Cor. 4.18] we infer that $\mathcal{W}(x_k)_{k \in \mathbb{N}} \subseteq \mathbf{fix} T$. The claim then follows from [15, Thm. 3.8]. \square

IV. GENERAL ABSTRACT FRAMEWORK

Unless differently specified, in the rest of the paper we work under the following assumption.

Assumption I. $T : \mathcal{H} \rightarrow \mathcal{H}$ is an α -averaged operator for some $\alpha \in (0, 1]$ and with $\mathbf{fix} T \neq \emptyset$. With $R := \text{id} - T$ we denote its (2α -Lipschitz continuous) fixed-point residual.

We also stick to this notation, so that, whenever mentioned, T , R , and α are as in Assumption I. Our goal is to find a fixed point of T , or, equivalently, a zero of R :

$$\text{find } x_\star \in \mathbf{fix} T = \mathbf{zer} R. \quad (3)$$

In this section we introduce Algorithm 1, an abstract procedure to solve problem (3). The scheme is not implementable in and of itself, as it gives no hint as to how to compute each of the iterates, but it rather serves as a comprehensive ground framework for a class of algorithms with global convergence guarantees. In Section VI we will derive the *Super-Mann scheme*, an implementable instance which also enjoys appealing asymptotic properties.

The general framework prescribes three kinds of updates.

K_0) **Blind updates.** Inspired from [17], whenever the residual $\|Rx_k\|$ at iteration k has *sufficiently* decreased with respect to past iterates we allow for an *uncontrolled* update. For an efficient implementation such guess should be somehow reasonable and not completely a “blind” guess; however, for the sake of global convergence the proposed scheme is robust to any choice.

K_1) **Educated updates.** To encourage favorable updates, similarly to what has been proposed in [9, §5.3.1] and [8, §8.3.2] an *educated guess* x_{k+1} is accepted whenever the candidate residual is *sufficiently* smaller than the current.

K_2) **Safeguard (Fejérian) updates.** This last kind of updates is similar to K_1 as it is also based on the goodness of x_{k+1} with respect to x_k . The difference is that instead of checking the residual, what needs to be *sufficiently* decreased is the distance from each point in $\mathbf{fix} T$. This is meant in a Fejérian fashion as in Definition III.2.

Blind K_0 - and *educated* K_1 -updates are somehow complementary: the former is enabled when enough progress has been made in the past, whereas the latter when the candidate update yields a sufficient improvement. *Progress* and *improvement* are meant in terms of a linear decrease of (the norm of) the residual; at iteration k , K_0 is enabled if $\|Rx_k\| \leq c_0 \|Rx_{\bar{k}}\|$, where $c_0 \in [0, 1)$ is a user-defined constant and \bar{k} is the last blind iteration before k ; K_1 is enabled if $\|Rx_{k+1}\| \leq c_1 \|Rx_k\|$ where $c_1 \in [0, 1)$ is another user-defined constant and x_{k+1} is the candidate next iterate. To ensure global convergence, *educated* updates are authorized only if the current residual $\|Rx_k\|$ is not larger than $\|Rx_{\bar{k}+1}\|$ (up to a linearly decreasing error $q^{\bar{k}}$); here \bar{k} denotes the last K_1 -update before k .

While *blind* K_0 - and *educated* K_1 -updates are in charge of the asymptotic behavior, what makes the algorithm convergent are *safeguard* K_2 -iterations.

A. Global weak convergence

To establish a notation, we partition the set of iteration indices $K \subseteq \mathbb{N}$ as $K_0 \cup K_1 \cup K_2$. Namely, relative to Algorithm 1, K_0 , K_1 and K_2 denote the sets of indices k passing the test at steps **2**, **3(a)** and **3(b)**, respectively. Furthermore, we index the sets K_0 and K_1 of *blind* and *educated* updates as

$$K_0 = \{k_1, k_2, \dots\}, \quad K_1 = \{k'_1, k'_2, \dots\}. \quad (5)$$

To rule out trivialities, throughout the paper we work under the assumption that a solution is not found in a finite number

Algorithm 1 *General framework* for finding a fixed point of the α -averaged operator T with residual $R = \text{id} - T$

REQUIRE $x_0 \in \mathcal{H}$, $c_0, c_1, q \in [0, 1)$, $\sigma > 0$

INITIALIZE $\eta_0 = r_{\text{safe}} = \|Rx_0\|$, $k = 0$

1. IF $Rx_k = 0$, THEN STOP.

2. IF $\|Rx_k\| \leq c_0\eta_k$, THEN set $\eta_{k+1} = \|Rx_k\|$, proceed with a *blind update* x_{k+1} and go to step 4.

3. Set $\eta_{k+1} = \eta_k$ and select x_{k+1} such that

3(a) EITHER the *safe condition* $\|Rx_k\| \leq r_{\text{safe}}$ holds, and x_{k+1} is *educated*:

$$\|Rx_{k+1}\| \leq c_1\|Rx_k\|$$

in which case update $r_{\text{safe}} = \|Rx_{k+1}\| + q^k$;

3(b) OR it is *Fejérian* with respect to $\text{fix } T$:

$$\|x_{k+1} - z\|^2 \leq \|x_k - z\|^2 - \sigma\|Rx_k\|^2 \quad \forall z \in \text{fix } T. \quad (4)$$

4. Set $k \leftarrow k + 1$ and go to step 1.

of steps, so that the residual of each iterate is always nonzero. As long as it is well defined, the algorithm therefore produces an infinite number of iterates.

Theorem IV.1 (Global convergence of the general framework Algorithm 1). *Consider the iterates generated by Algorithm 1 and suppose that for all k it is always possible to find a point x_{k+1} complying with the requirements of either step 2, 3(a) or 3(b), and further satisfying*

$$\|x_{k+1} - x_k\| \leq D\|Rx_k\| \quad \forall k \in K_0 \cup K_1 \quad (6)$$

for some constant $D \geq 0$. Then,

(i) $(x_k)_{k \in \mathbb{N}}$ is quasi-Fejér monotone with respect to $\text{fix } T$;

(ii) $Rx_k \rightarrow 0$ with $(\|Rx_k\|)_{k \in \mathbb{N}} \in \ell^2$;

(iii) $(x_k)_{k \in \mathbb{N}}$ converges weakly to a point $x_* \in \text{fix } T$;

(iv) if $c_0 > 0$ the number of blind updates at step 2 is infinite.

Proof. See Appendix A. \square

B. Local linear convergence

More can be said about the convergence rates if the mapping R possesses *metric subregularity*. Differently from (bounded) linear regularity [18], metric subregularity is a local property and as such it is more general. For a (possibly multivalued) operator R , metric subregularity at \bar{x} is equivalent to calmness of R^{-1} at $R\bar{x}$ [19, Thm 3.2], and is a weaker condition than metric regularity and Aubin property. We refer the reader to [20, §9] for an extensive discussion.

Definition IV.2 (Metric subregularity at zeros). *Let $R : \mathcal{H} \rightarrow \mathcal{H}$ and $\bar{x} \in \text{zer } R$. R is metrically subregular at \bar{x} if there exist $\varepsilon, \gamma > 0$ such that*

$$\text{dist}(x, \text{zer } R) \leq \gamma\|Rx\| \quad \forall x \in \mathbf{B}(\bar{x}; \varepsilon). \quad (7)$$

γ and ε are (one) modulus and (one) radius of subregularity of R at \bar{x} , respectively.

In finite-dimensional spaces, if R is differentiable at $\bar{x} \in \text{zer } R$ and \bar{x} is isolated in $\text{zer } R$ (e.g., if it is the unique zero), then metric subregularity is equivalent to nonsingularity of $JR\bar{x}$. Metric subregularity is however a much weaker property than nonsingularity of the Jacobian, firstly because it does not assume differentiability, and secondly because it can cope

with ‘wide’ regions of zeros; for instance, any piecewise linear mapping is globally metrically subregular [21].

If the residual $R = \text{id} - T$ of the α -averaged operator T is metrically subregular at $\bar{x} \in \text{zer } R = \text{fix } T$ with modulus γ and radius ε , then

$$\frac{1}{\gamma} \text{dist}(x, \text{fix } T) \leq \|Rx\| \leq 2\alpha \text{dist}(x, \text{fix } T) \quad (8)$$

for all $x \in \mathbf{B}(\bar{x}; \varepsilon)$. Consequently, if $\|Rx_k\| \rightarrow 0$ for some sequence $(x_k)_{k \in \mathbb{N}} \subset \mathcal{H}$, so does $\text{dist}(x_k, \text{fix } T)$ with the same asymptotic rate of convergence, and viceversa. Metric subregularity is the key property under which the residual in the classical KM scheme achieves linear convergence; in the next result we show that this asymptotic behavior is preserved in the general framework of Algorithm 1.

Theorem IV.3 (Linear convergence of the general framework Algorithm 1). *Suppose that the hypotheses of Theorem IV.1 hold, and suppose further that $(x_k)_{k \in \mathbb{N}}$ converges strongly to a point x_* (this being true if \mathcal{H} is finite dimensional) at which R is metrically subregular.*

Then, $(x_k)_{k \in \mathbb{N}}$ and $(Rx_k)_{k \in \mathbb{N}}$ are R -linearly convergent.

Proof. See Appendix A. \square

C. Main idea

Being interested in solving the nonlinear equation (3), one could think of implementing one of the many existing *fast* methods for nonlinear equations that achieve fast asymptotic rates, such as Newton-type schemes. At each iteration, such schemes compute an update direction d_k and prescribe steps of the form $x_{k+1} = x_k + \tau_k d_k$, where $\tau_k > 0$ is a stepsize that needs to be sufficiently small in order for the method to enjoy global convergence; on the other hand, fast asymptotic rates are ensured if $\tau_k = 1$ is eventually always accepted. The stepsize is a crucial feature of fast methods, and a feasible τ_k is usually backtracked with a linesearch on a smooth merit function. Unfortunately, in meaningful applications of the problem at hand arising from fixed-point theory the residual mapping R is nonsmooth, and the typical merit function $x \mapsto \|Rx\|^2$ does not meet the necessary smoothness requirement.

What we propose in this paper is a hybrid scheme that allows for the employment of any (fast) method for solving nonlinear equations, with global convergence guarantees that

do not require smoothness, and which is based only on the nonexpansiveness of T . Once *fast* directions d_k are selected, Algorithm 1 can be specialized as follows:

- 1) *blind* updates as in step 2 shall be of the form $x_{k+1} = x_k + d_k$;
- 2) *educated* updates as in step 3(a) shall be of the form $x_{k+1} = x_k + \tau_k d_k$, with τ_k small enough so as to ensure the acceptance condition $\|Rx_{k+1}\| \leq c_1 \|Rx_k\|$;
- 3) *safeguard* updates as in step 3(b) shall be employed as *last resort* both for globalization purposes and for well definedness of the scheme.

Ideally, the scheme should eventually reduce to the local scheme $x_{k+1} = x_k + d_k$ when *good* directions d_k are used.

In Section V we address the problem of providing explicit *safeguard* updates that comply with the quasi-Fejér monotonicity requirement of step 3(b). Because of the arbitrariness of the other two updates, once we succeed in this task Algorithm 1 will be of practical implementation. In Section VI we will then discuss specific K_0 - and K_1 -updates to be used at steps 2 and 3(a) that ensure global and fast convergence, yet maintaining the simplicity of fixed-point iterations of T (evaluations of T and direct linear algebra).

V. GENERALIZED MANN ITERATIONS

A. The classical Krasnosel'skiĭ-Mann scheme

Starting from a point $x_0 \in \mathcal{H}$, the classical Krasnosel'skiĭ-Mann scheme (KM) performs the following updates

$$x_{k+1} = T_{\lambda_k} x_k = (1 - \lambda_k)x_k + \lambda_k T x_k \quad (9)$$

and converges weakly to a fixed point of T provided that $\lambda_k \in [0, 1/\alpha]$ and $(\lambda_k(1/\alpha - \lambda_k))_{k \in \mathbb{N}} \notin \ell^1$ [14, Thm. 5.14]. The key property of KM iterations is Fejér monotonicity:

$$\|x_{k+1} - z\|^2 \leq \|x_k - z\|^2 - \lambda_k(1/\alpha - \lambda_k)\|Rx_k\|^2 \quad \forall z \in \mathbf{fix} T.$$

In particular, in Algorithm 1 KM iterations can be used as *safeguard* updates at step 3(b). The drawback of such a selection is that it completely discards the hypothetical fast update direction d_k that *blind* and *educated* updates try to enforce. This is particularly penalizing when the local method for computing the directions d_k is a *quasi-Newton* scheme; such methods are indeed very sensitive to past iterations, and discarding directions is neither theoretically sound nor beneficial in practice.

In this section we provide alternative *safeguard* updates that while ensuring the desirable Fejér monotonicity are also amenable to taking into account arbitrary directions. The key idea lies in interpreting KM iterations as projections onto suitable half-spaces (see Fig. 2), and then exploiting known properties of projections. These facts are shown in the next result. To this end, let us remark that the projection \mathbf{P}_C onto a nonempty closed and convex set C is FNE [14, Prop. 4.8], and that consequently its λ -averaging $\mathbf{P}_{C,\lambda}$ is $\lambda/2$ -averaged for any $\lambda \in (0, 2]$, as it follows from (2).

Proposition V.1 (KM iterations as projections). *For $x \in \mathcal{H}$, define*

$$C_x = C_x^{T,\alpha} := \{z \in \mathcal{H} \mid \|Rx\|^2 - 2\alpha \langle Rx, x - z \rangle \leq 0\}. \quad (10)$$

Then,

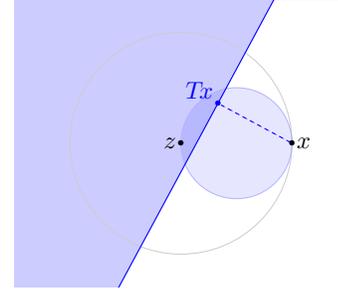


Figure 2: Mann iteration of a FNE operator T as projection on C_x (the blue half-space, as defined in (10) for $\alpha = 1/2$). The outer circle is the set of all possible images of a nonexpansive operator, given that z is a fixed point. The inner circle corresponds to the possible images of firmly nonexpansive operators. Notice that C_x separates x from z as long as Tx is contained in the small circle, which characterizes firm nonexpansiveness.

- (i) $x \in C_x$ iff $x \in \mathbf{fix} T$;
- (ii) $\mathbf{fix} T = \bigcap_{x \in \mathcal{H}} C_x$;
- (iii) for any $\lambda \in [0, 1/\alpha]$ it holds that $T_{\lambda} x = \mathbf{P}_{C_x, 2\alpha\lambda} x = (1 - 2\alpha\lambda)x + 2\alpha\lambda \mathbf{P}_{C_x} x$.

Proof. The set C_x can be equivalently expressed as

$$C_x = \{z \in \mathcal{H} \mid \langle x - T_{1/2\alpha} x, z - T_{1/2\alpha} x \rangle \leq 0\}.$$

V.1(i) is of immediate verification, and V.1(ii) then follows from [14, Cor. 4.16] combined with (2).

We now show V.1(iii). If $Rx = 0$, then $x \in \mathbf{fix} T$ and $C_x = \mathcal{H}$, and the claim is trivial. Otherwise, notice that

$$C_x = \left\{z \in \mathcal{H} \mid \langle Rx, z \rangle \leq \langle Rx, x - \frac{1}{2\alpha} Rx \rangle\right\}, \quad (11)$$

and the claim can be readily verified using the formula for the projection on a halfspace $H_{v,\beta} := \{z \in \mathcal{H} \mid \langle v, z \rangle \leq \beta\}$, namely

$$\mathbf{P}_{H_{v,\beta}} x = x - \frac{[\langle v, x \rangle - \beta]_+}{\|v\|^2} v, \quad (12)$$

defined for $v \in \mathcal{H} \setminus \{0\}$ and $\beta \in \mathbb{R}$ [14, Ex. 28.16(iii)]. \square

B. Generalized Mann projections

Though particularly attractive for its simplicity and global convergence properties, the KM scheme (9) finds its main drawback in its convergence rate, being it Q -linear at best and highly sensitive to ill conditioning of the problem. In response to these issues, Algorithm 1 allows for the integration of fast local methods still ensuring global convergence properties. The efficiency of the resulting scheme, which will be proven later on, is based on an ad hoc selection of *safeguard* updates for step 3(b) which is based on the following generalization of Proposition V.1.

Proposition V.2. *Suppose that $x, w \in \mathcal{H}$ are such that*

$$\rho := \|Rw\|^2 - 2\alpha \langle Rw, w - x \rangle > 0. \quad (13)$$

For $\lambda \in [0, 1/\alpha]$ let

$$x^+ := x - \lambda \frac{\rho}{\|Rw\|^2} Rw. \quad (14)$$

Then, the following hold:

(i) $x^+ = \mathbf{P}_{C_w, 2\alpha\lambda} x$ where $C_w = C_w^{T,\alpha}$ as in (10);

(ii) $\|x^+ - z\|^2 \leq \|x - z\|^2 - \lambda(1/\alpha - \lambda) \frac{\rho^2}{\|Rw\|^2} \quad \forall z \in \mathbf{fix} T$.

Proof. V.2(i) easily follows from (11) and (12), since by condition (13) the positive part in the formula may be omitted. In turn, V.2(ii) follows from [14, Prop. 4.25(iii)] by observing that $\Pi_{C_w, 2\alpha\lambda}$ is $\alpha\lambda$ -averaged due to [14, Prop.s 4.8] and (2), and that $\mathbf{fix} T \subseteq C_w$ as shown in Prop. V.1(ii). \square

Notice that condition (13) is equivalent to $x \notin C_w$. Therefore, Proposition V.2(ii) states that whenever a point x lies outside the half-space C_w for some $w \in \mathcal{H}$, since $\mathbf{fix} T \subseteq C_w$ (cf. Prop. V.1) the projection onto C_w moves x closer to $\mathbf{fix} T$. This means that after moving from x along a candidate direction d to the point $w = x + d$, even though w might be farther from $\mathbf{fix} T$ the point $x^+ = \Pi_{C_w} x$ is not. We may then use this projection as a safeguard step to prevent from diverging from the set of fixed points. Based on this, we define a *generalized KM update along a direction d* .

Definition V.3 (GKM update). A generalized KM update (GKM) at x along d for the α -averaged operator $T : \mathcal{H} \rightarrow \mathcal{H}$ with relaxation $\lambda \in [0, 1/\alpha]$ is

$$x^+ := \begin{cases} x & \text{if } w \in \mathbf{fix} T \\ x - \lambda \frac{\langle \rho \rangle_+}{\|Rw\|^2} Rw & \text{otherwise,} \end{cases}$$

where $w = x + d$ and $\rho := \|Rw\|^2 - 2\alpha\langle Rw, w - x \rangle$. In particular, $d = 0$ yields the classical KM update $x^+ = T_\lambda x$.

C. Linesearch for GKM

It is evident from Definition V.3 that a GKM update trivializes to $x^+ = x$ if either $w \in \mathbf{fix} T$ or $\rho \leq 0$. Having $w \in \mathbf{fix} T$ corresponds to having found a solution to problem (3), and the case deserves no further investigation. In this section we address the remaining case $\rho \leq 0$, showing how it can be avoided by simply introducing a suitable linesearch. In order to recover the same global convergence properties of the classical KM scheme we need something more than simply imposing $\rho > 0$. The next result addresses this requirement, showing further that it is achieved for any direction d by sufficiently small stepsizes.

Theorem V.4. Let $x, d \in \mathcal{H}$ and $\sigma \in [0, 1)$ be fixed, and consider

$$\bar{\tau} := \begin{cases} 1 & \text{if } d = 0 \\ \frac{1-\sigma}{4\alpha} \frac{\|Rx\|}{\|d\|} & \text{otherwise.} \end{cases}$$

Then, for all $\tau \in (0, \bar{\tau}]$ the point $w = x + \tau d$ satisfies

$$\rho := \|Rw\|^2 - 2\alpha\langle Rw, w - x \rangle \geq \sigma\|Rw\|\|Rx\|. \quad (15)$$

Proof. Let a constant $c \geq 0$ to be determined be such that

$$\tau\|d\| = \|w - x\| \leq c\|Rx\|.$$

Observe that $\rho = 4\alpha^2\langle w - T_{1/2\alpha}w, x - T_{1/2\alpha}w \rangle$, and recall from (1) and (2) that $T_{1/2\alpha}$ is FNE with residual $\text{id} - T_{1/2\alpha} = \frac{1}{2\alpha}R$. Then,

$$\rho = 4\alpha^2(\|w - T_{1/2\alpha}w\|^2 + \langle w - T_{1/2\alpha}w, x - w \rangle)$$

using Cauchy-Schwartz inequality,

$$\geq 4\alpha^2\|w - T_{1/2\alpha}w\|(\|w - T_{1/2\alpha}w\| - \|x - w\|)$$

the bound on $\|x - w\|$,

$$\geq 2\alpha\|Rw\|(\|w - T_{1/2\alpha}w\| - 2\alpha c\|x - T_{1/2\alpha}x\|)$$

the (reverse) triangular inequality,

$$\geq 2\alpha\|Rw\| \left((1 - 2\alpha c)\|x - T_{1/2\alpha}x\| - \|(\text{id} - T_{1/2\alpha})w - (\text{id} - T_{1/2\alpha})x\| \right)$$

the nonexpansiveness of $\text{id} - T_{1/2\alpha}$

$$\geq 2\alpha\|Rw\| \left(\frac{1-2\alpha c}{2\alpha}\|Rx\| - \|w - x\| \right)$$

and again the bound on $\|w - x\|$,

$$\geq (1 - 4\alpha c)\|Rw\|\|Rx\|$$

equating $\sigma = 1 - 4\alpha c$ the assert follows. \square

Notice that if $d = 0$, then $\rho = \|Rx\|^2 \geq \sigma\|Rx\|^2$ for any $\sigma \in [0, 1)$, and therefore the linesearch condition (15) is always satisfied; in particular, the classical KM step $x^+ = Tx$ is always accepted regardless of the value of σ .

Let us now observe how a GKM projection extends the classical KM depicted in Figure 2 and how the linesearch works. In the following we use the notation of Theorem V.4, and for the sake of simplicity we consider $\sigma = 0$ in (15) and a FNE operator T . Suppose that the fixed point z and the points x, Tx , and w are as in Figure 3a; due to firm nonexpansiveness, the image Tw of w is somewhere in the intersection of the orange circles. We want to avoid the unfavorable situation depicted in Figure 3b, where the couple (w, Tw) generates a halfspace C_w that contains x , i.e., such that $\rho \leq 0$: in fact, with simple algebra it can be seen that $\rho \leq 0$ iff Tw belongs to the dashed circle of Figure 3b:

$$B_{x,w} := \{\bar{w} \mid \langle w - \bar{w}, x - \bar{w} \rangle \leq 0\}. \quad (16)$$

Since the dashed orange circle (in which Tw must lie) is simply the translation by a vector $Tx - x$ of $B_{x,w}$, both having diameter $\tau\|d\|$, for sufficiently small τ the two have empty intersection, meaning that $\rho > 0$ regardless of where Tw is.

VI. THE SUPERMANN SCHEME

In this section we introduce the *SuperMann scheme* (Alg. 2), a special instance of the general framework of Algorithm 1 that employs GKM updates as safeguard K_2 -steps. While the global worst-case convergence properties of *SuperMann* are the same as for the classical KM scheme, its asymptotic behavior is determined by how *blind* K_0 - and *educated* K_1 -updates are selected. In Section VI-B we will characterize the “quality” of update directions and the mild requirements under which superlinear convergence rates are attained; in particular, Section VI-C is dedicated to the analysis of quasi-Newton Broyden’s directions.

The scheme follows the same philosophy of the general abstract framework. The main idea is globalizing a local method for solving the monotone equation $Rx = 0$, in such a way that when the iterates get close enough to a solution the fast convergence of the local method is automatically triggered. Approaching a solution is possible thanks to the generalized KM updates (step 5(b)), provided enough backtracking is performed, as ensured by Prop. V.2(ii) and Thm. V.4. When a basin of fast (i.e., superlinear) attraction for the local method is reached, the (norm of) Rx will decrease more than linearly, and the condition triggering the *educated* updates of step 5(a)

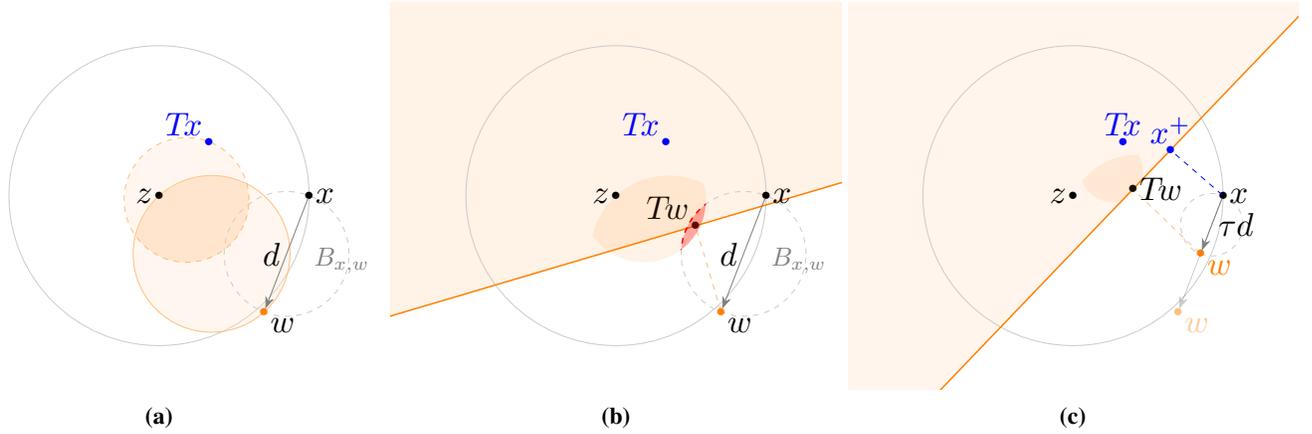


Figure 3: SuperMann iteration of a FNE operator T as projection on C_w .
 (a) the darker orange region represents the area in which Tw must lie given the points x , Tx and the fixed point z as prescribed by firm nonexpansiveness of T .
 (b) if Tw lies (also) in the ball $B_{x,w}$ as in (16), then the half-space C_w (shaded in orange) separates x from w , which is to be avoided.
 (c) when w is close enough to x the feasible region for Tw has empty intersection with $B_{x,w}$ and C_w does not contain x .

(which is checked first) will be verified without performing any backtracking.

To discuss its global and local convergence properties we stick to the same notation of the general framework of Algorithm 1, denoting the sets of *blind*, *educated*, and *safeguard* updates as K_0 , K_1 and K_2 , respectively.

A. Global and linear convergence

To comply with (6), we impose the following requirement on the magnitude of the directions (see also Rem. VI.9).

Assumption II. *There exists a constant $D \geq 0$ such that the directions $(d_k)_{k \in \mathbb{N}}$ in the SuperMann scheme (Alg. 2) satisfy*

$$\|d_k\| \leq D\|Rx_k\| \quad \forall k \in \mathbb{N}. \quad (17)$$

Theorem VI.1 (Global and linear convergence of the SuperMann scheme). *Consider the iterates generated by the SuperMann scheme (Alg. 2) with $(d_k)_{k \in \mathbb{N}}$ selected so as to satisfy Assumption II. Then,*

- (i) $(x_k)_{k \in \mathbb{N}}$ is quasi-Fejér monotone with respect to $\mathbf{fix} T$;
- (ii) $\tau_k = 1$ if $d_k = 0$, and $\tau_k \geq \min\{\beta^{\frac{1-\sigma}{4\alpha D}}, 1\}$ otherwise.
- (iii) $Rx_k \rightarrow 0$ with $(\|Rx_k\|)_{k \in \mathbb{N}} \in \ell^2$;
- (iv) $(x_k)_{k \in \mathbb{N}}$ converges weakly to a point $x_\star \in \mathbf{fix} T$;
- (v) if $c_0 > 0$ the number of blind updates at step 3 is infinite. Moreover, if $(x_k)_{k \in \mathbb{N}}$ converges strongly to a point x_\star (this being true if \mathcal{H} is finite dimensional) at which R is metrically subregular, then
- (vi) $(x_k)_{k \in \mathbb{N}}$ and $(Rx_k)_{k \in \mathbb{N}}$ are R -linearly convergent.

Proof. See Appendix B. □

B. Superlinear convergence

Though global convergence of the SuperMann scheme is independent of the choice of the directions d_k , its performance and tail convergence surely does. We characterize the *quality* of the directions d_k in terms of the following definition.

Definition VI.2 (Superlinear directions for the SuperMann scheme). *Relative to the sequence $(x_k)_{k \in \mathbb{N}}$ generated by the SuperMann scheme, we say that $(d_k)_{k \in \mathbb{N}} \subset \mathcal{H}$ are superlinear directions if the following limit holds*

$$\lim_{k \rightarrow \infty} \frac{\|R(x_k + d_k)\|}{\|Rx_k\|} = 0.$$

Remark VI.3. Definition VI.2 makes no mention of a limit point x_\star of the sequence $(x_k)_{k \in \mathbb{N}}$, differently from the definition in [8] which instead requires $\frac{\|x_k + d_k - x_\star\|}{\|x_k - x_\star\|}$ to be vanishing with no mention of R . Due to 2α -Lipschitz continuity of R , whenever the directions d_k are bounded as in (17) we have

$$\frac{\|R(x_k + d_k)\|}{\|Rx_k\|} \leq 2\alpha D \frac{\|x_k + d_k - x_\star\|}{\|d_k\|}.$$

Invoking [8, Lem. 7.5.7] it follows that Definition VI.2 is implied by the one in [8] and is therefore more general. □

Theorem VI.4. *Consider the iterates generated by the SuperMann scheme (Alg. 2) with either $c_0 > 0$ or $c_1 > 0$, and with $(d_k)_{k \in \mathbb{N}}$ being superlinear directions as in Definition VI.2. Then,*

- (i) eventually, stepsize $\tau_k = 1$ is always accepted and safeguard updates K_2 are deactivated (i.e., the scheme reduces to the local method $x_{k+1} = x_k + d_k$);
- (ii) $(Rx_k)_{k \in \mathbb{N}}$ converges Q -superlinearly;
- (iii) if the directions d_k satisfy Assumption II, then $(x_k)_{k \in \mathbb{N}}$ converges R -superlinearly;
- (iv) if $c_0 > 0$, then the complement of K_0 is finite.

Proof. See Appendix B. □

Theorem VI.4 shows that when the directions d_k are *good*, then eventually the SuperMann scheme reduces to the local method $x_{k+1} = x_k + d_k$ and consequently inherits its local convergence properties. The following result specializes to the choice of semismooth Newton directions.

Corollary VI.5 (Superlinear convergence for semismooth Newton directions). *Suppose that \mathcal{H} is finite dimensional, and*

Algorithm 2 *SuperMann* scheme for solving (3), given an α -averaged operator T with residual $R = \text{id} - T$

REQUIRE $x_0 \in \mathcal{H}$, $c_0, c_1, q \in [0, 1)$, $\beta, \sigma \in (0, 1)$, $\lambda \in (0, 1/\alpha)$.

INITIALIZE $\eta_0 = r_{\text{safe}} = \|Rx_0\|$, $k = 0$

1. IF $Rx_k = 0$, THEN STOP.

2. Choose an update direction $d_k \in \mathcal{H}$

3. (**K₀**) IF $\|Rx_k\| \leq c_0\eta_k$, THEN set $\eta_{k+1} = \|Rx_k\|$, proceed with a *blind update* $x_{k+1} = w_k := x_k + d_k$ and go to step 6.

4. Set $\eta_{k+1} = \eta_k$ and $\tau_k = 1$.

5. Let $w_k = x_k + \tau_k d_k$.

5(a) (**K₁**) IF the *safe condition* $\|Rx_k\| \leq r_{\text{safe}}$ holds and w_k is *educated*:

$$\|Rw_k\| \leq c_1 \|Rx_k\|$$

THEN set $x_{k+1} = w_k$, update $r_{\text{safe}} = \|Rw_k\| + q^k$, and go to step 6.

5(b) (**K₂**) IF $\rho_k := \|Rw_k\|^2 - 2\alpha \langle Rw_k, w_k - x_k \rangle \geq \sigma \|Rw_k\| \|Rx_k\|$

THEN set

$$x_{k+1} = x_k - \lambda \frac{\rho_k}{\|Rw_k\|^2} Rw_k$$

OTHERWISE set $\tau_k \leftarrow \beta \tau_k$ and go to step 5.

6. Set $k \leftarrow k + 1$ and go to step 1.

that R is semismooth. Consider the iterates generated by the *SuperMann* scheme (Alg. 2) with either $c_0 > 0$ or $c_1 > 0$ and directions d_k chosen as solutions of

$$(G_k + \mu_k \text{id})d_k = -Rx_k \quad \text{for some } G_k \in \partial Rx_k \quad (18)$$

where ∂R denotes the Clarke generalized Jacobian of R and $0 \leq \mu_k \rightarrow 0$. Suppose that the sequence $(x_k)_{k \in \mathbb{N}}$ converges to a point x_\star at which all the elements in ∂R are nonsingular.

Then, $(d_k)_{k \in \mathbb{N}}$ are superlinear directions as in Definition VI.2, and in particular all the claims of Theorem VI.4 hold.

Proof. Any $G_k \in \partial R$ is positive semidefinite due to the monotonicity of R , and therefore d_k as in (18) is well defined for any $\mu_k > 0$. The bound (17) holds due to [8, Thm. 7.5.2]. Moreover,

$$\frac{\|Rx_k + G_k d_k\|}{\|d_k\|} = \mu_k \rightarrow 0$$

as $k \rightarrow \infty$, and the proof follows by invoking [8, Thm. 7.5.8(a)] and Rem. VI.3. \square

Notice that since $\partial R = \text{id} - \partial T$, nonsingularity of the elements in $\partial R(x_\star)$ is equivalent to having $\|G\| < 1$ for all $G \in \partial T(x_\star)$, i.e., that T is a local contraction around x_\star .

Despite the favorable properties of semismooth Newton methods, in this paper we are oriented towards choices of directions that (1) are defined for any nonexpansive mapping, regardless of the (generalized) first-order properties, and that (2) require exactly the same oracle information as the original KM scheme. This motivates the investigation of quasi-Newton directions, whose superlinear behavior is based on the classical Dennis-Moré criterion, which we provide next. We first recall the notions of *semi-* and *strict differentiability*.

Definition VI.6. We say that $R : \mathcal{H} \rightarrow \mathcal{H}$ is

(i) strictly differentiable at \bar{x} if it is differentiable there with $JR(\bar{x})$ satisfying

$$\lim_{\substack{(y,z) \rightarrow (\bar{x}, \bar{x}) \\ y \neq z}} \frac{\|Ry - Rx - JR(\bar{x})(y - x)\|}{\|y - x\|} = 0; \quad (19)$$

(ii) semidifferentiable at \bar{x} if there exists a continuous and positively homogeneous function $DR(\bar{x}) : \mathcal{H} \rightarrow \mathcal{H}$, called the semiderivative of R at \bar{x} , such that

$$Rx = R\bar{x} + DR(\bar{x})[x - \bar{x}] + o(\|x - \bar{x}\|);$$

(iii) calmly semidifferentiable at \bar{x} if there exists a neighborhood $U_{\bar{x}}$ of \bar{x} in which R is semidifferentiable and such that for all $w \in \mathcal{H}$ with $\|w\| = 1$ the function $U_{\bar{x}} \ni x \mapsto DR(x)[w]$ is Lipschitz continuous at \bar{x} .

There is a slight ambiguity in the literature, as *strict differentiability* is sometimes referred to rather as *strong differentiability* [22], [23]. We choose to stick the proposed terminology, following [20]. Semidifferentiability is clearly a milder property than differentiability in that the mapping $DR(\bar{x})$ need not be linear. More precisely, since the residual R of a nonexpansive operator is (globally) Lipschitz continuous, then semidifferentiability is equivalent to directional differentiability [8, Prop. 3.1.3] and the semiderivative is sometimes called *B-derivative* [22], [8]. The three concepts in Definition VI.6 are related as (iii) \Rightarrow (i) \Rightarrow (ii) [23, Thm. 2] and neither requires the existence of the (classical) Jacobian around \bar{x} .

Theorem VI.7 (Dennis-Moré criterion for superlinear convergence). Consider the iterates generated by the *SuperMann* scheme (Alg. 2) and suppose that $(x_k)_{k \in \mathbb{N}}$ converges strongly to a point x_\star at which R is strictly differentiable. Suppose further that the update directions $(d_k)_{k \in \mathbb{N}}$ satisfy Assumption II and the Dennis-Moré condition

$$\lim_{k \rightarrow \infty} \frac{\|Rx_k + JR(x_\star)d_k\|}{\|d_k\|} = 0. \quad (20)$$

Then, the directions d_k are superlinear as in Definition VI.2. In particular, all the claims of Theorem VI.4 hold.

Proof.

$$\begin{aligned} 0 &= \lim_{k \rightarrow \infty}^{(20)} \frac{\|Rx_k + JR(x_\star)d_k + R(x_k + d_k) - R(x_k + d_k)\|}{\|d_k\|} \\ &= \lim_{k \rightarrow \infty} \frac{\|R(x_k + d_k)\|}{\|d_k\|} \stackrel{(17)}{\geq} \frac{1}{D} \lim_{k \rightarrow \infty} \frac{\|R(x_k + d_k)\|}{\|Rx_k\|}, \end{aligned}$$

where in the second equality we used strict differentiability of R at x_* . \square

C. A modified Broyden's direction scheme

In practical application the Hilbert space \mathcal{H} is finite dimensional, and consequently it can be identified with \mathbb{R}^n . Then, the computation of quasi-Newton directions d_k in the *SuperMann scheme* amounts to selecting

$$d_k = -B_k^{-1}Rx_k, \quad (21a)$$

where $B_k \in \mathbb{R}^{n \times n}$ are recursively defined by low-rank updates satisfying a *secant condition*, starting from an invertible matrix B_0 . The most popular quasi-Newton scheme is the 2-rank BFGS formula, which also enforces symmetricity. As such, BFGS is well performing only when the Jacobian at the solution JRx_* possesses this property, a requirement that is not met by the residual R of generic nonexpansive mappings.

For this reason we consider Broyden's method as a universal alternative. We adopt Powell's modification [10] to enforce nonsingularity and make (21a) well defined: for a fixed parameter $\bar{\vartheta} \in (0, 1)$, matrices B_k are recursively defined as

$$B_{k+1} = B_k + \frac{1}{\|s_k\|_2^2} (\tilde{y}_k - B_k s_k) s_k^\top, \quad (21b)$$

where for $\gamma_k := \frac{\langle B_k^{-1} y_k, s_k \rangle_2}{\|s_k\|_2^2}$ we have defined

$$\begin{cases} s_k = w_k - x_k \\ y_k = R w_k - R x_k \\ \tilde{y}_k = (1 - \vartheta_k) B_k s_k + \vartheta_k y_k \end{cases} \quad \vartheta_k := \begin{cases} 1 & \text{if } |\gamma_k| \geq \bar{\vartheta} \\ \frac{1 - \text{sgn}(\gamma_k) \bar{\vartheta}}{1 - \gamma_k} & \text{if } |\gamma_k| < \bar{\vartheta} \end{cases} \quad (21c)$$

with the convention $\text{sgn } 0 = 1$. Letting $H_k := B_k^{-1}$ and using the Sherman-Morrison identity, the inverse of B_k is given by

$$H_{k+1} = H_k + \frac{1}{\langle H_k \tilde{y}_k, s_k \rangle_2} (s_k - H_k \tilde{y}_k) (s_k^\top H_k). \quad (21d)$$

Consequently, there is no need to compute and store the matrices B_k and we can directly operate with their inverses H_k .

Theorem VI.8 (Superlinear convergence of the *SuperMann scheme* with Broyden's directions). *Suppose that \mathcal{H} is finite dimensional. Consider the sequence $(x_k)_{k \in \mathbb{N}}$ generated by the *SuperMann scheme* (Alg. 2), $(d_k)_{k \in \mathbb{N}}$ being selected with the modified Broyden's scheme (21) for some $\bar{\vartheta} \in (0, 1)$.*

Suppose that $(H_k)_{k \in \mathbb{N}}$ remains bounded, and that R is calmly semidifferentiable and metrically subregular at the limit x_ of $(x_k)_{k \in \mathbb{N}}$. Then, $(d_k)_{k \in \mathbb{N}}$ satisfies the Dennis-Moré condition (20). In particular, all the claims of Theorem VI.7 hold.*

Proof. See Appendix B. \square

Remark VI.9. It follows from Theorem VI.1(iv) that the *SuperMann scheme* is globally convergent as long as $\|d_k\| \leq D \|Rx_k\|$ for some constant D . To enforce it we may select a (large) constant $D > 0$ and as a possible choice truncate $d_k \leftarrow D \frac{\|Rx_k\|}{\|d_k\|} d_k$ whenever d_k does not satisfy (17). \square

Let us observe that in order to achieve superlinear convergence the *SuperMann scheme* does not require nonsingularity of the Jacobian at the solution. This is the standard requirement for asymptotic properties of quasi-Newton schemes, which is needed to show first that the method converges at least linearly. [24] generalizes this property invoking the concepts of (strong) metric (sub)regularity (see also [19] for an

extensive review on these properties). However, if R is strictly differentiable at x_* , then strong subregularity, regularity and strong regularity are equivalent to injectivity, surjectivity and invertibility of $JR(x_*)$, respectively, these conditions being all equivalent for mappings $\mathcal{H} \rightarrow \mathcal{H}$ with \mathcal{H} finite dimensional. In particular, contrary to the *SuperMann scheme* standard approaches require the solution x_* at least to be isolated.

Restarted (modified) Broyden's scheme: Broyden's scheme requires storing and operating with $n \times n$ matrices, where n is the dimension of the optimization variable, and is consequently feasible in practice only for small problems. Alternatively, one can restrict Broyden's update rule (21d) to only the most recent pairs of vectors (s_i, y_i) . As detailed in Algorithm 3, this can be done by keeping track of the last vectors s_i and some auxiliary vectors $\tilde{s}_i = \frac{s_i - H_i \tilde{y}_i}{\langle s_i, H_i \tilde{y}_i \rangle_2}$. These are stored in some buffers S and \tilde{S} , which are initially empty and can contain up to m vectors. The *memory* m is a small integer typically between 3 and 20; when the memory is full, the buffers are emptied and Broyden's scheme is *restarted*. The choice of a *restarted* rather than a *limited-memory* variant obviates the need of a nested for-loop to account for Powell's modification.

Algorithm 3 Restarted Broyden's scheme with memory m using Powell's modification

Input: old buffers S, \tilde{S} ; new pair (s, y) ; current Rx

Output: new buffers S, \tilde{S} ; update direction d

- 1: $d \leftarrow -Rx$, $\tilde{s} \leftarrow y$
- 2: **for** $i = 1 \dots \#S$ **do**
 $\tilde{s} \leftarrow \tilde{s} + \langle s_i, \tilde{s} \rangle_2 \tilde{s}_i$, $d \leftarrow d + \langle s_i, d \rangle_2 \tilde{s}_i$
- 3: **end for**
- 4: compute ϑ as in (21c) with $\gamma = \frac{1}{\|s\|_2^2} \langle \tilde{s}, s \rangle_2$
- 5: $\tilde{s} \leftarrow \frac{\vartheta}{(1 - \vartheta + \vartheta \gamma) \|s\|_2^2} (s - \tilde{s})$, $d \leftarrow d + \langle s, d \rangle_2 \tilde{s}$
- 6: **if** $\#S = m$ **then** $S, \tilde{S} \leftarrow []$ **else** $S \leftarrow [S, s]$, $\tilde{S} \leftarrow [\tilde{S}, \tilde{s}]$

D. Parameters selection in SuperMann

As shown in Theorem VI.4, the *SuperMann scheme* makes sense as long as either $c_0 > 0$ or $c_1 > 0$; indeed, *safeguard* K_2 -steps are only needed for globalization, while it is *blind* K_0 - and *educated* K_1 -steps that exploit the quality of the directions d_k . Evidently, K_1 -updates are more reliable than K_0 -updates in that they take into account the residual of the candidate next point. As such, it is advisable to select c_1 close to 1 and use small values of c_0 if more conservatism and robustness are desired. To further favor K_1 -updates, the parameter q used for updating the safeguard r_{safe} at step **5(a)** may be also chosen very close to 1.

As to *safeguard* K_2 -steps, a small value of σ makes condition (15) easier to satisfy and results in fewer backtrackings; the averaging factor λ may be chosen equal to 1 whenever possible, i.e., if $\alpha \lesssim 1$ (which is the typical case when, e.g., T comes from splitting schemes in convex optimization), or any close value otherwise. In the simulations of Section VII we used $c_0 = c_1 = q = 0.99$, $\sigma = 0.1$, $\lambda = 1$ and $\beta = 1/2$. For a matter of scaling, we multiplied the summable term q^k by $\|Rx^0\|$ in updating the parameter r_{safe} at step **5(a)**. The directions were computed according to the restarted modified Broy-

den's scheme (Alg. 3) with memory $m = 20$ and $\bar{\theta} = 0.2$; we applied the truncation rule as in Remark VI.9 with $D = 10^4$. We also imposed a maximum of 8 backtrackings after which a nominal fixed-point iteration would be executed.

E. Comparisons with other methods

1) *Hybrid global and local phase algorithms: Blind K_0 -updates* in the *SuperMann scheme* are inspired from [17, Alg. 1], and so is the notation $K_0 = \{k_0, k_1, \dots\}$.

Educated K_1 - and safeguard K_2 -updates instead play the role of *inner-* and *outer-phases* in the general algorithmic framework described in [9, §5.3] for finding a zero of a candidate merit function φ (e.g. $\varphi(x) = \frac{1}{2}\|Rx\|^2$ in our case). Differently from [9, Alg. 5.16] where all previous inner-phase iterations are discarded as soon as the required sufficient decrease is not met, the *SuperMann scheme* allows for an alternation of phases that eventually stabilizes on the fast local one, provided the solution is sufficiently regular. Our scheme is more in the flavor of [9, Alg. 5.19], although with less conservative requirements for triggering *inner K_1 -updates* ($\varphi(x_{k+1})$ is here compared with $\varphi(x_k)$, whereas in the cited scheme with the smallest past value).

2) *Inexact Newton methods for monotone equations:* The GKM updates are closely related to the extra-gradient steps described in [25, Alg. 2.1]. This work introduces an inexact Newton algorithm for solving systems of continuous monotone equations $Rx = 0$, where $\text{id} - R$ need not be nonexpansive. At a given point x , first a direction d is computed as (possibly approximate) solution of $Gd = -Rx$, where G is some positive definite matrix; then, an intermediate point $w = x + \tau d$ is retrieved with a linesearch on τ that ensures the condition

$$\|Rw\|^2 - \langle Rw, x - Tw \rangle \leq -\sigma\tau\|d\|^2 \quad (22)$$

for some $\sigma > 0$; here, we defined $T := \text{id} - R$ to highlight the symmetry with (10). Finally, the new iterate is given by $x^+ = \Pi_{H_w} x$, where

$$H_w := \{z \in \mathcal{H} \mid \|Rw\|^2 - \langle Rw, z - Tw \rangle \geq 0\}. \quad (23)$$

Letting C_w be the half-space as in Prop. V.2, so that $x_{\text{GKM}}^+ = \Pi_{C_w} x$ (for simplicity we set $\lambda = 1$), for the half-spaces (23) it holds that

$$\mathbf{zer} R \subseteq C_w \subseteq H_w,$$

the last inclusion holding as equality iff $Rw = 0$. This means that in the GKM scheme, the same w yields an iterate x_{GKM}^+ which is closer to any $z \in \mathbf{zer} R$ with respect to x^+ (cf. Fig. 4). Notice further that the hyperplanes delimiting the two half-spaces are parallel, with $\mathbf{bdry} C_w$ passing by Tw (or $T_{1/2\alpha} w$ for generic α 's) and $\mathbf{bdry} H_w$ by w .

The requirement of positive definiteness of matrix G in defining the update direction d is due to the fact that [25] addresses a broader class of monotone operators; we instead exploited at full the nonexpansiveness of $\text{id} - R$ and as a result have complete freedom in selecting d (Fig. 4a) and better projections (Fig. 4c).

Moreover, it can be easily verified that the proposed extra-gradient step does not extend the classical KM iteration unless T has a very peculiar property, namely that $Rx = \frac{\langle RTx, Rx \rangle}{\|Rx\|^2} RTx$ for every x . (In particular, for such a T necessarily $\|Rx\| =$

$\|RTx\|$ for all x , and consequently there cannot exist $\alpha \in (0, 1)$ for which T is α -averaged).

3) *Line-search for KM:* The recent work [13] proposes an acceleration of the classical KM scheme for finding a fixed point of an α -averaged operator T based on a linesearch on the relaxation parameter. Namely, instead of the *nominal* update $\bar{x} = T_\lambda x$ with $\lambda \in [0, 1/\alpha]$ as in (9), values $\lambda' > 1/\alpha$ are first tested and the update $x^+ = T_{\lambda'} x$ is accepted as long as $\|Rx^+\| \leq c_1 \|R\bar{x}\|$ holds for some constant $c_1 \in (0, 1)$.

In the setting of the *SuperMann scheme*, this corresponds to selecting $d_k = -Rx_k$, discarding *blind* updates (i.e., setting $c_0 = 0$), foretracking *educated* updates and using plain KM iterations as safeguard steps. Convergence can be enhanced and the method is attractive when $T = S_2 \circ S_1$ is the composition of an affine mapping S_1 and a cheap operator S_2 , in which case the linesearch is inexpensive. However, though preserving the same theoretical convergence guarantees of KM (hence of the *SuperMann scheme*), it does not improve its best-case local linear rate.

Although other choices d_k may also be considered, however fast directions such as Newton-type ones would be discarded and replaced by nominal KM updates every time the candidate point $x_k + d_k$ does not meet some requirements. Avoiding this take-it-or-leave-it behavior is exactly the primary goal of GKM iterations, so that candidate good directions are never discarded.

4) *Smooth optimization with envelope functions:* For solving nonsmooth minimization problems in composite form, [26], [27] introduced *forward-backward envelope* (FBE) and *Douglas-Rachford envelope* (DRE) functions. The original nonsmooth problem is recast into the minimization of continuous (possibly continuously differentiable) real-valued exact penalty functions closely related to FBS and DRS, named *envelopes* due to their kinship with the Moreau *envelope* and the proximal point algorithm. This paved the way for the employment of fast methods for smooth unconstrained minimization problems [26], [27], [28], or for globalizing convergence of fast methods for solving nonlinear equations [29], [30]. Although they have the advantage of being suited for nonconvex problems, however their employment is limited to composite operators as described above and they cannot handle, for instance, saddle-point convex-concave optimization problems typically arising from primal-dual splittings such as Vü-Condat [31]. The *SuperMann scheme* instead offers a unifying framework that is based uniquely on evaluations of the nonexpansive mapping T , regardless of their structure.

VII. SIMULATIONS – LINEAR OPTIMAL CONTROL

For matrices A_t and B_t of suitable size, $t = 0, \dots, N-1$, consider a state-input dynamical system

$$x_{t+1} = A_t x_t + B_t u_t, \quad t = 0, \dots, N-1, \quad (24a)$$

where the $x_0 \in \mathbb{R}^{n_s}$ is given, and the next states $x_t \in \mathbb{R}^{n_s}$ are determined by the user-defined inputs $u_\tau \in \mathbb{R}^{n_u}$, $\tau = 0, \dots, t-1$. States $\mathbf{x} = (x_1, \dots, x_N)$ can be expressed in terms of the inputs $\mathbf{u} = (u_0, \dots, u_{N-1})$ through a linear operator $L \in \mathbb{R}^{Nn_s \times Nn_u}$ as

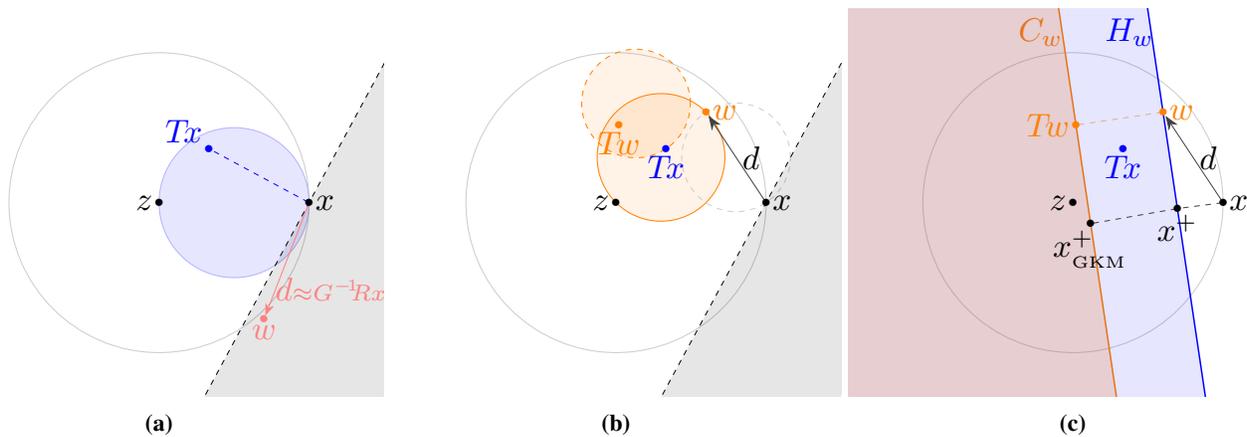


Figure 4: The positive definiteness of G prevents the update directions d in the scheme of [25] to point in the gray-shaded area. As a result, differently from the GKM scheme the cited algorithm is not robust to any choice of direction (e.g., it cannot accept the one as in Figure 3). In any case, the half-space C_w onto which x is projected according to the GKM scheme is properly contained in the half-space H_w corresponding to the update of [25]; consequently, the GKM update is always closer to any solution.

$x = Lu + b$ for some constant $b \in \mathbb{R}^{Nn_x}$. The goal is to choose inputs that minimize a cost

$$\ell(\mathbf{u}, \mathbf{x}) = \sum_{t=0}^{N-1} \ell_t(u_t, x_t) + \ell_N(x_N) \quad (24b)$$

subject to some constraints

$$x_{t+1} \in \mathcal{X}_{t+1}, \quad u_t \in \mathcal{U}_t, \quad t = 0, \dots, N-1. \quad (24c)$$

A. Vü-Condat splitting

The constraint sets in (24c) are typically simple and easy to project onto (boxes, Euclidean balls...). However, while simple input constraints can be easily handled, due to the coupling enforced by the dynamics (24a) expressing \mathcal{X}_{t+1} in terms of the optimization variable \mathbf{u} results in much more complicated sets (polyhedra, ellipsoids...). To avoid this complication we make use of the extremely versatile algorithm that Vü-Condat three-term splitting offers [31, Alg. 3.1]. In its general form, the algorithm addresses problems of the form

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad f(x) + g(x) + h(Lx), \quad (25)$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex with L_f -Lipschitz continuous gradient, $g : \mathbb{R}^n \rightarrow \bar{\mathbb{R}}$ and $h : \mathbb{R}^m \rightarrow \bar{\mathbb{R}}$ are convex, and $L \in \mathbb{R}^{n \times m}$, by iterating the following steps:

$$\begin{cases} x^+ = \text{prox}_{\tau g}(x - \tau(\nabla f(x) + L^T y)) \\ y^+ = \text{prox}_{\tau h}(y + \sigma L(2x^+ - x)). \end{cases} \quad (26)$$

Here, $0 < \tau < \frac{2}{L_f}$ and $0 < \sigma < \frac{1}{\|L\|^2}(\frac{1}{\tau} - \frac{L_f}{2})$ are stepsizes, and $y \in \mathbb{R}^m$ is a Lagrange multiplier. Vü-Condat splitting is a primal-dual method that generalizes FBS by allowing an extra nonsmooth term h and a linear operator L (by neglecting h and L one recovers the proximal gradient iterations of FBS).

The optimal control problem (24) can be cast into Vü-Condat splitting form (25) by letting $f(\mathbf{u}) = \ell(\mathbf{u}, L\mathbf{u})$, $g = \delta_{\mathcal{U}}$ and $h = \delta_{\mathcal{X}}(\cdot + b)$,¹ where $\mathcal{U} = \mathcal{U}_0 \times \dots \times \mathcal{U}_{N-1}$ and $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_N$ (in particular, $n = Nn_u$ and $m = Nn_x$). Then, $\text{prox}_{\tau g} = \Pi_{\mathcal{U}}$ and $\text{prox}_{\sigma h}(y) = y - \sigma \Pi_{\mathcal{X}}(\sigma^{-1}y + b)$.

¹ δ_C denotes the indicator function of the nonempty closed convex set C , namely $\delta_C(x) = 0$ if $x \in C$ and $\delta_C(x) = \infty$ otherwise.

Notice that $\Pi_{\mathcal{U}}$ and $\Pi_{\mathcal{X}}$ are fully decoupled as the projection of each input and state onto the corresponding constraint set. Moreover, there is no need to compute the full matrix L , as both L and L^T can be treated as abstract operators that simulate forward and backward dynamics.

Apparently, the appeal of Vü-Condat splitting in addressing the optimal control problem lies in the extreme simplicity of its operations and low memory requirements, making it particularly suited for medium-to-large-scale problems in which traditional interior point algorithms fail. However, like all first-order methods it is extremely sensitive to ill conditioning, which gets worse as the problem size increases. Fortunately, this splitting fits into the *SuperMann* framework. The operator T that maps (x, y) into (x^+, y^+) as in (26) is averaged in the Hilbert space \mathcal{H}_P , where \mathcal{H}_P is defined as $\mathbb{R}^n \times \mathbb{R}^m$ equipped with the scalar product $\langle z, z' \rangle_P := \langle z, Pz' \rangle$ with $P := \begin{pmatrix} \tau^{-1}I & -L^T \\ -L & \sigma^{-1}I \end{pmatrix}$ [31, proof of Thm. 3.1].

B. Oscillating masses experiment

We tried this approach on the benchmark problem of controlling a chain of oscillating masses connected by springs and with both ends attached to walls. The chain is composed of $2K$ bodies of unit mass subject to a viscous friction of 0.1, the springs have elastic constant 1 and no damping, and the system is controlled through K actuators, each being a force acting on a pair of masses, as depicted in Figure 5. Therefore $n_x = 4K$ (the states are the displacement from the rest position and velocity of each mass) and $n_u = K$. The inputs are constrained in $[-2, 2]$, while the position and velocity of each mass is constrained in $[-5, 5]$.

The continuous-time system was discretized with a sampling time $T_s = 0.1s$. We considered quadratic stage costs $\frac{1}{2}x^T Q x$ for the states and $\frac{1}{2}u^T u$ for the inputs, where Q is diagonal positive definite with random diagonal entries, and generated a random (feasible) initial state x_0 . Notice that a QP reformulation would require the computation of the full cost matrix, differently from the splitting approach where only the

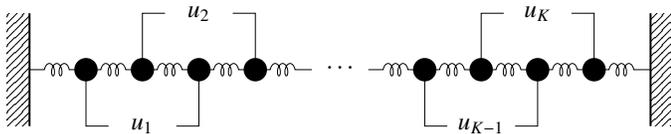


Figure 5: Oscillating masses

| | | Number of calls to L and L^\top ($\times 10^3$) | | | | | | | | | |
|----------|-----|---|-------|----------|-------|----------|------|----------|-------|----------|-------|
| $K = 8$ | | $N = 10$ | | $N = 20$ | | $N = 30$ | | $N = 40$ | | $N = 50$ | |
| | | avg | max | avg | max | avg | max | avg | max | avg | max |
| | VC | 19.0 | 337.1 | 15.0 | 174.4 | 25.0 | 400+ | 21.0 | 136.5 | 16.0 | 61.9 |
| | SVC | 1.0 | 5.5 | 1.0 | 4.3 | 2.0 | 19.3 | 2.0 | 10.9 | 2.0 | 6.6 |
| $K = 16$ | | $N = 10$ | | $N = 20$ | | $N = 30$ | | $N = 40$ | | $N = 50$ | |
| | | avg | max | avg | max | avg | max | avg | max | avg | max |
| | VC | 62.0 | 400+ | 30.0 | 344.9 | 30.0 | 400+ | 65.0 | 400+ | 29.0 | 318.6 |
| | SVC | 4.0 | 39.5 | 2.0 | 11.6 | 3.0 | 46.6 | 8.0 | 58.1 | 3.0 | 26.1 |

Table I: Comparison between Vü-Condat algorithm (VC) and its “super” enhancement (SuperVC) in solving the oscillating masses problem with $\|R x^k\| \leq 10^{-4} \|R x^0\|$ as termination criterion. Average and worst performances among 25 simulations with randomly generated starting point x_0 for each combination of $K \in \{8, 16\}$ and $N \in \{10, 20, 30, 40, 50\}$. The tables compare the number of calls to the operators L and L^\top , which are the expensive operations (the rest are projections on boxes). In four problems Vü-Condat exceeded $4 \cdot 10^5$ many calls (corresponding to 10^5 iterations) and was stopped prematurely.

small dynamics matrices A and B are needed, as L and L^\top can be abstract operators.

We simulated different scenarios for all combinations of $K \in \{8, 16\}$ and $N \in \{10, 20, 30, 40, 50\}$. We compared Vu-Condat splitting (VC) with its ‘super’ enhancement (SuperVC); parameters were set as detailed in Section VI-D. Figure 6 shows a comparison of the convergence rates for one problem instance, while Table I offers an overview of the whole experiment: SuperVC is roughly 13 times faster on average and 21 times better in worst-case performance than VC algorithm in reaching the termination criterion $\|R x^k\| \leq 10^{-4} \|R x^0\|$.

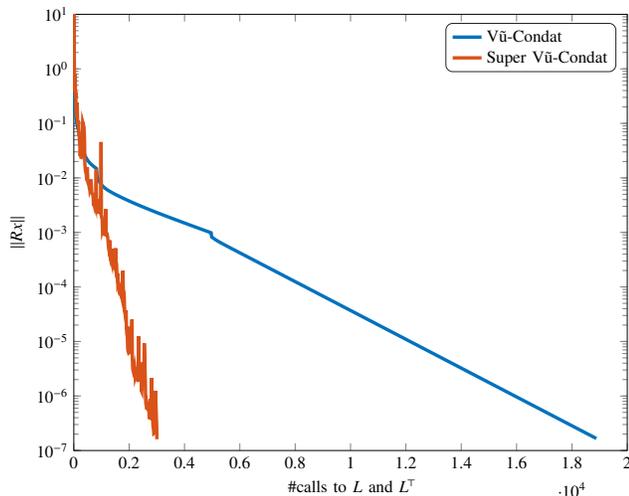


Figure 6: Random simulation of the spring-mass control problem with $K = 8$ and $N = 30$.

VIII. CONCLUSIONS

We proposed the *SuperMann scheme* (Alg. 2), a novel algorithm for finding fixed points of a nonexpansive operator T that generalizes and greatly improves the classical Krasnosel’skiĭ-Mann (KM) scheme, enjoying the same favorable properties: global convergence with worst-case sublinear rate, cheap iterations based solely on evaluations of T , and easy codability. The *SuperMann scheme* is an extremely versatile algorithm, its flexibility being twofold: on one hand it works for any nonexpansive operator T by requiring only the oracle $x \mapsto Tx$; on the other hand it allows for the integration of any fast local method for solving nonlinear equations, leaving much freedom for trading off cheap iterations or faster convergence. The remarkable performance of the method is supported both in practice with promising simulations and in theory where the employment of quasi-Newton directions is shown to yield asymptotic superlinear convergence rates provided a condition analogous to the famous result by Dennis and Moré is satisfied. Most importantly, superlinear convergence does not require nonsingularity of the Jacobian of the residual at the solution but merely metric subregularity, and as such can be achieved even when the solution is not isolated.

We encourage the employment of the *SuperMann scheme* to improve and robustify convex splitting algorithms; in particular, we strongly believe that its integration in generic solvers which are based on fixed-point iterations of nonexpansive operators such as SCS [32] would be extremely beneficial.

REFERENCES

- [1] N. Parikh and S. Boyd, “Proximal algorithms,” *Found. Trends Optim.*, vol. 1, no. 3, pp. 127–239, Jan. 2014.
- [2] G. Stathopoulos, H. Shukla, A. Szucs, Y. Pu, and C. N. Jones, “Operator splitting methods in control,” *Foundations and Trends in Systems and Control*, vol. 3, no. 3, pp. 249–362, 2016.
- [3] P. L. Combettes and J.-C. Pesquet, *Proximal Splitting Methods in Signal Processing*. New York, NY: Springer New York, 2011, pp. 185–212.
- [4] W. Ben-Ameur, P. Bianchi, and J. Jakubowicz, “Robust distributed consensus using total variation,” *IEEE Transactions on Automatic Control*, vol. 61, no. 6, pp. 1550–1564, June 2016.
- [5] F. Iutzeler, P. Bianchi, P. Ciblat, and W. Hachem, “Explicit convergence rate of a distributed alternating direction method of multipliers,” *IEEE Transactions on Automatic Control*, vol. 61, no. 4, pp. 892–904, April 2016.
- [6] P. L. Combettes and B. C. Vü, “Variable metric forward-backward splitting with applications to monotone inclusions in duality,” *Optimization*, vol. 63, no. 9, pp. 1289–1318, 2014.
- [7] Y. Zheng, G. Fantuzzi, A. Papachristodoulou, P. Goulart, and A. Wynn, “Fast ADMM for semidefinite programs with chordal sparsity,” in *2017 American Control Conference (ACC)*, May 2017, pp. 3335–3340.
- [8] F. Facchinei and J.-S. Pang, *Finite-dimensional variational inequalities and complementarity problems*. Springer, 2003, vol. II.
- [9] A. F. Izmailov and M. V. Solodov, *Newton-type methods for optimization and variational problems*. Springer, 2014.
- [10] M. Powell, “A hybrid method for nonlinear equations,” *Numerical Methods for Nonlinear Algebraic Equations*, pp. 87–144, 1970.
- [11] M. A. Krasnosel’skiĭ, “Two remarks on the method of successive approximations,” *Uspekhi Matematicheskikh Nauk*, vol. 10, no. 1, pp. 123–127, 1955.
- [12] W. R. Mann, “Mean value methods in iteration,” *Proceedings of the American Mathematical Society*, vol. 4, no. 3, pp. 506–510, 1953.
- [13] P. Giselsson, M. Fält, and S. Boyd, “Line search for averaged operator iteration,” in *2016 IEEE 55th Conference on Decision and Control (CDC)*, Dec 2016, pp. 1015–1022.
- [14] H. H. Bauschke and P. L. Combettes, *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2011.

- [15] P. L. Combettes, “Quasi-Fejérian analysis of some optimization algorithms,” *Studies in Computational Mathematics*, vol. 8, pp. 115–152, 2001.
- [16] J. Ermol’ev and A. D. Tuniev, “Random Fejér and quasi-Fejér sequences,” *Select. Translat. math. Statist. Probab.* 13, 143–148, 1973.
- [17] X. Chen and M. Fukushima, “Proximal quasi-Newton methods for nondifferentiable convex optimization,” *Mathematical Programming*, vol. 85, no. 2, pp. 313–334, 1999.
- [18] H. H. Bauschke, D. Noll, and H. M. Phan, “Linear and strong convergence of algorithms involving averaged nonexpansive operators,” *Journal of Mathematical Analysis and Applications*, vol. 421, no. 1, pp. 1–20, 2015.
- [19] A. L. Dontchev and R. T. Rockafellar, “Regularity and conditioning of solution mappings in variational analysis,” *Set-Valued Analysis*, vol. 12, no. 1-2, pp. 79–109, 2004.
- [20] R. T. Rockafellar and R. J.-B. Wets, *Variational analysis*. Springer Science & Business Media, 2011, vol. 317.
- [21] S. M. Robinson, *Some continuity properties of polyhedral multifunctions*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1981, pp. 206–214.
- [22] C.-M. Ip and J. Kyriaris, “Local convergence of quasi-Newton methods for B-differentiable equations,” *Mathematical Programming*, vol. 56, no. 1-3, pp. 71–89, 1992.
- [23] J.-S. Pang, “Newton’s method for B-differentiable equations,” *Mathematics of Operations Research*, vol. 15, no. 2, pp. 311–341, 1990.
- [24] F. J. Aragón Artacho, A. Belyakov, A. L. Dontchev, and M. López, “Local convergence of quasi-Newton methods under metric regularity,” *Computational Optimization and Applications*, vol. 58, no. 1, pp. 225–247, 2014.
- [25] M. V. Solodov and B. F. Svaiter, “A globally convergent inexact Newton method for systems of monotone equations,” in *Reformulation: Nonsmooth, Piecewise Smooth, Semismooth and Smoothing Methods*. Springer, 1998, pp. 355–369.
- [26] P. Patrinos and A. Bemporad, “Proximal Newton methods for convex composite optimization,” in *IEEE Conference on Decision and Control*, 2013, pp. 2358–2363.
- [27] P. Patrinos, L. Stella, and A. Bemporad, “Douglas-Rachford splitting: Complexity estimates and accelerated variants,” in *53rd IEEE Conference on Decision and Control*, Dec 2014, pp. 4234–4239.
- [28] L. Stella, A. Themelis, and P. Patrinos, “Forward-backward quasi-Newton methods for nonsmooth optimization problems,” *Computational Optimization and Applications*, vol. 67, no. 3, pp. 443–487, Jul 2017.
- [29] A. Themelis, L. Stella, and P. Patrinos, “Forward-backward envelope for the sum of two nonconvex functions: Further properties and nonmonotone linesearch algorithms,” *SIAM Journal on Optimization*, vol. 28, no. 3, pp. 2274–2303, 2018.
- [30] L. Stella, A. Themelis, P. Sotasakis, and P. Patrinos, “A simple and efficient algorithm for nonlinear model predictive control,” in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, Dec 2017, pp. 1939–1944.
- [31] L. Condat, “A primal-dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms,” *Journal of Optimization Theory and Applications*, vol. 158, no. 2, pp. 460–479, 2013.
- [32] B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd, “Conic optimization via operator splitting and homogeneous self-dual embedding,” *Journal of Optimization Theory and Applications*, pp. 1–27, 2016.

APPENDIX A PROOFS OF SECTION IV

Proof of Theorem IV.1.

♣ IV.1(i): we start observing that because of (6) and the triangular inequality, for all $k \in K_0 \cup K_1$ we have

$$\|x_{k+1} - z\| \leq \|x_k - z\| + D\|Rx_k\| \quad \forall z \in \mathbf{fix} T, \quad (27)$$

and since R is 2α -Lipschitz continuous we also have that

$$\|Rx_{k+1}\| \leq (1 + 2\alpha D)\|Rx_k\|. \quad (28)$$

By combining [15, Prop. 3.2(i)] with (4) and (27), it follows that in order to prove quasi-Fejér monotonicity it suffices to show that the sequence $(\|Rx_k\|)_{k \in K_0 \cup K_1}$ is summable. Let K_0 and

K_1 be indexed as in (5). Since η_k is kept constant whenever $k \notin K_0$,

$$\eta_{k_\ell} = \|Rx_{k_\ell}\| \leq c_0 \eta_{k_{\ell-1}} \cdots \leq c_0^{\ell-1} \eta_{k_1} = c_0^{\ell-1} \eta_0. \quad (29)$$

In particular, $(\|Rx_{k_\ell}\|)_{k_\ell \in K_0}$ is summable (regardless of whether K_0 is finite or not).

As for $k'_\ell \in K_1$, the safeguard parameter r_{safe} ensures that

$$\begin{aligned} \|Rx_{k'_\ell}\| &\leq \|Rx_{k'_{\ell-1}+1}\| + q^{k'_\ell-1} \leq c_1 \|Rx_{k'_{\ell-1}}\| + q^{k'_\ell-1} \\ &\leq c_1 \|Rx_{k'_{\ell-1}}\| + q^{\ell-1}. \end{aligned}$$

By iterating the inequality, for any $\rho \in (0, 1)$ such that $\rho > \max\{c_1, q\}$ we have

$$\|Rx_{k'_\ell}\| \leq \rho^{\ell-1} \|Rx_{k'_1}\| + \sum_{i=1}^{\ell-1} c_1^{i-1} \rho^{\ell-i} \leq C \rho^\ell, \quad (30)$$

where $C := \frac{1}{\rho} (\|Rx_{k'_1}\| + \sum_{i \in \mathbb{N}} (c_1/\rho)^i) < \infty$. In particular, also $(\|Rx_k\|)_{k \in K_1}$ is summable.

♣ IV.1(ii): due to quasi-Fejér monotonicity, for all $z \in \mathbf{fix} T$ there exists $(\varepsilon_k(z))_{k \in \mathbb{N}} \in \ell_1^+$ such that

$$\|x_{k+1} - z\|^2 \leq \|x_k - z\|^2 + \varepsilon_k(z).$$

By combining this with (4) and telescoping the inequalities, we obtain that for all $z \in \mathbf{fix} T$

$$\|x_0 - z\|^2 \geq \sigma \sum_{k \in K_2} \|Rx_k\|^2 - \sum_{k \in K_0 \cup K_1} \varepsilon_k(z). \quad (31)$$

Since the sequence $(\varepsilon_k(z))_{k \in K_0 \cup K_1}$ is summable, then so is $(\|Rx_k\|^2)_{k \in K_2}$. In turn, since $(\|Rx_k\|)_{k \in K_0 \cup K_1}$ is also summable it follows that the whole sequence of residuals is square-summable.

♣ IV.1(iii): follows by combining IV.1(ii) with Thm. III.3.

♣ IV.1(iv): trivially follows from the already proven point IV.1(ii), together with the observation that since η_k is kept constant whenever $k \notin K_0$, the condition $\|Rx_k\| \leq c_0 \eta_k$ will be satisfied infinitely often if $c_0 > 0$. \square

We now state two lemmas which will be needed in the proof of Theorem IV.3.

Lemma A.1 (Asymptotic properties of K_0 and K_1). *Suppose the hypotheses of Theorem IV.1 hold and let $(x_k)_{k \in \mathbb{N}}$ be the sequence generated by Algorithm 1. Then,*

- (i) $(\|Rx_k\|)_{k \in K_0}$ is Q -linearly convergent;
- (ii) $(\|Rx_k\|)_{k \in K_1}$ is R -linearly convergent;
- (iii) if $c_0 > 0$ then for some $\varrho \in (0, 1]$ and $\beta \in \mathbb{R}$

$$\ell_0(k) \geq \varrho \ell_1(k) - \beta \quad \forall k \in \mathbb{N},$$

where $\ell_j(k) := \#\{k' \in K_j \mid k' \leq k\}$, $j = 0, 1, 2$, is the number of times K_j was visited up to iteration k .

Proof.

♣ A.1(i) and A.1(ii): already shown in (29) and (30).

♣ A.1(iii): if $c_1 = 0$, then $K_1 = \emptyset$ and the claim trivially holds with $\varrho = 1$ and $\beta = 0$. Otherwise, from (30) and due to the definition of $\ell_1(k)$ there exist $C > 0$ and $\rho \in (0, 1)$ such that

$$\|Rx_k\| \leq C \rho^{\ell_1(k)} \quad \forall k \in K_1.$$

If $k \in K_1$, then $\|Rx_k\|$ didn’t pass the test at step 2, therefore

$$C \rho^{\ell_1(k)} \geq \|Rx_k\| \geq \eta_k = \|Rx_0\| c_0^{\ell_0(k)}.$$

The proof now follows by simply taking the logarithm on the outer inequality. \square

Lemma A.2. Let $(u_k)_{k \in \mathbb{N}} \subset [0, +\infty)$ be a sequence, and let $K_1, K_2 \subseteq \mathbb{N}$ be such that $\mathbb{N} = K_1 \cup K_2$. Let K_1 be indexed as $K_1 = \{k'_0, k'_1, \dots\}$, and suppose that there exist $a, b > 0$ and $\rho \in (0, 1)$ such that

$$\begin{cases} u_{k+1} \leq au_k & \text{for all } k \in \mathbb{N}, \\ u_{k'_\ell} \leq b\rho^\ell & \text{for all } k'_\ell \in K_1, \\ u_{k+1} \leq \rho u_k & \text{for all } k \in K_2. \end{cases}$$

Then, there exists $\sigma \in (0, 1)$ such that $u_k \leq ab\sigma^k$.

Proof. To exclude trivialities we assume that K_1 and K_2 are both infinite. To arrive to a contradiction, for all $\sigma \in (0, 1)$ let $k = k(\sigma)$ be the minimum such that $u_k > ab\sigma^k$. Let $\sigma \geq \rho$ be fixed. If $k-1 \in K_2$, then

$$\rho u_{k-1} \geq u_k > ab\sigma^k \geq ab\rho\sigma^{k-1}$$

and therefore $u_{k-1} > ab\sigma^{k-1}$ which contradicts minimality of k . It follows that necessarily $k-1 \in K_1$, hence $k-1 = k'_\ell \in K_1$ for some $\ell \in \mathbb{N}$. For all $n \in \mathbb{N}$, let $k'_{\ell_n} = k(\rho^{1/n}) - 1$, i.e., the minimum such that $u_{k'_{\ell_n}+1} > ab\rho^{\frac{k'_{\ell_n}+1}{n}}$. By combining with the property of K_1 we obtain

$$ab\rho^{\frac{k'_{\ell_n}+1}{n}} < u_{k'_{\ell_n}+1} \leq au_{k'_{\ell_n}} \leq ab\rho^{\ell_n}, \quad (32)$$

and in particular $\ell_n \leq \frac{k'_{\ell_n}}{n}$. This means that up to $k = k'_{\ell_n}$ there are at most k/n elements in K_1 , and consequently at least $k-k/n$ in K_2 . Therefore,

$$b\rho^{\frac{k+1}{n}} < u_k \leq a^{k/n}\rho^{k-k/n}u_0.$$

After taking the k -th square root on the outer inequality we are left with

$$(1/\rho)^{1-2/n-1/nk} < (u_0/b)^{1/k} a^{1/n}.$$

By letting $n \rightarrow +\infty$, so that also $k \rightarrow +\infty$, we arrive to the contradiction $\rho \geq 1$. \square

Proof of Theorem IV.3. Letting $e_k := \mathbf{dist}(x_k, \mathbf{fix} T)$, because of (28) and (8) there exists $B > 1$ such that

$$\|Rx_{k+1}\| \leq B\|Rx_k\| \quad \text{and} \quad e_{k+1} \leq Be_k \quad \forall k \in \mathbb{N}. \quad (33)$$

Suppose that R is metrically subregular at x_\star with radius $\varepsilon > 0$ and modulus $\gamma > 0$; since $x_k \rightarrow x_\star$, up to an index shifting without loss of generality we may assume that $(x_k)_{k \in \mathbb{N}} \subset \mathbf{B}(x_\star; \varepsilon)$. Let $z_k = \mathbf{P}_{\mathbf{fix} T} x_k$, so that $e_k = \|x_k - z_k\|$; by combining (4) and (8) we obtain that for all $k \in K_2$

$$e_{k+1}^2 \leq \|x_{k+1} - z_k\|^2 \leq \|x_k - z_k\|^2 - \sigma\|Rx_k\|^2 \leq \rho^2 e_k^2, \quad (34)$$

where $\rho := \sqrt{1 - \sigma/\gamma^2} \in (0, 1)$. By possibly enlarging ρ we may assume $\rho \geq \max\{c_0, c_1\}$.

If $c_0 = 0$, then $K_0 = \emptyset$ and using Lem. A.1(ii) and (33) we may invoke Lem. A.2 to infer R -linear convergence of the sequence $(e_k)_{k \in \mathbb{N}}$ and conclude the proof.

Therefore, let us suppose that $c_0 > 0$, so that by Thm. IV.1(iv) the set K_0 contains infinite many indices. We now show that there exists $n \in \mathbb{N}$ such that every n consecutive indices at least one is in K_0 . Let $k \in K_0$ be fixed and suppose that $k+1 \dots k+n+1 \notin K_0$.

- If $c_1 = 0$ then $K_1 = \emptyset$ and all such indices belong to K_2 . Then,

$$\begin{aligned} \|Rx_{k+n+1}\| &\stackrel{(8)}{\leq} 2\alpha e_{k+n+1} \stackrel{(34)}{\leq} 2\alpha\rho^n e_{k+1} \stackrel{(33)}{\leq} 2\alpha B\rho^n e_k \\ &\leq 2\alpha\gamma B\rho^n \|Rx_k\|. \end{aligned}$$

Since $k+n+1 \notin K_0$, then $\|Rx_{k+n+1}\|$ failed the test at step 3 and therefore

$$c_0\|Rx_k\| = c_0\eta_{k+n+1} < \|Rx_{k+n+1}\| \leq 2\alpha\gamma B\rho^n \|Rx_k\|,$$

which proves that n cannot be arbitrarily large.

- If instead $c_1 > 0$, let n_1 be the number of indices among $k+1 \dots k+n$ that belong to K_1 , and $n_2 = n - n_1$ those belonging to K_2 . Then, from iteration $k+1$ to $k+n+1$ the distance from the fixed set has reduced n_2 times (at least) by a factor ρ and, due to (33), increased at most by a factor B the remaining n_1 times:

$$\begin{aligned} \|Rx_{k+n+1}\| &\stackrel{(8)}{\leq} 2\alpha e_{k+n+1} \leq 2\alpha\rho^{n_2} B^{n_1} e_{k+1} \\ &\stackrel{(33)}{\leq} 2\alpha\rho^{n_2} B^{n_1+1} e_k \stackrel{(8)}{\leq} 2\alpha\gamma\rho^{n_2} B^{n_1+1} \|Rx_k\|. \end{aligned}$$

Again, since $k+n+1 \notin K_0$ we have $c_0\|Rx_k\| < 2\alpha\gamma\rho^{n_2} B^{n_1+1} \|Rx_k\|$, and therefore

$$n_1 > \frac{\ln c_0 / 2\alpha\gamma}{\ln B} - 1 + \frac{\ln 1/\rho}{\ln B} n_2.$$

In particular, for large n the number n_1 of indices in K_1 grows proportionally with respect to n , and from Lem. A.1(iii) we conclude once again that n cannot be arbitrarily large (since the number of visits to K_0 does not change from $k+1$ to $k+n$).

So far we proved that there exists $n \in \mathbb{N}$ such that every n indices at least one belongs to K_0 . In particular, indexing $K_0 = \{k_0, k_1, \dots\}$ we have that $k_\ell \leq n\ell$, hence for all $k_\ell \in K_0$

$$\|Rx_{k_\ell}\| \leq c_0^\ell \|Rx_0\| \leq (c_0^{1/n})^{k_\ell} \|Rx_0\|. \quad (35)$$

Moreover, any $k \in \mathbb{N}$ is at most $n-1$ indices away from the nearest previous index $k_\ell \in K_0$; combined with (35) and invoking (33) we obtain

$$\|Rx_k\| \leq B^{n-1} \|Rx_0\| (c_0^{1/n})^{k_\ell} \leq B^{n-1} \|Rx_0\| (c_0^{1/n})^k,$$

proving the sought R -linear convergence of $(\|Rx_k\|)_{k \in \mathbb{N}}$. It follows that for some $b > 0$ and $r \in (0, 1)$ we have $\|Rx_k\| \leq br^k$ for all $k \in \mathbb{N}$; then,

$$\begin{aligned} \|x_k - x_\star\| &\leq \sum_{j \geq k} \|x_{j+1} - x_j\| \leq D \sum_{j \geq k} \|Rx_j\| \leq bD \sum_{j \geq k} r^j \\ &= \frac{bD}{1-r} r^k, \end{aligned}$$

where in the second inequality we used the bound (6), which also holds for $k \in K_2$ (up to possibly enlarging D) due to the fact that for $k \in K_2$ under metric subregularity we have

$$\|x_{k+1} - x_k\| \leq \|x_{k+1} - z_k\| + \|x_k - z_k\| \leq 2e_k \stackrel{(8)}{\leq} 2\gamma\|Rx_k\|.$$

This shows that $(x_k)_{k \in \mathbb{N}}$ is R -linearly convergent too. \square

APPENDIX B

PROOFS OF SECTION VI

Proof of Theorem VI.1. Because of Thm. V.4 we know that for any direction d_k a feasible stepsize τ_k complying with the requirements of step 5(b) will eventually be found, lower bounded as in VI.1(ii) due to Thm. V.4 and Assumption II. In particular, the scheme is well defined. Moreover, from Prop. V.2(ii) we have that there exists a constant $\underline{\sigma} > 0$ such that

$$\|x_{k+1} - z\|^2 \leq \|x_k - z\|^2 - \underline{\sigma}\|Rx_k\|^2$$

for all $k \in K_2$ and $z \in \mathbf{fix} T$. It follows that the *SuperMann scheme* is a special case of Alg. 1 and the proof entirely follows from Thm.s IV.1 and IV.3. \square

Proof of Theorem VI.4.

♣ VI.4(i): let $w_k^0 := x_k + d_k$. Superlinear convergence of $(d_k)_{k \in \mathbb{N}}$ then reads $\frac{\|Rw_k^0\|}{\|Rx_k\|} \rightarrow 0$. In particular, if $c_1 > 0$ then there exists $\bar{k} \in \mathbb{N}$ such that $\|Rw_k^0\| \leq c_1 \|Rx_k\|$ for all $k \geq \bar{k}$, i.e., the point $w_k^0 = x_k + d_k$ will always pass condition at step 5(a) resulting in $x_{k+1} = w_k^0 = x_k + d_k$ for all $k \geq \bar{k}$.

Similarly, if $c_0 > 0$ then K_0 is infinite as shown in Thm. VI.1(v); moreover, for $\ell \in \mathbb{N}$

$$\frac{\|Rx_{k_\ell+1}\|}{\eta_{k_\ell+1}} = \frac{\|Rx_{k_\ell+1}\|}{\|Rx_{k_\ell}\|} = \frac{\|R(x_{k_\ell} + d_{k_\ell})\|}{\|Rx_{k_\ell}\|} \rightarrow 0$$

as $\ell \rightarrow \infty$, and therefore the ratio eventually is always smaller than c_0 , resulting in $k_\ell + 1 \in K_0$ for ℓ large enough. Consequently, the sequence will eventually reduce to $x_{k+1} = x_k + d_k$.

♣ VI.4(ii) and VI.4(iii): Q -superlinear convergence of the sequence $(Rx_k)_{k \in \mathbb{N}}$ follows from the fact that $x_{k+1} = x_k + d_k$ for $k \geq \bar{k}$. In particular, $(\|Rx_k\|)_{k \in \mathbb{N}}$ is summable and there exists a sequence $(\delta_k)_{k \in \mathbb{N}} \rightarrow 0$ such that $\|Rx_{k+1}\| \leq \delta_k \|Rx_k\|$ for all k . If $\|d_k\| \leq D \|Rx_k\|$ for some $D > 0$, then

$$\sum_{k \geq \bar{k}} \|x_{k+1} - x_k\| \leq D \sum_{k \geq \bar{k}} \|Rx_k\| < \infty,$$

which implies that $(x_k)_{k \in \mathbb{N}}$ is a Cauchy sequence, and hence converges to a point, be it x_* . Moreover, by possibly enlarging D so as to account for the iterates $k < \bar{k}$, we have

$$\begin{aligned} \|x_k - x_*\| &\leq \sum_{j \geq k} \|x_{j+1} - x_j\| \leq D \sum_{j \geq k} \|Rx_j\| \\ &\leq D \delta_0 \delta_1 \cdots \delta_{k-1} \sum_{j \in \mathbb{N}} \|Rx_j\| =: \Delta_k. \end{aligned}$$

This shows that $(x_k)_{k \in \mathbb{N}}$ is R -superlinearly convergent, since $\Delta_{k+1}/\Delta_k = \delta_k \rightarrow 0$.

♣ VI.4(iv): already shown in the proof of VI.4(i). \square

Proof of Theorem VI.8. Let $G_* := JRx_* \in \mathbb{R}^{n \times n}$ and let $\|\cdot\|$ denote the Euclidean norm. From [22, Lem. 2.2] we have that there exist a constant L and a neighborhood U_{x_*} of x_* such that

$$\begin{aligned} \frac{\|y_k - G_* s_k\|}{\|s_k\|} &= \frac{\|Rw_k - Rx_k - G_*(w_k - x_k)\|}{\|w_k - x_k\|} \\ &\leq L \max\{\|x_k - x_*\|, \|w_k - x_*\|\}. \end{aligned}$$

Because of (17), the fact that $\tau_k \leq 1$, and the triangular inequality we have $\|w_k - x_*\| \leq \|x_k - x_*\| + D \|Rx_k\|$ and consequently

$$\sum_{k \in \mathbb{N}} \frac{\|y_k - G_* s_k\|}{\|s_k\|} \leq L \sum_{k \in \mathbb{N}} (\|x_k - x_*\| + D \|Rx_k\|) < \infty,$$

where the last inequality follows from Thm. VI.1(vi).

Let $E_k := B_k - G_*$ and let $\|\cdot\|_F$ denote the Frobenius norm. With a simple modification of the proofs of [22, Thm. 4.1] and [24, Lem. 4.4] that takes into account the scalar $\vartheta_k \in [\bar{\vartheta}, 2 - \bar{\vartheta}]$ we obtain

$$\begin{aligned} \|E_{k+1}\|_F &\leq \left\| E_k \left(\text{id} - \vartheta_k \frac{s_k s_k^\top}{\|s_k\|^2} \right) \right\|_F + \vartheta_k \frac{\|y_k - G_* s_k\|}{\|s_k\|} \\ &\leq \|E_k\|_F - \frac{\bar{\vartheta}(2 - \bar{\vartheta})}{2\|E_k\|_F} \frac{\|E_k s_k\|^2}{\|s_k\|^2} + (2 - \bar{\vartheta}) \frac{\|y_k - G_* s_k\|}{\|s_k\|}. \end{aligned}$$

The last term on the right-hand side, be it σ_k , is summable and therefore the sequence $(E_k)_{k \in \mathbb{N}}$ is bounded. Let $\bar{E} := \sup(\|E_k\|_F)_{k \in \mathbb{N}}$, then

$$\|E_{k+1}\|_F - \|E_k\|_F \leq \sigma_k - \frac{\bar{\vartheta}(2 - \bar{\vartheta})}{2\bar{E}} \left(\frac{\|(B_k - G_*)s_k\|}{\|s_k\|} \right)^2.$$

Telescoping the above inequality, summability of σ_k ensures that of $\frac{\|(B_k - G_*)s_k\|^2}{\|s_k\|^2}$ proving in particular the claimed Dennis-Moré condition (20). \square



Andreas Themelis is a postdoctoral researcher at the Department of Electrical Engineering (ESAT) of KU Leuven, Belgium. He received a joint PhD degree in Control and Optimization from KU Leuven and IMT Lucca (Italy) in 2018, and Bachelor and Master degrees in Mathematics from the University of Florence, Italy, in 2010 and 2013, respectively. His research focuses on nonconvex nonsmooth optimization, with particular interest in splitting algorithms, monotone operator theory, and variational analysis.



Panagiotis (Panos) Patrinos is assistant professor at the Department of Electrical Engineering (ESAT) of KU Leuven, Belgium. In 2014 he was a visiting professor at Stanford. He received his PhD in Control and Optimization, M.S. in Applied Mathematics and M.Eng. from the National Technical University of Athens in 2010, 2005 and 2003, respectively. After his PhD he held postdoc positions at the University of Trento and IMT Lucca, Italy, where he became an assistant professor in 2012. His current research interests are in the theory and algorithms of structured convex and nonconvex optimization, learning and predictive control with a focus on large-scale, distributed, stochastic and embedded optimization and a range of application areas including automotive, aerospace, machine learning, signal processing and energy.