Linear Convergence and Metric Selection for Douglas-Rachford Splitting and ADMM

Pontus Giselsson[*] and Stephen Boyd[†]

April 12, 2016

Abstract

Recently, several convergence rate results for Douglas-Rachford splitting and the alternating direction method of multipliers (ADMM) have been presented in the literature. In this paper, we show global linear convergence rate bounds for Douglas-Rachford splitting and ADMM under strong convexity and smoothness assumptions. We further show that the rate bounds are tight for the class of problems under consideration for all feasible algorithm parameters. For problems that satisfy the assumptions, we show how to select step-size and metric for the algorithm that optimize the derived convergence rate bounds. For problems with a similar structure that do not satisfy the assumptions, we present heuristic step-size and metric selection methods.

1 Introduction

Optimization problems of the form

$$\text{minimize} \quad f(x) + g(x) \tag{1}$$

where x is the variable and f and g are proper closed and convex, arise in numerous applications ranging from compressed sensing [9] and statistical estimation [29] to model predictive control [45] and medical imaging [37].

There exist a variety of operator splitting methods for solving problems of the form (1), see [12, 39]. In this paper, we focus on Douglas-Rachford splitting [17, 41]. Douglas-Rachford splitting is given by the iteration

$$z^{k+1} = ((1-\alpha)\text{Id} + \alpha R_{\gamma f} R_{\gamma g})z^k, \tag{2}$$

---

[*]Department of Automatic Control, Lund University. Email: `pontusg@control.lth.se`.
[†]Electrical Engineering Department, Stanford University. Email: `boyd@stanford.edu`.

where α in general is constrained to satisfy $\alpha \in (0,1)$. Here Id is the identity operator, $R_{\gamma f}$ is the reflected proximal operator $R_{\gamma f} = 2\mathrm{prox}_{\gamma f} - \mathrm{Id}$, and $\mathrm{prox}_{\gamma f}$ is the proximal operator defined by

$$\mathrm{prox}_{\gamma f}(z) := \underset{x}{\mathrm{argmin}}\{\gamma f(x) + \tfrac{1}{2}\|x - z\|^2\}. \tag{3}$$

The variable z^k in (2) converges to a fixed-point of $R_{\gamma f} R_{\gamma g}$ and the sequence $x^k = \mathrm{prox}_{\gamma g} z^k$ converges to a solution of (1) (if it exists), see [1, Proposition 25.6].

When (2) is applied to the Fenchel dual problem of (1), this algorithm is equivalent to ADMM, see [28, 21, 6] for ADMM and [20, 18] for the connection to Douglas-Rachford splitting. These methods have long been known to converge for any $\alpha \in (0,1)$ and $\gamma > 0$ under mild assumptions, [21, 36, 18]. However, the rate of convergence in the general case has just recently been shown to be $O(1/k)$, [30, 15, 13].

For a restricted class of monotone inclusion problems, Lions and Mercier showed in [36] that the Douglas-Rachford algorithm (with $\alpha = 0.5$) enjoys a linear convergence rate. To the authors' knowledge, this was the sole linear convergence rate results for a long period of time for these methods. Recently, however, many works have shown linear convergence rates for Douglas-Rachford splitting and ADMM in different settings [31, 42, 15, 14, 16, 22, 40, 32, 34, 5, 47, 2, 44, 38].

The works in [31, 15, 5, 42] concern local linear convergence under different assumptions. The works in [34, 47] consider distributed formulations and [32] considers multi-block ADMM, while the works in [14, 16, 22, 40, 36, 2, 44, 38] show global linear convergence. Of these, the work in [2] shows linear convergence for Douglas-Rachford splitting when solving a subspace intersection problem and the work in [44] (which appeared online during the submission procedure of this paper) shows linear convergence for equality constrained problems with upper and lower bounds on the variables. The other works that show global linear convergence, [14, 16, 22, 40, 36, 38], show this for Douglas-Rachford splitting or ADMM under strong convexity and smoothness assumptions.

In this paper, we provide explicit linear convergence rate bounds for the case where f in (1) is strongly convex and smooth and we show how to select algorithm parameters to optimize the bounds. We also show that the bounds are tight for the class of problems under consideration for all feasible algorithm parameters γ and α. Our results generalize and/or improve corresponding results in [14, 16, 22, 40, 36, 38]. A detailed comparison to the rate bounds in these works is found in Section 5.

We also show that a relaxation factor $\alpha \geq 1$ can be used in the algorithm. We provide explicit lower and upper bounds on α where the upper bound depends on the smoothness and strong convexity parameters of the problem. We further show that the bounds are tight, meaning that if an α is chosen that is outside the allowed interval, there exists a problem from the considered class that the algorithm does not solve.

The provided convergence rate bounds for Douglas-Rachford splitting and ADMM depend on the smoothness and strong convexity parameters of the problem to be solved. These parameters depend on the space on which the problem is defined. This space can therefore be chosen by optimizing the convergence rate bound. In this paper, we show how to do this in Euclidean settings. We select a space by choosing a (metric) matrix M that defines an inner-product $\langle x, y \rangle_M = x^T M y$ and its induced norm $\|x\|_M = \sqrt{\langle x, x \rangle_M}$. This matrix M is chosen to optimize the convergence rate bound (typically subject to M being diagonal). Letting the problem and algorithm be defined on this space can considerably improve theoretical and practical convergence.

The metric selection can be interpreted as a preconditioning with the objective to reduce the number of iterations. A similar preconditioning is presented in [22] for ADMM applied to solve quadratic programs with linear inequality constraints. Our results generalize these in several directions, see Section 5 for a detailed comparison.

Real-world problems rarely have the properties needed to ensure linear convergence of Douglas-Rachford splitting or ADMM. Therefore, we provide heuristic metric and parameter selection methods for such problems. The heuristics cover many problems that arise, e.g., in model predictive control [45], statistical estimation [29, 48], compressed sensing [9], and medical imaging [37]. We provide two numerical examples that show the efficiency of the theoretically justified and heuristic parameter and metric selection methods.

This paper extends and generalizes the conference papers [25, 24].

1.1   Notation

We denote by $\mathbb{R}$ the set of real numbers, $\mathbb{R}^n$ the set of real column-vectors of length n, and $\mathbb{R}^{m \times n}$ the set of real matrices with m rows and n columns. Further $\bar{\mathbb{R}} := \mathbb{R} \cup \{\infty\}$ denotes the extended real line. Throughout this paper $\mathcal{H}$, $\mathcal{H}_1$, $\mathcal{H}_2$, and $\mathcal{K}$ denote real Hilbert spaces. Their inner products are denoted by $\langle \cdot, \cdot \rangle$, the induced norm by $\| \cdot \|$, and the identity operator by Id. We specifically consider finite-dimensional Hilbert-spaces $\mathbb{H}_H$ with inner product $\langle x, y \rangle = x^T H y$ and induced norm $\|x\| = \sqrt{x^T H x}$. We denote these by $\langle \cdot, \cdot \rangle_H$ and $\| \cdot \|_H$. We also denote the Euclidean inner-product by $\langle x, y \rangle_2 = x^T y$ and the induced norm by $\| \cdot \|_2$. The conjugate function is denoted and defined by $f^*(y) \triangleq \sup_x \{\langle y, x \rangle - f(x)\}$ and the adjoint operator to a linear operator $\mathcal{A} : \mathcal{H} \to \mathcal{K}$ is defined as the unique operator $\mathcal{A}^* : \mathcal{K} \to \mathcal{H}$ that satisfies $\langle \mathcal{A}x, y \rangle = \langle x, \mathcal{A}^* y \rangle$. The range and kernel of $\mathcal{A}$ are denoted by $\mathrm{ran}\mathcal{A}$ and $\mathrm{ker}\mathcal{A}$ respectively. The orthogonal complement of a subspace $\mathcal{X}$ is denoted by $\mathcal{X}^\perp$. The power set of a set $\mathcal{X}$, i.e., the set of all subsets of $\mathcal{X}$, is denoted by $2^{\mathcal{X}}$. The graph of an (set-valued) operator $A : \mathcal{X} \to 2^{\mathcal{Y}}$ is defined and denoted by $\mathrm{gph}A = \{(x, y) \in \mathcal{X} \times \mathcal{Y} \mid y \in Ax\}$. The inverse operator A^{-1} is defined through its graph by $\mathrm{gph}A^{-1} = \{(y, x) \in \mathcal{Y} \times \mathcal{X} \mid y \in Ax\}$. The set of fixed-points to an operator $T : \mathcal{H} \to \mathcal{H}$ is denoted and defined as $\mathrm{fix}T = \{x \in \mathcal{H} \mid Tx = x\}$. Finally, the class of closed, proper, and convex functions $f : \mathcal{H} \to \bar{\mathbb{R}}$ is denoted by $\Gamma_0(\mathcal{H})$.

2   Background

In this section, we introduce some standard definitions that can be found, e.g.
in [1, 46].

Definition 1 (Strong monotonicity) *An operator* $A \ : \ \mathcal{H} \to 2^{\mathcal{H}}$ *is* σ-
strongly monotone *with* $\sigma > 0$ *if*

$$\langle u - v, x - y \rangle \geq \sigma \|x - y\|^2$$

for all $(x, u) \in \mathrm{gph}(A)$ *and* $(y, v) \in \mathrm{gph}(A)$.

The definition of *monotonicity* is obtained by setting $\sigma = 0$ in the above
definition.

In the following definitions, we suppose that $\mathcal{D} \subseteq \mathcal{H}$ is a nonempty subset of
$\mathcal{H}$.

Definition 2 (Lipschitz mappings) *A mapping* $T \ : \ \mathcal{D} \to \mathcal{H}$ *is* β-Lipschitz
continuous *with* $\beta \geq 0$ *if*

$$\|Tx - Ty\| \leq \beta \|x - y\|$$

holds for all $x, y \in \mathcal{D}$. *If* $\beta = 1$ *then* T *is* nonexpansive *and if* $\beta \in [0, 1)$ *then* T
is β-contractive.

Definition 3 (Averaged mappings) *A mapping* $T \ : \ \mathcal{D} \to \mathcal{H}$ *is* α-averaged
if there exists a nonexpansive mapping $S \ : \ \mathcal{D} \to \mathcal{H}$ *and* $\alpha \in (0, 1)$ *such that*
$T = (1 - \alpha)\mathrm{Id} + \alpha S$.

Definition 4 (Cocoercivity) *A mapping* $T \ : \ \mathcal{D} \to \mathcal{H}$ *is* β-cocoercive with
$\beta > 0$ *if* βT *is* $\frac{1}{2}$-averaged.

This definition implies that cocoercive mappings T can be expressed as

$$T = \tfrac{1}{2\beta}(\mathrm{Id} + S) \tag{4}$$

for some nonexpansive operator S.

Definition 5 (Strong convexity) *A function* $f \in \Gamma_0(\mathcal{H})$ *is* σ-strongly convex
with $\sigma > 0$ *if* $f - \frac{\sigma}{2}\| \cdot \|^2$ *is convex.*

A strongly convex function has a minimum curvature that is at least σ. If
f is differentiable, strong convexity can equivalently be defined as that

$$f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \tfrac{\sigma}{2}\|x - y\|^2 \tag{5}$$

holds for all $x, y \in \mathcal{H}$. Functions with a maximal curvature are called smooth.
Next, we present a smoothness definition for convex functions.

Definition 6 (Smoothness for convex functions) *A function $f \in \Gamma_0(\mathcal{H})$ is β-smooth with $\beta \geq 0$ if it is differentiable and $\frac{\beta}{2}\|\cdot\|^2 - f$ is convex, or equivalently that*

$$f(x) \leq f(y) + \langle \nabla f(y), x - y \rangle + \frac{\beta}{2}\|x - y\|^2 \tag{6}$$

holds for all $x, y \in \mathcal{H}$.

Remark 1 *It can be seen from (5) and (6) that for a function that is σ-strongly convex and β-smooth, we always have $\beta \geq \sigma$.*

3 Douglas-Rachford splitting

The Douglas-Rachford algorithm can be applied to solve composite convex optimization problems of the form

$$\text{minimize} \quad f(x) + g(x) \tag{7}$$

where $f, g \in \Gamma_0(\mathcal{H})$. The solutions to (7) are characterized by the following optimality conditions, [1, Proposition 25.1]

$$z = R_{\gamma g} R_{\gamma f} z, \qquad\qquad x = \text{prox}_{\gamma f}(z) \tag{8}$$

where $\gamma > 0$, the prox operator $\text{prox}_{\gamma f}$ is defined in (3), and the reflected proximal operator $R_{\gamma f} = 2\text{prox}_{\gamma f} - \text{Id}$. In other words, a solution to (7) is found by applying the proximal operator on z, where z is a fixed-point to $R_{\gamma g} R_{\gamma f}$.

One approach to find a fixed-point to $R_{\gamma g} R_{\gamma f}$ is to iterate the composition

$$z^{k+1} = R_{\gamma g} R_{\gamma f} z^k.$$

This algorithm is sometimes referred to as Peaceman-Rachford splitting, [41]. However, since $R_{\gamma f}$ and $R_{\gamma g}$ are in general nonexpansive, so is their composition, and convergence of this algorithm cannot be guaranteed in the general case.

The Douglas-Rachford splitting algorithm is obtained by iterating the averaged map of the nonexpansive Peaceman-Rachford operator $R_{\gamma g} R_{\gamma f}$. That is, it is given by the iteration

$$z^{k+1} = ((1 - \alpha)\text{Id} + \alpha R_{\gamma g} R_{\gamma f})z^k \tag{9}$$

where $\alpha \in (0, 1)$ to guarantee convergence in the general case, see [19]. (We will, however, see that when additional regularity assumptions are introduced, $\alpha = 1$, i.e. Peaceman-Rachford splitting, and even some $\alpha > 1$ can be used and convergence can still be guaranteed.)

The Douglas-Rachford algorithm (9) can more explicitly be written as

$$x^k = \text{prox}_{\gamma f}(z^k) \tag{10}$$
$$y^k = \text{prox}_{\gamma g}(2x^k - z^k) \tag{11}$$
$$z^{k+1} = z^k + 2\alpha(y^k - x^k) \tag{12}$$

Note that sometimes, the algorithm obtained by letting $\alpha = \frac{1}{2}$ is called Douglas-Rachford splitting [18]. Here we use the name Douglas-Rachford splitting for all feasible values of α.

3.1 Linear convergence

Under some regularity assumptions, the convergence of the Douglas-Rachford algorithm is linear. We will analyze its convergence under the following set of assumptions:

Assumption 1 *Suppose that:*

(i) *f and g are proper, closed, and convex.*

(ii) *f is σ-strongly convex and β-smooth.*

To show linear convergence rates of the Douglas-Rachford algorithm under these regularity assumptions on f, we need to characterize some properties of the proximal and reflected proximal operators to f. Specifically, we will show that the reflected proximal operator of f is contractive (as opposed to nonexpansive in the general case). We will also provide a tight contraction factor.

The key to arriving at this contraction factor is the following, to the authors' knowledge, novel (but straightforward) interpretation of the proximal operator.

Proposition 1 *Assume that $f \in \Gamma_0(\mathcal{H})$ and that $\gamma \in (0, \infty)$ and define f_γ as*

$$f_\gamma := \gamma f + \tfrac{1}{2}\|\cdot\|^2. \tag{13}$$

Then $\mathrm{prox}_{\gamma f}(y) = \nabla f_\gamma^*(y)$.

Proof. Since the proximal operator is the resolvent of $\gamma \partial f$, see [1, Example 23.3], we have $\mathrm{prox}_{\gamma f}(y) = (\mathrm{Id} + \gamma \partial f)^{-1} y = (\partial f_\gamma)^{-1} y$. Since $f \in \Gamma_0(\mathcal{H})$ and $\gamma \in (0, \infty)$ also $f_\gamma \in \Gamma_0(\mathcal{H})$. Therefore [1, Corollary 16.24] implies that $\mathrm{prox}_{\gamma f}(y) = (\partial f_\gamma)^{-1} y = \nabla f_\gamma^*(y)$, where differentiability of f_γ^* follows from [1, Theorem 18.15], since f_γ is 1-strongly convex, and since $f = (f^*)^*$ for proper, closed, and convex functions, see [1, Theorem 13.32]. This concludes the proof. $\square$

This interpretation of the proximal operator is used to prove the following proposition. The proof is found in Appendix B.

Proposition 2 *Assume that $f \in \Gamma_0(\mathcal{H})$ is σ-strongly convex and β-smooth and that $\gamma \in (0, \infty)$. Then $\mathrm{prox}_{\gamma f} - \frac{1}{1+\gamma\beta}\mathrm{Id}$ is $\frac{1}{\frac{1}{1+\gamma\sigma} - \frac{1}{1+\gamma\beta}}$-cocoercive if $\beta > \sigma$ and 0-Lipschitz if $\beta = \sigma$.*

This result is used to show the following contraction properties of the reflected proximal operator. A proof to this result, which is one of the main results of the paper, is found in Appendix C.

Theorem 1 *Suppose that $f \in \Gamma_0(\mathcal{H})$ is σ-strongly convex and β-smooth and that $\gamma \in (0, \infty)$. Then $R_{\gamma f}$ is $\max\left(\frac{\gamma\beta-1}{\gamma\beta+1}, \frac{1-\gamma\sigma}{\gamma\sigma+1}\right)$-contractive.*

For future reference, we let this contraction factor be denoted by δ, i.e.,

$$\delta := \max\left(\tfrac{\gamma\beta-1}{\gamma\beta+1}, \tfrac{1-\gamma\sigma}{\gamma\sigma+1}\right). \tag{14}$$

Theorem 1 lays the foundation for the linear convergence rate result in the following theorem, which is proven in Appendix D.

Theorem 2 *Suppose that Assumption 1 holds, that $\gamma \in (0,\infty)$, that $\alpha \in (0, \frac{2}{1+\delta})$ with δ in (14). Then the Douglas Rachford algorithm (9) converges linearly towards a fixed-point $\bar{z} \in \operatorname{fix}(R_{\gamma g} R_{\gamma f})$ at least with rate $|1-\alpha|+\alpha\delta$, i.e.,:*

$$\|z^{k+1} - \bar{z}\| \leq (|1-\alpha|+\alpha\delta)\,\|z^k - \bar{z}\|.$$

Remark 2 *Note that $\alpha > 1$ can be used in the Douglas-Rachford algorithm when solving problems that satisfy Assumption 1. (This also holds for general relaxed iteration of a contractive mapping.) That α-values greater than 1 can be used is reported in [38] for ADMM (i.e., for dual Douglas-Rachford). They solve a small SDP to assess whether ADMM converges with a specific rate, for specific problem parameters β and σ, and specific algorithm parameters α and γ. This SDP gives an affirmative answer for some problems and some $\alpha > 1$. As opposed to here, no explicit bounds on α are provided.*

Also the x^k iterates in (10) converge linearly. The following corollary is proven in Appendix E.

Corollary 1 *Let $x^\star$ be the (unique) solution to (7). Then the x^k iterates in (10) satisfy*

$$\|x^{k+1} - x^\star\| \leq (|1-\alpha|+\delta\alpha)^{k+1}\,\tfrac{1}{1+\gamma\sigma}\|z^0 - \bar{z}\|. \tag{15}$$

Remark 3 *Note that Corollary 1 and Theorem 2 still hold if the order of $R_{\gamma g}$ and $R_{\gamma f}$ is swapped in the Douglas-Rachford algorithm (9).*

We can choose the algorithm parameters γ and α to optimize the bound on the convergence rate in Theorem 2. This is done in the following proposition.

Proposition 3 *Suppose that Assumption 1 holds. Then the optimal parameters for the Douglas-Rachford algorithm in (9) are $\alpha = 1$ and $\gamma = \frac{1}{\sqrt{\sigma\beta}}$. Further, the optimal rate is $\frac{\sqrt{\beta/\sigma}-1}{\sqrt{\beta/\sigma}+1}$.*

Proof. Since $\delta \in [0,1)$, the rate in Theorem 2, $|1-\alpha|+\alpha\delta$, is a decreasing function of α for $\alpha \leq 1$ and increasing for $\alpha \geq 1$. Therefore the rate factor is optimized by $\alpha = 1$. The γ parameter should be chosen to minimize the

max-expression $\max\left(\frac{\gamma\beta-1}{\gamma\beta+1}, \frac{1-\gamma\sigma}{1+\gamma\sigma}\right)$ defining δ in (14). This is done by setting the arguments equal, which gives $\gamma = 1/\sqrt{\beta\sigma}$. Inserting these values into the rate factor expression gives $\frac{\sqrt{\beta/\sigma}-1}{\sqrt{\beta/\sigma}+1}$. $\qquad\square$

Remark 4 *Note that $\alpha = 1$ is optimal in Proposition 3. So the Peaceman-Rachford algorithm gives the best bound on the convergence rate under Assumption 1, even though it is not guaranteed to converge in the general case. The reason is that $R_{\gamma f}$ becomes contractive under Assumption 1 (as opposed to non-expansive in the general case).*

3.2 Tightness of bounds

In this section, we show that the linear convergence rate bounds in Theorem 2 are tight. We consider a two dimensional example of the form (7) in a standard Euclidean space to show tightness. Let $x = (x_1, x_2)$, then the functions used are

$$f(x) = \tfrac{1}{2}(\beta x_1^2 + \sigma x_2^2), \tag{16}$$
$$g_1(x) = 0, \tag{17}$$
$$g_2(x) = \iota_{x=0}(x), \tag{18}$$

where $\beta \geq \sigma > 0$ and $\iota_{x=0}$ is the indicator function for $x = 0$, i.e., $\iota_{x=0}(x) = 0$ iff $x = 0$, otherwise ∞.

The function f is β-smooth since $\frac{\beta}{2}\|x\|_2^2 - f(x) = \frac{\beta-\sigma}{2}x_2^2$ is convex. It is σ-strongly convex since $f(x) - \frac{\sigma}{2}\|x\|_2^2 = \frac{\beta-\sigma}{2}x_1^2$ is convex. So the function f satisfies Assumption 1(ii).

The proximal operator to f is

$$\begin{aligned}
\mathrm{prox}_{\gamma f}(y) &= \underset{x}{\mathrm{argmin}}\{\gamma f(x) + \tfrac{1}{2}\|x - y\|_2^2\} \\
&= \underset{x}{\mathrm{argmin}}\{\tfrac{1}{2}\left((\beta\gamma + 1)x_1^2 + (\sigma\gamma + 1)x_2^2\right) + x_1 y_1 + x_2 y_2\} \\
&= (\tfrac{1}{1+\gamma\beta}y_1, \tfrac{1}{1+\gamma\sigma}y_2)
\end{aligned} \tag{19}$$

where $y = (y_1, y_2)$. The reflected proximal operator is

$$\begin{aligned}
R_{\gamma f}(y) &= 2\mathrm{prox}_{\gamma f}(y) - y \\
&= 2(\tfrac{1}{1+\gamma\beta}y_1, \tfrac{1}{1+\gamma\sigma}y_2) - (y_1, y_2) \\
&= (\tfrac{1-\gamma\beta}{1+\gamma\beta}y_1, \tfrac{1-\gamma\sigma}{1+\gamma\sigma}y_2).
\end{aligned} \tag{20}$$

The proximal and reflected proximal operators to g_1 in (17) are $\mathrm{prox}_{\gamma g_1} = R_{\gamma g_1} = \mathrm{Id}$. The proximal and reflected proximal operators to g_2 in (18) are $\mathrm{prox}_{\gamma g_2}(x) = 0$ and $R_{\gamma g_2} = 2\mathrm{prox}_{\gamma g_2} - \mathrm{Id} = -\mathrm{Id}$.

Next, these results are used to show tightness of the convergence rate bounds in Theorem 2. The tightness result is proven in Appendix F.

Theorem 3 *The convergence rate bound in Theorem 2 for the Douglas-Rachford splitting algorithm (9) is tight for the class of problems that satisfy Assumption 1 for all algorithm parameters specified in Theorem 2, namely $\alpha \in (0, \frac{2}{1+\delta})$ with δ in (14) and $\gamma \in (0, \infty)$. If instead $\alpha \notin (0, \frac{2}{1+\delta})$, there exist a problem that satisfies Assumption 1 that the Douglas-Rachford algorithm does not solve.*

Remark 5 *This result on tightness can be generalized to any real Hilbert space. This is obtained by considering the same g_1 and g_2 as in (17) and (18) but with*

$$f(x) = \sum_{i=1}^{|\mathcal{H}|} \tfrac{\lambda_i}{2} \langle x, \phi_i \rangle^2.$$

Here $\{\phi_i\}_{i=1}^{|\mathcal{H}|}$ is an orthonormal basis for $\mathcal{H}$, $|\mathcal{H}|$ is the dimension of the space $\mathcal{H}$ (possibly infinite), and $\lambda_i = \sigma > 0$ or $\lambda_i = \beta \geq \sigma$ for all i, where at least one $\lambda_i = \sigma$ and one $\lambda_i = \beta$.

4 ADMM

In this section, we apply the results from the previous section to the Fenchel dual problem. Since ADMM applied to the primal problem is equivalent to the Douglas-Rachford algorithm applied to the dual problem, the results obtained in this section hold for ADMM.

We consider solving problems of the form

$$\begin{array}{ll} \text{minimize} & f(x) + g(y) \\ \text{subject to} & \mathcal{A}x + \mathcal{B}y = c \end{array} \tag{21}$$

where $f \in \Gamma_0(\mathcal{H}_1)$, $g \in \Gamma_0(\mathcal{H}_2)$, $\mathcal{A} : \mathcal{H}_1 \to \mathcal{K}$ and $\mathcal{B} : \mathcal{H}_2 \to \mathcal{K}$ are bounded linear operators and $c \in \mathcal{K}$. In addition, we assume:

Assumption 2

(i) $f \in \Gamma_0(\mathcal{H}_1)$ is β-smooth and σ-strongly convex.

(ii) $\mathcal{A} : \mathcal{H}_1 \to \mathcal{K}$ is surjective.

The assumption that $\mathcal{A}$ is a surjective bounded linear operator reduces to that $\mathcal{A}$ is a real matrix with full row rank in the Euclidean case.

Problems of the form (21) cannot be directly efficiently solved using Douglas-Rachford splitting. Therefore, we instead solve the (negative) Fenchel dual problem, which is

$$\text{minimize} \quad d_1(\mu) + d_2(\mu) \tag{22}$$

where $d_1, d_2 \in \Gamma_0(\mathcal{K})$ are

$$d_1(\mu) := f^*(-\mathcal{A}^*\mu) + \langle c, \mu \rangle, \qquad\qquad d_2(\mu) = g^*(-\mathcal{B}^*\mu). \tag{23}$$

The Douglas-Rachford algorithm when solving the dual problem becomes

$$z^{k+1} = ((1 - \alpha)\mathrm{Id} + \alpha R_{\gamma d_1} R_{\gamma d_2})z^k. \tag{24}$$

This formulation will be used when analyzing ADMM since for $\alpha = \frac{1}{2}$ it is equivalent to ADMM and for $\alpha \in (0, \frac{1}{2}]$ and $\alpha \in [\frac{1}{2}, 1)$ it is equivalent to under- and over-relaxed ADMM, see [20, 18, 19]. Note that we only use (24) as an analysis algorithm for ADMM. The algorithm should still be implemented as the standard ADMM algorithm (with over- or under-relaxation). Here we state ADMM with scaled dual variables u, see [6]:

$$x^{k+1} = \underset{x}{\mathrm{argmin}}\{f(x) + \tfrac{\gamma}{2}\|Ax + By^k - c + u^k\|_2^2\} \tag{25}$$

$$x_A^{k+1} = 2\alpha Ax^{k+1} - (1 - 2\alpha)(By^k - c) \tag{26}$$

$$y^{k+1} = \underset{y}{\mathrm{argmin}}\{g(y) + \tfrac{\gamma}{2}\|x_A^{k+1} + By - c + u^k\|_2^2\} \tag{27}$$

$$u^{k+1} = u^k + (x_A^{k+1} + By^{k+1} - c) \tag{28}$$

where z^k in (24) satisfies $z^k = \gamma(u^k - By^k)$, see [19, Theorem 8] and [27, Appendix B].

4.1 Linear convergence

To show linear convergence rate results in this dual setting, we need to quantify the strong convexity and smoothness parameters for d_1 in (23). This is done in the following proposition.

Proposition 4 *Suppose that Assumption 2 holds. Then $d_1 \in \Gamma_0(\mathcal{K})$ in (23) is $\frac{\|\mathcal{A}^*\|^2}{\sigma}$-smooth and $\frac{\theta^2}{\beta}$-strongly convex, where $\theta > 0$ always exists and satisfies $\|\mathcal{A}^*\mu\| \geq \theta\|\mu\|$ for all $\mu \in \mathcal{K}$.*

Proof. We define $d(\mu) := f^*(-\mathcal{A}^*\mu)$, so $d_1(\mu) = d(\mu) + \langle c, \mu \rangle$. We first note that the linear term $\langle c, \mu \rangle$ does not affect the strong convexity or smoothness parameters. So, we proceed by showing d is $\frac{\|\mathcal{A}^*\|^2}{\sigma}$-smooth and $\frac{\theta^2}{\beta}$-strongly convex.

Since f is σ-strongly convex, [1, Theorem 18.15] gives that f^* is $\frac{1}{\sigma}$-smooth and that ∇f^* is $\frac{1}{\sigma}$-Lipschitz continuous. Therefore, ∇d satisfies

$$\|\nabla d(\mu) - \nabla d(\nu)\| = \|\mathcal{A}\nabla f^*(-\mathcal{A}^*\mu) - \mathcal{A}\nabla f^*(-\mathcal{A}^*\nu)\|$$
$$\leq \tfrac{\|\mathcal{A}\|}{\sigma}\|\mathcal{A}^*(\mu - \nu)\|$$
$$\leq \tfrac{\|\mathcal{A}^*\|^2}{\sigma}\|\mu - \nu\|$$

since $\|\mathcal{A}\| = \|\mathcal{A}^*\|$. This is equivalent to that d is $\frac{\|\mathcal{A}^*\|^2}{\sigma}$-smooth, see [1, Theorem 18.15].

Next, we show the strong convexity result for d. Since f is β-smooth, f^* is $\frac{1}{\beta}$-strongly convex, and ∇f^* is $\frac{1}{\beta}$-strongly monotone, see [1, Theorem 18.15]. Therefore, ∇d satisfies

$$\begin{aligned} \langle \nabla d(\mu) - \nabla d(\nu), \mu - \nu \rangle &= \langle -\mathcal{A}(\nabla f^*(-\mathcal{A}^*\mu) - \nabla f^*(-\mathcal{A}^*\nu)), \mu - \nu \rangle \\ &= \langle \nabla f^*(-\mathcal{A}^*\mu) - \nabla f^*(-\mathcal{A}^*\nu), -\mathcal{A}^*\mu + \mathcal{A}^*\nu \rangle \\ &\geq \tfrac{1}{\beta}\|\mathcal{A}^*(\mu - \nu)\|^2 \\ &\geq \tfrac{\theta^2}{\beta}\|\mu - \nu\|^2. \end{aligned}$$

This, by [1, Theorem 18.15], is equivalent to d being $\frac{\theta^2}{\beta}$-strongly convex. That $\theta > 0$ follows from [1, Fact 2.18 and Fact 2.19]. Specifically, [1, Fact 2.18] says that $\ker\mathcal{A}^* = (\mathrm{ran}\mathcal{A})^\perp = \emptyset$, since $\mathcal{A}$ is surjective. Since $\mathrm{ran}\mathcal{A} = \mathcal{K}$ (again by surjectivity), it is closed. Then [1, Fact 2.19] states that there exists $\theta > 0$ such that $\|\mathcal{A}^*\mu\| \geq \theta\|\mu\|$ for all $\mu \in (\ker\mathcal{A}^*)^\perp = (\emptyset)^\perp = \mathcal{K}$. $\qquad\square$

Combining this result with Theorem 1 implies that $R_{\gamma d_1}$ is $\hat{\delta}$-contractive with

$$\hat{\delta} := \max\left(\tfrac{\gamma\hat{\beta}-1}{1+\gamma\hat{\beta}}, \tfrac{1-\gamma\hat{\sigma}}{1+\gamma\hat{\sigma}}\right) \tag{29}$$

where $\hat{\beta} = \frac{\|\mathcal{A}^*\|^2}{\sigma}$ and $\hat{\sigma} = \frac{\theta^2}{\beta}$. This gives the following immediate corollary to Theorem 2 and Proposition 3.

Corollary 2 *Suppose that Assumption 2 holds, that $\gamma \in (0, \infty)$, and that $\alpha \in (0, \frac{2}{1+\hat{\delta}})$ with $\hat{\delta}$ in (29). Then algorithm (24) (or equivalently ADMM applied to solve (21)) converges at least with rate $|1 - \alpha| + \alpha\hat{\delta}$, i.e.,*

$$\|z^{k+1} - \bar{z}\| \leq \left(|1 - \alpha| + \alpha\hat{\delta}\right)\|z^k - \bar{z}\|$$

where $\bar{z} \in \mathrm{fix}(R_{\gamma d_1} R_{\gamma d_2})$.

Further, the algorithm parameters γ and α that optimize the rate bound are $\alpha = 1$ and $\gamma = \frac{1}{\sqrt{\hat{\beta}\hat{\sigma}}} = \frac{\sqrt{\beta\sigma}}{\|\mathcal{A}^*\|\theta}$. The optimized linear convergence rate bound factor is $\frac{\sqrt{\hat{\kappa}}-1}{\sqrt{\hat{\kappa}}+1}$, where $\hat{\kappa} = \frac{\hat{\beta}}{\hat{\sigma}} = \frac{\|\mathcal{A}^*\|^2\beta}{\theta^2\sigma}$.

Remark 6 *Similarly to in the primal setting in Corollary 1, R-linear convergence for the x^k update in (25) can be shown. The proof is more elaborate than the corresponding proof for primal Douglas-Rachford. This result is therefore omitted due to space considerations.*

4.2 Tightness of bounds

Also in this dual setting, the convergence rate estimates are tight. Consider again the functions (16), (17), and (18). Further let $c = 0$, $\mathcal{B} = -I$, and $\mathcal{A} := \begin{bmatrix} \theta & 0 \\ 0 & \zeta \end{bmatrix}$ with $\zeta \geq \theta > 0$.

Since $f^*(\nu) = \frac{1}{2}(\frac{1}{\beta}\nu_1^2 + \frac{1}{\sigma}\nu_2^2)$ for f in (16), we get

$$d_1(\mu) = f^*(-\mathcal{A}^T\mu) = \tfrac{1}{2}(\tfrac{\theta^2}{\beta}\mu_1^2 + \tfrac{\zeta^2}{\sigma}\mu_2^2).$$

Since d_1 is on the same form as f in Section 3.2, d_1 is $\hat{\sigma} := \frac{\theta^2}{\beta}$-strongly convex and $\hat{\beta} := \frac{\zeta^2}{\sigma}$-smooth, or equivalently $\hat{\beta} := \frac{\|\mathcal{A}^T\|^2}{\sigma}$-smooth since $\zeta = \|\mathcal{A}\| = \|\mathcal{A}^T\|$. Further $d_2(\mu) = g^*(-\mathcal{B}^T\mu) = g^*(\mu)$ where either $g = g_1$ with g_1 in (17) or $g = g_2$ with g_2 in (18). The functions g_1 and g_2 in (17) and (18) are each other's conjugates, i.e., $g_1^* = g_2$ and $g_2^* = g_1$. Therefore, the dual problem

$$\text{minimize} \quad d_1(\mu) + d_2(\mu)$$

(with $d_2 = g_1$ or $d_2 = g_2$) is the same as the problem considered in Section 3.2 but with the smoothness parameter β in Section 3.2 replaced by $\hat{\beta} = \frac{\|\mathcal{A}^*\|^2}{\sigma}$ and the strong convexity modulus σ in Section 3.2 replaced by $\hat{\sigma} = \frac{\theta^2}{\beta}$. Therefore, we can state the following immediate corollary to Theorem 3.

Corollary 3 *The convergence rate bound in Corollary 2 for ADMM applied to the primal or, equivalently, the Douglas-Rachford algorithm applied to the dual (24), is tight for the class of problems that satisfy Assumption 2 for all algorithm parameters specified in Corollary 2, namely for $\alpha \in (0, \frac{2}{1+\hat{\delta}})$ with $\hat{\delta}$ in (29) and $\gamma \in (0, \infty)$. If instead $\alpha \notin (0, \frac{2}{1+\hat{\delta}})$, there exist a problem that satisfies Assumption 2 that the algorithm does not solve.*

Remark 7 *As in the primal Douglas-Rachford case, tightness can be shown in general real Hilbert spaces. This is obtained with the same f, g_1, and g_2 as in Remark 5 and with $\mathcal{B} = -\mathrm{Id}$, $c = 0$, and*

$$\mathcal{A}x = \sum_{i=1}^{|\mathcal{H}|} \nu_i \langle x, \phi_i \rangle \phi_i$$

where $\nu_i = \theta > 0$ if $\lambda_i = \beta$ in Remark 5 and $\nu_i = \zeta \geq \theta$ if $\lambda_i = \sigma$ in Remark 5.

5 Related work

Recently, many works have appeared that show linear convergence rate bounds for Douglas-Rachford splitting and ADMM under different sets of assumptions. In this section, we will discuss and compare some of these linear convergence rate bounds that hold with the assumptions stated in this paper. For simplicity, we will compare the various optimal rates, i.e., rates obtained by using optimal algorithm parameters.

Specifically, we will compare our results in Proposition 3 and Corollary 2 to the previously known linear convergence rate results in [14, 16, 22, 40, 36] and the linear convergence rate [38] that appeared online during the submission procedure of this paper.
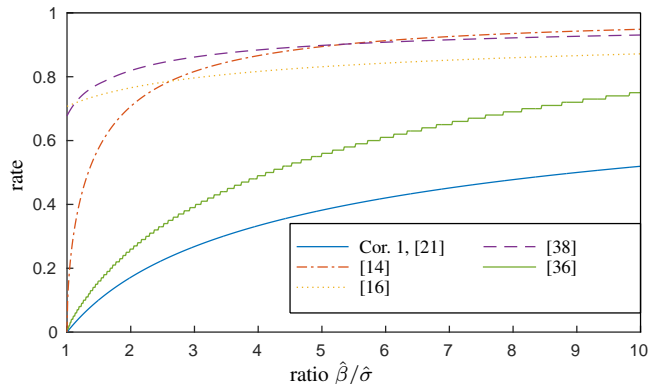
Figure 1: Comparison between linear convergence rate bounds for Douglas-Rachford splitting/ADMM provided in [14, 16, 22, 40, 38] and in Corollary 2.

Some of these results consider the Douglas-Rachford algorithm while others consider ADMM. A Douglas-Rachford rate result becomes an ADMM rate result by replacing the strong convexity and smoothness parameters σ and β with the dual problem counterparts $\hat{\sigma} = \frac{\theta^2}{\beta}$ and $\hat{\beta} = \frac{\|\mathcal{A}^*\|^2}{\sigma}$, see Proposition 4. All rate bounds in this comparison that directly analyze ADMM also depend on the dual problem strong convexity and smoothness parameters $\hat{\sigma}$ and $\hat{\beta}$. Therefore, the comparison can be carried out in dual problem parameters $\hat{\sigma}$ and $\hat{\beta}$. Obviously, since our bounds are tight, none of the other bounds can be better. This comparison is merely to show that we improve and/or generalize previous results.

In [36, Proposition 4, Remark 10], the linear convergence rate for Douglas-Rachford splitting with $\alpha = \frac{1}{2}$ when solving monotone inclusion problems with one operator being $\hat{\sigma}$-strongly monotone and $\hat{\beta}$-Lipschitz continuous is shown to be $\sqrt{1 - \frac{\hat{\sigma}}{2\hat{\beta}}}$. Note that the setting considered in [36] is more general than the setting in this paper. Recently, this bound was improved in [23] where tight estimates are provided.

In [14, Theorem 6], dual Douglas-Rachford splitting with $\alpha = 1$ is shown to converge linearly at least with rate $\sqrt{1 - \frac{\hat{\sigma}}{\hat{\beta}}}$. For $\alpha = \frac{1}{2}$ they recover the bound in [36]. In Figure 1, the former (better) of these rate bounds is plotted. We see that it is more conservative than the one in Corollary 2.

The optimal convergence rate bound in [16, Corollary 3.6] is $\sqrt{1/(1 + 1/\sqrt{\hat{\beta}/\hat{\sigma}})}$ and is analyzed directly in the ADMM setting. Figure 1 reveals that our rate bound is tighter.

In [40], the authors show that if the γ parameter is small enough and if f is a quadratic function, then Douglas-Rachford splitting is equivalent to a gradient method applied to a smooth convex function named the Douglas-Rachford

envelope. Convergence rate estimates then follow from the gradient method rate estimates. Also, accelerated variants of Douglas-Rachford splitting are proposed, based on fast gradient methods. In Figure 1, we compare to the rate bound of the fast Douglas-Rachford splitting in [40, Theorem 6] when applied to the dual. This rate bound is better than the rate bound for standard Douglas-Rachford splitting in [40, Theorem 4]. We note that Corollary 2 gives better rate bounds.

The convergence rate bound provided in [22] coincides with the bound provided in Corollary 2. The rate bound in [22] holds for ADMM applied to Euclidean quadratic problems with linear inequality constraints. We generalize these results, using a different machinery, to arbitrary real Hilbert spaces (also infinite dimensional), to both Douglas-Rachford splitting and ADMM, to general smooth and strongly convex functions f, and, perhaps most importantly, to any proper, closed, and convex function g.

Finally, we compare our rate bound to the rate bound in [38]. Figure 1 shows that our bound is tighter. As opposed to all the other rate bounds in this comparison, the rate bound in [38] is not explicit. Rather, a sweep over different rate bound factors is needed. For each guess, a small semi-definite program is solved to assess whether the algorithm is guaranteed to converge with that rate. The quantization level of this sweep is the cause of the steps in the rate curve in Figure 1.

6 Metric selection

In this section, we consider problems of the form (21) in an Euclidean setting. We assume $f \in \Gamma_0(\mathbb{H}_M)$, $g \in \Gamma_0(\mathbb{H}_{\hat{K}})$, $\mathcal{A} : \mathbb{H}_M \to \mathbb{H}_K$, $\mathcal{B} : \mathbb{H}_{\hat{K}} \to \mathbb{H}_K$, $c \in \mathbb{H}_K$ and that:

Assumption 3

(i) $f \in \Gamma_0(\mathbb{H}_M)$ is 1-strongly convex if defined on $\mathbb{H}_H$ (i.e., $M = H$) and 1-smooth if defined on $\mathbb{H}_L$ (i.e., $M = L$).

(ii) The bounded linear operator $\mathcal{A} : \mathbb{H}_M \to \mathbb{H}_K$ is surjective.

We solve (21) by applying Douglas-Rachford splitting on the dual problem (22) (or equivalently by applying ADMM directly on the primal (21)). This algorithm behaves differently depending on which space $\mathbb{H}_K$ the problem is defined and the algorithm is run. We will show how to select a space $\mathbb{H}_K$ on which the algorithm rate bound is optimized. To aid in this selection, we show in the following proposition how the strong convexity modulus and smoothness constant of $d_1 \in \Gamma_0(\mathbb{H}_K)$ depend on the space on which it is defined.

Proposition 5 *Suppose that Assumption 3 holds, that $A \in \mathbb{R}^{m \times n}$ satisfies $Ax = \mathcal{A}x$ for all x, and that $K = E^T E$. Then $d_1 \in \mathbb{H}_K$ in (23) is $\|EAH^{-1}A^T E^T\|_2$-smooth and $\lambda_{\min}(EAL^{-1}A^T E^T)$-strongly convex, where $\lambda_{\min}(EAL^{-1}A^T E^T) > 0$.*

Proof. First, we note that $d_1(\mu) = d(\mu) + \langle c, \mu \rangle$ where $d(\mu) = f^*(-\mathcal{A}^*\mu)$ and that a linear function does not change the strong convexity of smoothness parameter of a problem. Therefore, we show the result for d.

First, we relate $\mathcal{A}^* : \mathbb{H}_K \to \mathbb{H}_M$ to A, M, and K. We have

$$\langle \mathcal{A}x, \mu \rangle_K = \langle Ax, K\mu \rangle_2 = \langle x, A^T K\mu \rangle_2 = \langle x, M^{-1} A^T K\mu \rangle_M$$
$$= \langle M^{-1} A^T K\mu, x \rangle_M = \langle \mathcal{A}^*\mu, x \rangle_M.$$

Thus, $\mathcal{A}^*\mu = M^{-1} A^T K\mu$ for all $\mu \in \mathbb{H}_K$.

Next, we show that the space $\mathbb{H}_M$ on which f and f^* are defined does not influence the shape of d. We denote by f_H, f_L, and f_e the function f defined on $\mathbb{H}_H$, $\mathbb{H}_L$ and $\mathbb{R}^n$ respectively and by $\mathcal{A}_H^* : \mathbb{H}_K \to \mathbb{H}_H$, $\mathcal{A}_L^* : \mathbb{H}_K \to \mathbb{H}_L$, and $A^T : \mathbb{R}^m \to \mathbb{R}^n$ the operator $\mathcal{A}^*$ defined on the different spaces. Further, let $d_H := f_H^* \circ -\mathcal{A}_H^*$, $d_L := f_L^* \circ -\mathcal{A}_L^*$, and $d_e := f_e^* \circ -A^T$. With these definitions both d_L and d_H are defined on $\mathbb{H}_K$, while d_e is defined on $\mathbb{R}^m$. Next we show that d_L and d_H are identical for any μ:

$$d_H(\mu) = f_H^*(-\mathcal{A}_H^*\mu)$$
$$= \sup_x \left\{ \langle -\mathcal{A}_H^*\mu, x \rangle_H - f_H(x) \right\}$$
$$= \sup_x \left\{ \langle -HH^{-1} A^T K\mu, x \rangle_2 - f_e(x) \right\}$$
$$= \sup_x \left\{ \langle -LL^{-1} A^T K\mu, x \rangle_2 - f_e(x) \right\}$$
$$= \sup_x \left\{ \langle -\mathcal{A}_L^*\mu, x \rangle_L - f_L(x) \right\} = d_L(\mu)$$

where $\mathcal{A}_M^*\mu = M^{-1} A^T K\mu$ is used. This implies that we can show properties of $d \in \mathbb{H}_K$ by defining f on any space $\mathbb{H}_M$. Thus, Proposition 4 gives that 1-strong convexity of f when defined on $\mathbb{H}_H$ implies $\|\mathcal{A}^*\|^2$-smoothness of d, where

$$\|\mathcal{A}^*\| = \sup_\mu \left\{ \|\mathcal{A}^*\mu\|_H \mid \|\mu\|_K \leq 1 \right\}$$
$$= \sup_\mu \left\{ \|H^{-1} A^T K\mu\|_H \mid \|\mu\|_K \leq 1 \right\}$$
$$= \sup_\mu \left\{ \|H^{-1/2} A^T E^T E\mu\|_2 \mid \|E\mu\|_2 \leq 1 \right\}$$
$$= \sup_\nu \left\{ \|H^{-1/2} A^T E^T \nu\|_2 \mid \|\nu\|_2 \leq 1 \right\}$$
$$= \|H^{-1/2} A^T E^T\|_2.$$

Squaring this gives the smoothness claim. To show the strong-convexity claim, we use that 1-smoothness of f when defined on $\mathbb{H}_L$ implies θ^2-strong convexity of d where $\theta > 0$ satisfies $\|\mathcal{A}^*\mu\|_L \geq \theta\|\mu\|_K$ for all $\mu \in \mathbb{H}_K$, see Proposition 4.

We have

$$\begin{aligned}
\|\mathcal{A}^*\mu\|_L^2 &= \|L^{-1}A^T K\mu\|_L^2 \\
&= \|L^{-1/2}A^T E^T (E\mu)\|_2^2 \\
&= \|E\mu\|_{EAL^{-1}A^T E^T}^2 \\
&\geq \lambda_{\min}(EAL^{-1}A^T E^T)\|E\mu\|_2^2 \\
&= \lambda_{\min}(EAL^{-1}A^T E^T)\|\mu\|_K^2,
\end{aligned}$$

i.e, $\theta^2 = \lambda_{\min}(EAL^{-1}A^T E^T)$. The smallest eigenvalue $\lambda_{\min}(EAL^{-1}A^T E^T) > 0$ since A is surjective, i.e. has full row rank, and E and L are positive definite matrices. This concludes the proof. $\qquad\square$

This result shows how the smoothness constant and strong convexity modulus of $d_1 \in \Gamma_0(\mathbb{H}_K)$ change with the space $\mathbb{H}_K$ on which it is defined. Combining this with Proposition 3, we get that the bound on the convergence rate for Douglas-Rachford splitting applied to the dual problem (22) (or equivalently ADMM applied to the primal (21)) is optimized by choosing $K = E^T E$ where E is chosen to

$$\text{minimize} \quad \frac{\lambda_{\max}(EAH^{-1}A^T E^T)}{\lambda_{\min}(EAL^{-1}A^T E^T)} \tag{30}$$

and by choosing $\gamma = \frac{1}{\sqrt{\lambda_{\max}(EAH^{-1}A^T E^T)\lambda_{\min}(EAL^{-1}A^T E^T)}}$. Using a non-diagonal K usually gives prohibitively expensive prox-evaluations. Therefore, we propose to select a diagonal $K = E^T E$ that minimizes (30). The reader is referred to [26, Section 6] for different methods to achieve this exactly and approximately.

Remark 8 *It can be shown (details are omitted for space considerations) that solving the dual problem on space $\mathbb{H}_K$ using Douglas-Rachford splitting is equivalent to solving the preconditioned problem*

$$\begin{aligned}
&\textit{minimize} \quad f(x) + g(y) \\
&\textit{subject to} \quad E(Ax + By) = Ec.
\end{aligned} \tag{31}$$

using ADMM. Therefore, the metric selection presented here covers the preconditioning suggestion in [22] as a special case.

Remark 9 *Metric selection and preconditioning is not new for iterative methods. It can be used with the aim to reduce the computational burden within each iteration [11, 43, 8, 7, 33], and/or with the aim to reduce the total number of iterations like in our analysis and in [22, 26, 4, 7, 33]. It is interesting to note that in our paper (if we restrict ourselves to quadratic f with Hessian H) and in [22, 26, 7, 33], different algorithms are analyzed with different methods, but all analyses suggest to make $AH^{-1}A^T$ well conditioned (by choosing a metric or doing preconditioning). The analyzed algorithms are ADMM here and in [22], (fast) dual forward-backward splitting in [26], and Uzawa's method to*

solve linear systems of the form $\begin{bmatrix} H & A^T \\ A & 0 \end{bmatrix} (x, \mu) = (-q, b)$ in [7, 33]. This linear system is the KKT-condition for the problem $\min_x \{f(x) + g(y) \mid Ax = y\}$ where $f(x) = \frac{1}{2}x^T H x + q^T x$ and $g(y) = \iota_{y=b}(y)$. The dual functions are $d_1(\mu) = (q + A^T \mu)^T H^{-1}(A^T \mu + q)$ with Hessian $AH^{-1}A^T$ and $d_2(\mu) = g^*(\mu) = b^T \mu$. The functions d_1 in the dual problems in all these analyses are the same (if we restrict ourselves to quadratic f with Hessian H in our setting), but the functions d_2 are different. So all the different metric selections (preconditionings) try to make d_1, and its Hessian $AH^{-1}A^T$, well conditioned.

6.1 Heuristic metric selection

Many problems do not satisfy all assumptions in Assumption 2, but for many interesting problems a couple of them are typically satisfied. Therefore, we will here discuss metric and parameter selection heuristics for ADMM (Douglas-Rachford applied to the dual (22)) when some of the assumptions are not met. We focus on quadratic problems of the form

$$\text{minimize} \quad \underbrace{\frac{1}{2}x^T Q x + q^T x + \hat{f}(x)}_{f(x)} + g(y) \tag{32}$$
$$\text{subject to} \quad Ax + By = c$$

where $Q \in \mathbb{R}^{n \times n}$ is positive semi-definite, $q \in \mathbb{R}^n$, $\hat{f} \in \Gamma_0(\mathbb{R}^n)$, $g \in \Gamma_0(\mathbb{R}^m)$, $A \in \mathbb{R}^{p \times n}$, $B \in \mathbb{R}^{p \times m}$, and $c \in \mathbb{R}^p$.

One set of assumptions that guarantee linear convergence for dual Douglas-Rachford splitting is that Q is positive definite, that $\hat{f}$ is nonsmooth, and that A has full row rank. By inverting these, we get a set of assumptions that if anyone of these are satisfied, we cannot guarantee linear convergence:

(i) Q is not positive definite, but positive semi-definite.

(ii) $\hat{f}$ is nonsmooth, e.g., the indicator function of a convex constraint set or a piece-wise affine function

(iii) A does not have full row rank.

In the first case, we loose strong convexity in f and smoothness in the dual d_1. In the second case, we loose smoothness in f and strong convexity d_1. In the third case, we loose strong convexity in the dual d_1.

We directly tackle the case where the assumptions needed to get linear convergence are violated by all points (i), (ii), and (iii). For the dual case we consider, i.e., the ADMM case, we propose to select the metric as if $\hat{f} \equiv 0$ in (32). To do this, we define the quadratic part of f in (32) to be $f_{\mathrm{pc}}(x) := \frac{1}{2}x^T Q x + q^T x$ and introduce the function $d_{\mathrm{pc}} := f_{\mathrm{pc}}^* \circ -A^T$. The conjugate function of f_{pc} is given by

$$f_{\mathrm{pc}}^*(y) = \sup_x \left\{ \langle y, x \rangle - f_{\mathrm{pc}}(x) \right\} = \begin{cases} \frac{1}{2}(y - q)^T Q^\dagger (y - q) & \text{if } (y - q) \in \mathcal{R}(Q) \\ \infty & \text{else} \end{cases}$$

where $Q^\dagger$ is the pseudo-inverse of Q and $\mathcal{R}$ denotes the range space. This gives

$$d_{\mathrm{pc}}(\mu) = \begin{cases} \frac{1}{2}(A^T\mu + q)^T Q^\dagger (A^T\mu - q) & \text{if } (A^T\mu + q) \in \mathcal{R}(Q) \\ \infty & \text{else} \end{cases}$$

with Hessian $AQ^\dagger A^T$ on its domain. We approximate the dual function $d_1(\mu) = f^*(-A^T\mu) + c^T\mu$ with $d_{\mathrm{pc}}(\mu) + c^T\mu$. Then we propose to select a diagonal metric $K = E^T E$ such that this approximation is well conditioned in directions where it has curvature. That is, we propose to select a metric $K = E^T E$ such that the pseudo condition number of $AQ^\dagger A^T$ is minimized. This is achieved by finding an E to

$$\text{minimize} \quad \frac{\lambda_{\max}(EAQ^\dagger A^T E^T)}{\lambda_{\min>0}(EAQ^\dagger A^T E^T)}$$

where $\lambda_{\min>0}$ denotes the smallest non-zero eigenvalue. (See [26, Section 6] for methods to achieve this exactly and approximately.) This heuristic reduces to the optimal metric choice in the case where linear convergence is achieved. The γ-parameter is also chosen in accordance with the above reasoning and Corollary 2 as $\gamma = \frac{1}{\sqrt{\lambda_{\max}(EAQ^\dagger A^T E^T)\lambda_{\min>0}(EAQ^\dagger A^T E^T)}}$.

If instead $\hat{f}$ in (32) is the indicator function of an affine subspace, i.e., $\hat{f} = \iota_{Lx=b}$ and Q is strongly convex on the null-space of L. Then $d_1(\mu) = \frac{1}{2}\mu^T AP_{11}A^T\mu + \xi^T\mu + \chi + c^T\mu$ where $\xi \in \mathbb{R}^n$, $\chi \in \mathbb{R}$, and $\begin{bmatrix} Q & L^T \\ L & 0 \end{bmatrix}^{-1} = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix}$. Then we can choose metric by minimizing the pseudo condition number of $AP_{11}A^T$ (which is the Hessian of d_1) and select γ as $\gamma = \frac{1}{\sqrt{\lambda_{\max}(EAP_{11}A^T E^T)\lambda_{\min>0}(EAP_{11}A^T E^T)}}$.

7 Numerical example

7.1 A problem with linear convergence

We consider a weighted Lasso minimization problem formulation with inspiration from [10] of the form

$$\text{minimize} \quad \tfrac{1}{2}\|Ax - b\|_2^2 + \|Wx\|_1 \tag{33}$$

with $x \in \mathbb{R}^{200}$, the data matrix $A \in \mathbb{R}^{300 \times 200}$ is a sparse with on average 10 nonzero entries per row, and $b \in \mathbb{R}^{300}$. Each nonzero element in A and b is drawn from a Gaussian distribution with zero mean and unit variance. Further, $W \in \mathbb{R}^{200 \times 200}$ is a diagonal matrix with positive diagonal elements drawn from a uniform distribution on the interval $[0, 1]$.

We solve the problem using ADMM in the standard Euclidean setting and in the optimal metric setting according to Section 6. We use the optimal $\alpha = 1$ and different parameters γ. The γ parameter is varied in the range $[10^{-2}\gamma^\star, 10^2\gamma^\star]$ where $\gamma^\star$ is the theoretically optimal γ in the respective settings (Euclidean and
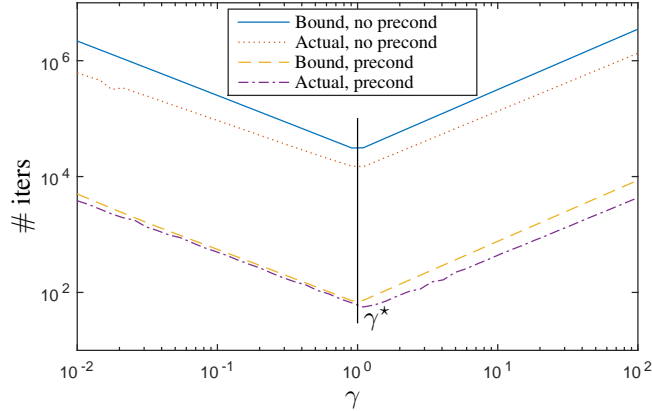
Figure 2: Theoretical bounds on and actual number of iterations to achieve a prespecified relative accuracy of 10^{-5} when solving (33) using ADMM with $\alpha = 1$ for different γ. We present results with and without preconditioning (metric selection).

optimal metric). In Figure 2, the actual number of iterations needed to obtain a specific accuracy (a relative tolerance of 10^{-5}) as well as the theoretical worst case number of iterations are plotted, both for the Euclidean and optimal metric setting.

Figure 2 reveals that, for this particular example, the actual numbers of iterations are fairly consistent with the iteration bounds. We also see that there is a lot to gain by running the algorithm with an appropriately chosen metric. Also the optimal parameters $\gamma^\star$ give close to optimal performance in both settings.

7.2 A problem without linear convergence

Here, we evaluate the heuristic metric and parameter selection method by applying ADMM to the (small-scale) aircraft control problem from [35, 3]. As in [3], the continuous time model from [35] is sampled using zero-order hold every 0.05 s. The system has four states $x = (x_1, x_2, x_3, x_4)$, two outputs $y = (y_1, y_2)$, two inputs $u = (u_1, u_2)$, and obeys the following dynamics

$$x(t+1) = \begin{bmatrix} 0.999 & -3.008 & -0.113 & -1.608 \\ -0.000 & 0.986 & 0.048 & 0.000 \\ 0.000 & 2.083 & 1.009 & -0.000 \\ 0.000 & 0.053 & 0.050 & 1.000 \end{bmatrix} x(t) + \begin{bmatrix} -0.080 & -0.635 \\ -0.029 & -0.014 \\ -0.868 & -0.092 \\ -0.022 & -0.002 \end{bmatrix} u(t),$$

$$y(t) = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} x(t).$$

The system is unstable, the magnitude of the largest eigenvalue of the dynamics matrix is 1.313. The outputs are the attack and pitch angles, while the inputs
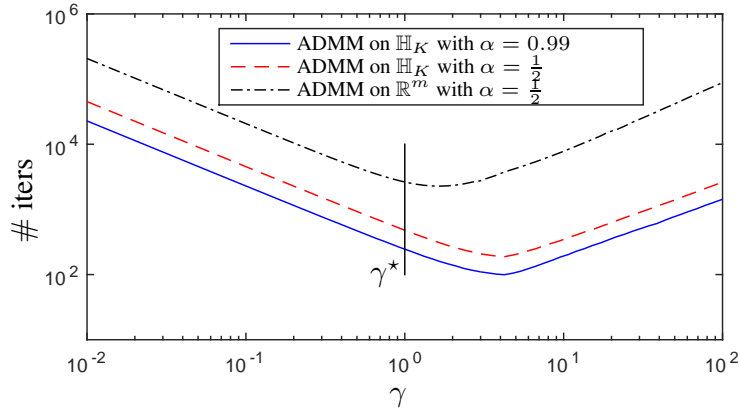
Figure 3: Average number of iterations for different γ-values, different metrics, and different relaxations α.

are the elevator and flaperon angles. The inputs are physically constrained to satisfy $|u_i| \leq 25°$, $i = 1, 2$. The outputs are soft constrained and modeled using the piece-wise linear cost function

$$h(y, l, u, s) = \begin{cases} 0 & \text{if } l \leq y \leq u \\ s(y - u) & \text{if } y \geq u \\ s(l - y) & \text{if } y \leq l \end{cases}$$

Especially, the first output is penalized using $h(y_1, -0.5, 0.5, 10^6)$ and the second is penalized using $h(y_2, -100, 100, 10^6)$. The states and inputs are penalized using

$$\ell(x, u, s) = \tfrac{1}{2}\big((x - x_r)^T Q(x - x_r) + u^T R u\big)$$

where x_r is a reference, $Q = \text{diag}(0, 10^2, 0, 10^2)$, and $R = 10^{-2}I$. Further, the terminal cost is Q, and the control and prediction horizons are $N = 10$. The numerical data in Figure 3 is obtained by following a reference trajectory on the output. The objective is to change the pitch angle from $0°$ to $10°$ and then back to $0°$ while the angle of attack satisfies the (soft) output constraints $-0.5° \leq y_1 \leq 0.5°$. The constraints on the angle of attack limit the rate on how fast the pitch angle can be changed. By stacking vectors and forming appropriate matrices, the full optimization problem can be written on the form

$$\min \quad \underbrace{\tfrac{1}{2}z^T Q z + r_t^T z + I_{Lz = bx_t}(z)}_{f(z)} + \underbrace{\sum_{i=1}^{m} h(z_i', \underline{d}_i, \bar{d}_i, 10^6)}_{g(z')}$$

$$\text{s. t.} \quad Cz = z'$$

where x_t and r_t may change from one sampling instant to the next. This fits the optimization problem formulation discussed in Section 6.1. For this problem all items (i)–(iii) violate the assumptions that guarantee linear convergence.

Since the numerical example treated here is a model predictive control application, we can spend much computational effort offline to compute a metric that will be used in all samples in the online controller. We compute a diagonal metric $K = E^T E$, where E minimizes the pseudo condition number of $ECP_{11}C^T E^T$ and P_{11} is implicitly defined by $\begin{bmatrix} Q & L^T \\ L & 0 \end{bmatrix}^{-1} = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix}$ (as suggested in Section 6.1). This matrix $K = E^T E$ defines the space $\mathbb{H}_K$ on which the algorithm is applied.

In Figure 3, the performance of ADMM when applied on $\mathbb{H}_K$ with relaxations $\alpha = \frac{1}{2}$ and $\alpha = 0.99$, and ADMM applied on $\mathbb{R}^m$ with $\alpha = \frac{1}{2}$ is shown. In this particular example, improvements of about one order of magnitude are achieved when applied on $\mathbb{H}_K$ compared to when applied on $\mathbb{R}^m$. Figure 3 also shows that ADMM with over-relaxation performs better than standard ADMM. The proposed γ-parameter selection is denoted by $\gamma^\star$ in Figure 3 (E or C is scaled to get $\gamma^\star = 1$ for all examples). Figure 3 shows that $\gamma^\star$ does not coincide with the empirically found best γ, but still gives a reasonable choice of γ in all cases.

8   Conclusions

We have shown tight global linear convergence rate bounds for Douglas-Rachford splitting and ADMM. Based on these results, we have presented methods to select metric and algorithm parameters for these methods. We have also provided numerical examples to evaluate the proposed metric and parameter selection methods for ADMM.

9   Acknowledgments

References

[1] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces.* Springer, 2011.

[2] H. H. Bauschke, J. Y. B. Cruz, T. T. A. Nghia, H. M. Phan, and X. Wang. The rate of linear convergence of the Douglas-Rachford algorithm for subspaces is the cosine of the Friedrichs angle. *Journal of Approximation Theory*, 185(0):63–79, 2014.

[3] A. Bemporad, A. Casavola, and E. Mosca. Nonlinear control of constrained linear systems via predictive reference management. *IEEE Transactions on Automatic Control*, 42(3):340–349, 1997.

[4] M. Benzi. Preconditioning techniques for large linear systems: A survey. *Journal of Computational Physics*, 182(2):418–477, 2002.

[5] D. Boley. Local linear convergence of the alternating direction method of multipliers on quadratic or linear programs. *SIAM Journal on Optimization*, 23(4):2183–2207, 2013.

[6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[7] J. H. Bramble, J. E. Pasciak, and A. T. Vassilev. Analysis of the inexact Uzawa algorithm for saddle point problems. *SIAM Journal on Numerical Analysis*, 34(3):1072–1092, 1997.

[8] K. Bredies and H. Sun. Preconditioned Douglas–Rachford splitting methods for convex-concave saddle-point problems. *SIAM Journal on Numerical Analysis*, 53(1):421–444, 2015.

[9] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, Feb 2006.

[10] E. J. Candes, M. B. Wakin, and S. Boyd. Enhancing sparsity by reweighted l1 minimization. *Journal of Fourier Analysis and Applications*, 14(5):877–905, Dec 2008. Special issue on sparsity.

[11] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.

[12] P. L. Combettes and J-C. Pesquet. Proximal splitting methods in signal processing. In H. H. Bauschke, R. S. Burachik, P. L. Combettes, V. Elser, D. R. Luke, and H. Wolkowicz, editors, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, volume 49 of *Springer Optimization and Its Applications*, pages 185–212. Springer New York, 2011.

[13] D. Davis and W. Yin. Convergence rate analysis of several splitting schemes. Available `http://arxiv.org/abs/1406.4834`, August 2014.

[14] D. Davis and W. Yin. Faster convergence rates of relaxed Peaceman-Rachford and ADMM under regularity assumptions. Available: `http://arxiv.org/abs/1407.5210`, July 2014.

[15] L. Demanet and X. Zhang. Eventual linear convergence of the Douglas-Rachford iteration for basis pursuit. Available: `http://arxiv.org/abs/1301.0542`, May 2013.

[16] W. Deng and W. Yin. On the global and linear convergence of the generalized alternating direction method of multipliers. Technical Report CAAM 12-14, Rice University, 2012.

[17] J. Douglas and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Trans. Amer. Math. Soc.*, 82:421–439, 1956.

[18] J. Eckstein. *Splitting methods for monotone operators with applications to parallel optimization.* PhD thesis, MIT, 1989.

[19] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1):293–318, 1992.

[20] D. Gabay. Applications of the method of multipliers to variational inequalities. In M. Fortin and R. Glowinski, editors, *Augmented Lagrangian Methods: Applications to the Solution of Boundary-Value Problems.* North-Holland: Amsterdam, 1983.

[21] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers and Mathematics with Applications*, 2(1):17–40, 1976.

[22] E. Ghadimi, A. Teixeira, I. Shames, and M. Johansson. Optimal parameter selection for the alternating direction method of multipliers (ADMM): Quadratic problems. *IEEE Transactions on Automatic Control*, 60(3):644–658, March 2015.

[23] P. Giselsson. Tight global linear convergence rate bounds for Douglas-Rachford splitting. 2015. Submitted. Available: `http://arxiv.org/abs/1506.01556`.

[24] P. Giselsson. Tight linear convergence rate bounds for Douglas-Rachford splitting and ADMM. In *Proceedings of 54th Conference on Decision and Control*, Osaka, Japan, Dec 2015.

[25] P. Giselsson and S. Boyd. Diagonal scaling in Douglas-Rachford splitting and ADMM. In *53rd IEEE Conference on Decision and Control*, pages 5033–5039, Los Angeles, CA, December 2014.

[26] P. Giselsson and S. Boyd. Metric selection in fast dual forward-backward splitting. *Automatica*, 62:1–10, 2015.

[27] P. Giselsson, M. Fält, and S. Boyd. Line search for averaged operator iteration. Available: `http://arxiv.org/abs/1603.06772`, 2016.

[28] R. Glowinski and A. Marroco. Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problémes de dirichlet non linéaires. *ESAIM: Mathematical Modelling and*

Numerical Analysis - Modélisation Mathématique et Analyse Numérique, 9:41–76, 1975.

[29] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference and prediction.* Springer, 2nd edition, 2009.

[30] B. He and X. Yuan. On the $o(1/n)$ convergence rate of the Douglas-Rachford alternating direction method. *SIAM Journal on Numerical Analysis*, 50(2):700–709, 2012.

[31] R. Hesse, D. R. Luke, and P. Neumann. Alternating projections and Douglas-Rachford for sparse affine feasibility. *IEEE Transactions on Signal Processing*, 62(18):4868–4881, September 2014.

[32] M. Hong and Z.-Q. Luo. On the linear convergence of the alternating direction method of multipliers. Available: `http://arxiv.org/abs/1208.3922`, March 2013.

[33] Q. Hu and J. Zou. Nonlinear inexact Uzawa algorithms for linear and nonlinear saddle-point problems. *SIAM Journal on Optimization*, 16(3):798–825, 2006.

[34] F. Iutzeler, P. Bianchi, P. Ciblat, and W. Hachem. Explicit convergence rate of a distributed alternating direction method of multipliers. Available: `http://arxiv.org/abs/1312.1085`, December 2013.

[35] P. Kapasouris, M. Athans, and G. Stein. Design of feedback control systems for unstable plants with saturating actuators. In *Proceedings of the IFAC Symposium on Nonlinear Control System Design*, pages 302–307. Pergamon Press, 1990.

[36] P. L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.

[37] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007.

[38] R. Nishihara, L. Lessard, B. Recht, A. Packard, and M. Jordan. A general analysis of the convergence of ADMM. February 2015. Available: `http://arxiv.org/abs/1502.02009`.

[39] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2014.

[40] P. Patrinos, L. Stella, and A. Bemporad. Douglas-Rachford splitting: Complexity estimates and accelerated variants. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, Los Angeles, CA, December 2014.

[41] D. W. Peaceman and H. H. Rachford. The numerical solution of parabolic and elliptic differential equations. *Journal of the Society for Industrial and Applied Mathematics*, 3(1):28–41, 1955.

[42] H. M. Phan. Linear convergence of the Douglas-Rachford method for two closed sets. Available: `http://arxiv.org/abs/1401.6509`, October 2014.

[43] T. Pock and A. Chambolle. Diagonal preconditioning for first order primal-dual algorithms in convex optimization. In *2011 IEEE International Conference on Computer Vision (ICCV)*, pages 1762–1769, Nov 2011.

[44] A. Raghunathan and S. Di Cairano. ADMM for convex quadratic programs: Linear convergence and infeasibility detection. November 2014. Available: `http://arxiv.org/abs/1411.7288`.

[45] J. B. Rawlings and D. Q. Mayne. *Model Predictive Control: Theory and Design.* Nob Hill Publishing, Madison, WI, 2009.

[46] R. T. Rockafellar and R. J-B. Wets. *Variational Analysis.* Springer, Berlin, 1998.

[47] W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin. On the linear convergence of the ADMM in decentralized consensus optimization. *IEEE Transactions on Signal Processing*, 62(7):1750–1761, April 2014.

[48] R Tibshirani. Regression shrinkage and selection via the lasso. *J. Royal. Statist. Soc B.*, 58(1):267–288, 1996.

A   A lemma

We state the following lemma that is used in several of the proofs in this appendix.

Lemma 1 *Let $\gamma, \beta, \sigma \in (0, \infty)$ and $\beta \geq \sigma$. Then*

$$\max(\tfrac{1-\gamma\sigma}{1+\gamma\sigma}, \tfrac{\gamma\beta-1}{1+\gamma\beta}) = \begin{cases} \frac{1-\gamma\sigma}{1+\gamma\sigma} & \text{if } \gamma \leq \frac{1}{\sqrt{\beta\sigma}} \\ \frac{\gamma\beta-1}{1+\gamma\beta} & \text{if } \gamma \geq \frac{1}{\sqrt{\beta\sigma}} \end{cases} \tag{34}$$

Further $\max(\tfrac{1-\gamma\sigma}{1+\gamma\sigma}, \tfrac{\gamma\beta-1}{1+\gamma\beta}) \in [0, 1)$

Proof.   Let $\psi(x) = \frac{1-x}{1+x}$ for $x \geq 0$. This implies $\psi'(x) = \frac{-2}{(x+1)^2}$, i.e., ψ is (strictly) monotonically deceasing for $x \geq 0$. So for $x \leq 1/y$ and $x \geq 0$, we have

$$\psi(x) = \tfrac{1-x}{1+x} \geq \tfrac{1-1/y}{1+1/y} = -\tfrac{1-y}{y+1} = -\psi(y).$$

Similarly if $x \geq 1/y$ and $x \geq 0$, we have

$$\psi(x) = \tfrac{1-x}{1+x} \leq \tfrac{1-1/y}{1+1/y} = -\tfrac{1-y}{y+1} = -\psi(y).$$

So for $x \geq 0$, $\psi(x) \leq -\psi(y)$ if and only if $x \geq 1/y$. Therefore

$$
\begin{aligned}
\max(\tfrac{1-\gamma\sigma}{1+\gamma\sigma}, \tfrac{\gamma\beta-1}{1+\gamma\beta}) &= \max(\psi(\gamma\sigma)), -\psi(\gamma\beta)) \\
&= \begin{cases} \psi(\gamma\sigma) & \text{if } 0 < \gamma \leq \frac{1}{\sqrt{\beta\sigma}} \\ -\psi(\gamma\beta) & \text{if } \gamma \geq \frac{1}{\sqrt{\beta\sigma}} \end{cases} \\
&= \begin{cases} \frac{1-\gamma\sigma}{1+\gamma\sigma} & \text{if } 0 < \gamma \leq \frac{1}{\sqrt{\beta\sigma}} \\ \frac{\gamma\beta-1}{1+\gamma\beta} & \text{if } \gamma \geq \frac{1}{\sqrt{\beta\sigma}} \end{cases}
\end{aligned}
\tag{35}
$$

Finally, since $\psi(0) = 1$ and $\psi(1) = 0$, and ψ is strictly monotonically decreasing, $\psi(x) \in [0,1)$ for $x \in (0,1]$. Also $\psi(x) \to -1$ as $x \to \infty$, so $\psi(x) \in (-1,0]$ for $x \geq 1$. Therefore, since $\beta \geq \sigma$, $\psi(\gamma\sigma) \in [0,1)$ if $0 \leq \gamma \leq \frac{1}{\sigma}$ and $-\psi(\gamma\beta) \in [0,1)$ if $\gamma \geq \frac{1}{\beta}$. Since $\frac{1}{\beta} \leq \frac{1}{\sigma}$, (35) implies that $\max(\frac{1-\gamma\sigma}{1+\gamma\sigma}, \frac{\gamma\beta-1}{1+\gamma\beta}) = \max(\psi(\gamma\sigma), -\psi(\gamma\beta)) \in [0,1)$. $\square$

B   Proof to Proposition 2

Proof. Since f is σ-strongly convex and β-smooth and since $\gamma \in (0,\infty)$, f_γ is $(1+\gamma\sigma)$-strongly convex and $(1+\gamma\beta)$-smooth. Therefore [1, Theorem 18.15] and [1, Theorem 13.32] directly imply that f_γ^* is $\frac{1}{1+\gamma\sigma}$-smooth and $\frac{1}{1+\gamma\beta}$-strongly convex. From the smoothness definition in Definition 6, we get that

$$
\tfrac{1}{2(1+\gamma\sigma)}\|\cdot\|^2 - f_\gamma^* = \left(\tfrac{1}{2(1+\gamma\sigma)} - \tfrac{1}{2(1+\gamma\beta)}\right)\|\cdot\|^2 - (f_\gamma^* - \tfrac{1}{2(1+\gamma\beta)}\|\cdot\|^2)
\tag{36}
$$

is convex. Further, Definition 5 implies that $f_\gamma^* - \frac{1}{2(1+\gamma\beta)}\|\cdot\|^2$ is convex (since f_γ^* is $\frac{1}{1+\gamma\beta}$-strongly convex), and therefore (36) is the definition of $\frac{1}{2(1+\gamma\sigma)} - \frac{1}{2(1+\gamma\beta)}$-smoothness of $f_\gamma^* - \frac{1}{2(1+\gamma\beta)}\|\cdot\|^2$.

Now, let $\beta = \sigma$, then $f_\gamma^* - \frac{1}{2(1+\gamma\beta)}\|\cdot\|^2$ is 0-smooth, or equivalently by applying [1, Theorem 18.15], $\nabla f_\gamma^* - \frac{1}{1+\gamma\beta}\mathrm{Id} = \mathrm{prox}_{\gamma f} - \frac{1}{1+\gamma\beta}\mathrm{Id}$ is 0-Lipschitz. Let $\beta > \sigma$, then [1, Theorem 18.15] implies that $\frac{1}{2(1+\gamma\sigma)} - \frac{1}{2(1+\gamma\beta)}$-smoothness of $f_\gamma^* - \frac{1}{2(1+\gamma\beta)}\|\cdot\|^2$ is equivalent to $\frac{1}{\frac{1}{1+\gamma\sigma} - \frac{1}{1+\gamma\beta}}$-cocoercivity of $\nabla f_\gamma^* - \frac{1}{1+\gamma\beta}\mathrm{Id} = \mathrm{prox}_{\gamma f} - \frac{1}{1+\gamma\beta}\mathrm{Id}$. This concludes the proof. $\square$

C   Proof to Theorem 1

Proof. First assume that $\beta > \sigma$. By Proposition 2, $\mathrm{prox}_{\gamma f} - \frac{1}{1+\gamma\beta}\mathrm{Id}$ is $\frac{1}{\frac{1}{1+\gamma\sigma} - \frac{1}{1+\gamma\beta}}$-cocoercive. Therefore, according to (4):

$$
\mathrm{prox}_{\gamma f} - \tfrac{1}{1+\gamma\beta}\mathrm{Id} = \tfrac{1}{2}(\tfrac{1}{1+\gamma\sigma} - \tfrac{1}{1+\gamma\beta})(\mathrm{Id} + C)
\tag{37}
$$

for some nonexpansive mapping C.

When $\beta = \sigma$, Proposition 2 implies that $\text{prox}_{\gamma f} = \frac{1}{1+\gamma\beta}\text{Id}$. Therefore, also in this case, it can be represented by (37) (since the right hand side is 0). We get

$$
\begin{aligned}
\|R_{\gamma f}x - R_{\gamma f}y\| &= \|(2\text{prox}_{\gamma f} - \text{Id})x - (2\text{prox}_{\gamma g} - \text{Id})y\| \\
&= \|(\tfrac{1}{1+\gamma\sigma} - \tfrac{1}{1+\gamma\beta} + \tfrac{2}{1+\gamma\beta} - 1)(x-y) \\
&\quad + (\tfrac{1}{1+\gamma\sigma} - \tfrac{1}{1+\gamma\beta})(Cx - Cy)\| \\
&= \|(\tfrac{1}{1+\gamma\sigma} + \tfrac{1}{1+\gamma\beta} - 1)(x-y) \\
&\quad + (\tfrac{1}{1+\gamma\sigma} - \tfrac{1}{1+\gamma\beta})(Cx - Cy)\| \\
&\leq (|\tfrac{1}{1+\gamma\sigma} + \tfrac{1}{1+\gamma\beta} - 1| + \tfrac{1}{1+\gamma\sigma} - \tfrac{1}{1+\gamma\beta})\|x-y\|.
\end{aligned}
$$

So $R_{\gamma f}$ is Lipschitz continuous with constant $|\frac{1}{1+\gamma\sigma} + \frac{1}{1+\gamma\beta} - 1| + \frac{1}{1+\gamma\sigma} - \frac{1}{1+\gamma\beta}$. The kink in the absolute value term is when

$$
\begin{aligned}
0 &= \tfrac{1}{1+\gamma\sigma} + \tfrac{1}{1+\gamma\beta} - 1 \\
\Leftrightarrow \qquad 0 &= 1 + \gamma\beta + 1 + \gamma\sigma - (1+\gamma\sigma)(1+\gamma\beta) \\
\Leftrightarrow \qquad 0 &= 1 - \gamma^2\sigma\beta,
\end{aligned}
$$

i.e, when $\gamma = \frac{1}{\sqrt{\beta\sigma}}$. For $\gamma \in (0, \frac{1}{\sqrt{\beta\sigma}}]$, the expression in the absolute value is positive, and the Lipschitz constant is

$$
\tfrac{1}{1+\gamma\sigma} + \tfrac{1}{1+\gamma\beta} - 1 + \tfrac{1}{1+\gamma\sigma} - \tfrac{1}{1+\gamma\beta} = \tfrac{2}{1+\gamma\sigma} - 1 = \tfrac{1-\gamma\sigma}{1+\gamma\sigma}.
$$

For $\gamma \in [\frac{1}{\sqrt{\beta\sigma}}, \infty)$, the expression in the absolute value is negative, and the Lipschitz constant is

$$
1 - \tfrac{1}{1+\gamma\sigma} - \tfrac{1}{1+\gamma\beta} + \tfrac{1}{1+\gamma\sigma} - \tfrac{1}{1+\gamma\beta} = 1 - \tfrac{2}{1+\gamma\beta} = \tfrac{\gamma\beta-1}{1+\gamma\beta}.
$$

Lemma 1 in Appendix A implies that the Lipschitz constant therefore can be written as

$$
\max(\tfrac{1-\gamma\sigma}{1+\gamma\sigma}, \tfrac{\gamma\beta-1}{\gamma\beta+1})
$$

and that $\max(\frac{1-\gamma\sigma}{1+\gamma\sigma}, \frac{\gamma\beta-1}{\gamma\beta+1}) \in [0,1)$, i.e. that $R_{\gamma f}$ is a contraction. This concludes the proof.

$\square$

D Proof to Theorem 2

Proof. Since $\gamma \in (0,\infty)$, [1, Corollary 23.10] implies that $R_{\gamma g}$ is nonexpansive and by Theorem 1, $R_{\gamma f}$ is $\delta = \max(\frac{1-\gamma\sigma}{1+\gamma\sigma}, \frac{\gamma\beta-1}{\gamma\beta+1})$-contractive. Therefore the composition $R_{\gamma g}R_{\gamma f}$ is also δ-contractive since

$$
\|R_{\gamma g}R_{\gamma f}z_1 - R_{\gamma g}R_{\gamma f}z_2\| \leq \|R_{\gamma f}z_1 - R_{\gamma f}z_2\| \leq \delta\|z_1 - z_2\| \tag{38}
$$

for any $z_1, z_2 \in \mathcal{H}$. Now, let $T = (1 - \alpha)I + \alpha R_{\gamma g} R_{\gamma f}$ be the Douglas-Rachford operator in (9). Since $\bar{z}$ is a fixed-point to $R_{\gamma g} R_{\gamma f}$ it is also a fixed-point to T, i.e., $\bar{z} = T\bar{z}$. Thus

$$
\begin{aligned}
\|z^{k+1} - \bar{z}\| &= \|Tz^k - T\bar{z}\|^2 \\
&= \|(1 - \alpha)(z^k - \bar{z}) + \alpha(R_{\gamma g} R_{\gamma f} z^k - R_{\gamma g} R_{\gamma f} \bar{z})\| \\
&\leq |1 - \alpha|\|z^k - \bar{z}\| + \alpha\|R_{\gamma g} R_{\gamma f} z^k - R_{\gamma g} R_{\gamma f} \bar{z}\| \\
&\leq (|1 - \alpha| + \alpha\delta)\|z^k - \bar{z}\| \\
&= \left(|1 - \alpha| + \alpha \max(\tfrac{1-\gamma\sigma}{1+\gamma\sigma}, \tfrac{\gamma\beta-1}{\gamma\beta+1})\right)\|z^k - \bar{z}\|
\end{aligned}
$$

where (38) is used in the second inequality. For any $\gamma \in (0, \infty)$, Lemma 1 says that $\delta = \max(\tfrac{1-\gamma\sigma}{1+\gamma\sigma}, \tfrac{\gamma\beta-1}{\gamma\beta+1}) \in [0, 1)$ and straightforward manipulations show that the factor

$$
|1 - \alpha| + \alpha\delta < 1
$$

if and only if $\alpha \in (0, \frac{2}{1+\delta})$. This concludes the proof. $\qquad\square$

E   Proof of Corollary 1

Proof.   Since f is σ-strongly convex, $f_\gamma = \gamma f + \frac{1}{2}\|\cdot\|^2$ is $(1 + \gamma\sigma)$-strongly convex, and ∇f_γ^* is $\frac{1}{1+\gamma\sigma}$-Lipschitz continuous, see [1, Proposition 18.15].

Now, recall from Proposition 1 that $\nabla f_\gamma^*(z) = \mathrm{prox}_{\gamma f}(z)$. Further, let $\bar{z}$ be a fixed-point to $R_{\gamma g} R_{\gamma f}$, i.e., $\bar{z}$ satisfies $\bar{z} = R_{\gamma g} R_{\gamma f} \bar{z}$ and $x^\star = \mathrm{prox}_{\gamma f}(\bar{z})$, see (8). Therefore

$$
\begin{aligned}
\|x^{k+1} - x^\star\| &= \|\mathrm{prox}_{\gamma f}(z^{k+1}) - \mathrm{prox}_{\gamma f}(\bar{z})\| \\
&\leq \tfrac{1}{1+\gamma\sigma}\|z^{k+1} - \bar{z}\| \\
&\leq \tfrac{1}{1+\gamma\sigma}(|1 - \alpha| + \delta\alpha)\|z^k - \bar{z}\| \\
&\leq \tfrac{1}{1+\gamma\sigma}(|1 - \alpha| + \delta\alpha)^{k+1}\|z^0 - \bar{z}\|.
\end{aligned}
$$

where the second and third inequalities come from Theorem 2. This concludes the proof. $\qquad\square$

F   Proof of Theorem 3

Let $z^0 = (1, 0)$ or $z^0 = (0, 1)$. Then separability of $R_{\gamma f}$ (20) and $R_{\gamma g_1} = \mathrm{Id}$ implies that the Douglas-Rachford algorithm (9) for minimizing $f + g_1$ is

$$
z^{k+1} = \left((1 - \alpha) + \alpha\tfrac{1-\gamma\lambda_i}{1+\gamma\lambda_i}\right)z^0 \tag{39}
$$

with $\lambda_i = \beta$ if $z^0 = (1,0)$ and $\lambda_i = \sigma$ if $z^0 = (0,1)$. The exact rate is

$$\left|1 - \alpha + \alpha \frac{1-\gamma\lambda_i}{1+\gamma\lambda_i}\right|. \tag{40}$$

Similarly, when minimizing $f + g_2$ we get (since $R_{\gamma g_2} = -\mathrm{Id}$ and $R_{\gamma f}$ is linear)

$$z^{k+1} = \left((1-\alpha) - \alpha\frac{1-\gamma\lambda_i}{1+\gamma\lambda_i}\right)z^0 \tag{41}$$

with exact rate

$$\left|1 - \alpha - \alpha\frac{1-\gamma\lambda_i}{1+\gamma\lambda_i}\right|. \tag{42}$$

Let's consider the following four cases

(i) $\alpha \le 1$, $\gamma \in (0, \frac{1}{\sqrt{\beta\sigma}}]$, $g = g_1$, $z_0 = (0,1) \Rightarrow \lambda_i = \sigma$

(ii) $\alpha \ge 1$, $\gamma \ge \frac{1}{\sqrt{\beta\sigma}}$, $g = g_1$, $z_0 = (1,0) \Rightarrow \lambda_i = \beta$

(iii) $\alpha \le 1$, $\gamma \ge \frac{1}{\sqrt{\beta\sigma}}$, $g = g_2$, $z_0 = (1,0) \Rightarrow \lambda_i = \beta$

(iv) $\alpha \ge 1$, $\gamma \in (0, \frac{1}{\sqrt{\beta\sigma}}]$, $g = g_2$, $z_0 = (0,1) \Rightarrow \lambda_i = \sigma$

where λ_i refers to the λ_i in (39)-(42).

Using Lemma 1 in Appendix A, it is straightforward to verify that the exact rate factors (40) for (i) and (ii) and (42) for (iii) and (iv) equals the upper bound on the rate in Theorem 2 for $\alpha \in (0, \frac{2}{1+\delta})$ with δ in (14).

It is also straightforward to verify that for $\alpha \notin (0, \frac{2}{1+\delta})$, the exact rate is greater than or equal to one, so it does not converge.

We show this for the first case, and leave the three other very similar cases for the reader to verify. For case (i) with $\alpha \in (0, 1]$, Lemma 1 implies that the rate (40) is

$$\left|1 - \alpha + \alpha\frac{1-\gamma\sigma}{1+\gamma\sigma}\right| = 1 - \alpha + \alpha\max(\frac{1-\gamma\sigma}{1+\gamma\sigma}, \frac{\gamma\beta-1}{1+\gamma\beta})$$
$$= |1 - \alpha| + \alpha\max(\frac{1-\gamma\sigma}{1+\gamma\sigma}, \frac{\gamma\beta-1}{1+\gamma\beta})$$

which is the optimal rate in Theorem 2. If instead $\alpha \le 0$, we get

$$\left|1 - \alpha + \alpha\frac{1-\gamma\sigma}{1+\gamma\sigma}\right| = 1 + |\alpha|(1 - \max(\frac{1-\gamma\sigma}{1+\gamma\sigma}, \frac{\gamma\beta-1}{1+\gamma\beta})) \ge 1$$

since $\max(\frac{1-\gamma\sigma}{1+\gamma\sigma}, \frac{\gamma\beta-1}{1+\gamma\beta}) \in [0, 1)$ by Lemma 1. So it does not converge to the solution.

Repeating similar arguments for the three other cases gives the result.