



UNIVERSITÀ
DEGLI STUDI
FIRENZE

FLORE

Repository istituzionale dell'Università degli Studi di Firenze

Updating constraint preconditioners for KKT systems in quadratic programming via low-rank corrections

Questa è la Versione finale referata (Post print/Accepted manuscript) della seguente pubblicazione:

Original Citation:

Updating constraint preconditioners for KKT systems in quadratic programming via low-rank corrections / Bellavia, Stefania; De Simone, Valentina; di Serafino Daniela; Morini, Benedetta. - In: SIAM JOURNAL ON OPTIMIZATION. - ISSN 1052-6234. - STAMPA. - 25:(2015), pp. 1787-1808. [10.1137/130947155]

Availability:

The webpage <https://hdl.handle.net/2158/1004217> of the repository was last updated on 2021-04-01T16:40:47Z

Published version:

DOI: 10.1137/130947155

Terms of use:

Open Access

La pubblicazione è resa disponibile sotto le norme e i termini della licenza di deposito, secondo quanto stabilito dalla Policy per l'accesso aperto dell'Università degli Studi di Firenze (<https://www.sba.unifi.it/upload/policy-oa-2016-1.pdf>)

Publisher copyright claim:

La data sopra indicata si riferisce all'ultimo aggiornamento della scheda del Repository FloRe - The above-mentioned date refers to the last update of the record in the Institutional Repository FloRe

(Article begins on next page)

On the update of constraint preconditioners for regularized KKT systems

Stefania Bellavia · Valentina De Simone ·
Daniela di Serafino · Benedetta Morini

Revised version – September 18, 2015

Abstract We address the problem of preconditioning sequences of regularized KKT systems, such as those arising in interior point methods for convex quadratic programming. In this case, constraint preconditioners (CPs) are very effective and widely used; however, when solving large-scale problems, the computational cost for their factorization may be high, and techniques for approximating them appear as a convenient alternative. Here, given a block LDL^T factorization of the CP associated with a KKT matrix of the sequence, called seed matrix, we present a technique for updating the factorization and building inexact CPs for subsequent matrices of the sequence. We have recently proposed an updating procedure that performs a low-rank correction of the Schur complement of the (1,1) block of the CP for the seed matrix. Now we focus on KKT sequences with nonzero (2,2) blocks and make a step further, by enriching the low-rank correction of the Schur complement by an additional cheap update. The latter update takes into account information not included in the former one and expressed as a diagonal modification of the low-rank correction. Theoretical results and numerical experiments show that the new strategy can be more effective than the procedure based on the low-rank modification alone.

Keywords KKT systems · Constraint preconditioners · Matrix updates · Interior point methods

Work partially supported by INdAM-GNCS, 2015 Projects *Metodi di regolarizzazione per problemi di ottimizzazione e applicazioni* and *Metodi numerici per l'ottimizzazione non convessa o non differenziabile e applicazioni*.

Stefania Bellavia · Benedetta Morini

Dipartimento di Ingegneria Industriale, Università degli Studi di Firenze, viale Morgagni 40, I-50134 Firenze, Italy. E-mail: stefania.bellavia@unifi.it, benedetta.morini@unifi.it.

Valentina De Simone · Daniela di Serafino

Dipartimento di Matematica e Fisica, Seconda Università degli Studi di Napoli, viale A. Lincoln 5, I-81100 Caserta, Italy. E-mail: valentina.desimone@unina2.it, daniela.diserafino@unina2.it.

AMS subject classification 65F08 · 65F10 · 90C20

1 Introduction

We consider the problem of preconditioning a sequence of regularized KKT systems of the following form:

$$\begin{bmatrix} B_k & A^T \\ A & -\Theta_k \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}, \quad k = 1, 2, \dots, \quad (1)$$

where $B_k \in \mathbb{R}^{n \times n}$ is symmetric positive definite, $A \in \mathbb{R}^{m \times n}$ is full rank, $m \leq n$, and $\Theta_k \in \mathbb{R}^{m \times m}$ is diagonal positive semidefinite. We focus on problems where $\Theta_k \neq 0$ and its sparsity pattern does not change throughout the sequence. We further assume that the above linear systems are large and possibly sparse. Linear systems of this form arise, e.g., in interior point (IP) methods for convex quadratic programming problems [25, 34]:

$$\begin{aligned} & \text{minimize} \quad \frac{1}{2} x^T Q x + c^T x, \\ & \text{subject to} \quad A_1 x - s = b_1, \quad A_2 x = b_2, \quad x + v = u, \quad (x, s, v) \geq 0, \end{aligned} \quad (2)$$

where $A_1 \in \mathbb{R}^{m_1 \times n}$, $A_2 \in \mathbb{R}^{m_2 \times n}$, and s and v are slack variables, used to transform the inequality constraints $A_1 x \geq b_1$ and $x \leq u$ into equality constraints. In fact, the application of an IP method to problem (2) gives rise to a sequence of linear systems of the form (1), where

$$B_k = Q + \Phi_k, \quad \Phi_k = X_k^{-1} W_k + V_k^{-1} T_k, \quad \Theta_k = \begin{bmatrix} Y_k^{-1} S_k & 0 \\ 0 & 0 \end{bmatrix}, \quad A = \begin{bmatrix} A_1 \\ A_2 \end{bmatrix},$$

X_k , W_k , S_k , Y_k , V_k and T_k are suitable diagonal matrices, and $m = m_1 + m_2$. More precisely, letting (x_k, w_k) , (s_k, y_k) and (v_k, t_k) be the pairs of complementary variables of problem (2) evaluated at the k -th iteration, we have that X_k , W_k , S_k , Y_k , V_k and T_k are the diagonal matrices where the diagonal is equal to the corresponding lowercase vector, according to the usual IP notation. A nonzero block Θ_k arises whenever the problem has linear inequality constraints, i.e., $m_1 \neq 0$, as in this case Θ_k admits m_1 positive diagonal entries corresponding to the slack variables for the linear inequality constraints; furthermore, the number and the position of the nonzero entries of Θ_k does not change throughout the IP iterations. It is well known that the nonzero entries of Φ_k and Θ_k usually display a drastic difference of magnitude: some of them tend to zero while others go to infinity. The nonzero matrix Θ_k may also arise when regularized IP methods are applied to quadratic programming problems with no linear inequality constraints. In this case a diagonal positive definite (2,2) block appears and its form depends on the dual regularization adopted (see, e.g., [23, 26] and the references therein).

Typically, the solution of sequence (1) represents the major computational burden of the IP procedure. Therefore, if the linear systems are solved iteratively, the effectiveness and the efficiency of the preconditioners used affect

the overall performance of the IP method. Here we address the case where the systems are solved by an iterative linear solver and constraint preconditioners (CPs) are employed (see, e.g., [10, 11, 15, 19, 20, 24, 28, 31]).

Whenever the computational cost for building a CP for each system of the sequence is high, a preconditioner updating strategy can offer a tradeoff between cost and efficiency and can enhance the overall procedure for the solution of problem (1). Given a factorization of the preconditioner for a matrix of the sequence, such a strategy builds the preconditioners for some successive systems of the sequence by suitably modifying the available factorization to take into account the matrix changes. The aim is to form a preconditioner that is computationally cheaper than the one computed from scratch, though preserving efficiency in the solution of the linear systems.

A first proposal for an updating strategy for CP preconditioners has been presented in [5]. It relies on a low-rank modification of the factorized Schur complement of the (1,1) block, which has been designed by exploiting results from [1, 33]. Clearly, to keep the computational cost low, only small-rank changes are allowed. In this paper, in order to partially recover the information discarded by the low-rank correction, we propose to perform a further update. We note that a part of the discarded information can be seen as a diagonal positive semidefinite modification of the matrix resulting from the low-rank modification. Then we compute an approximate Cholesky factorization of the diagonally modified matrix by updating the Cholesky factorization of the preconditioner resulting from the low-rank correction. This step is accomplished by using a procedure in the framework given in [4]. Theoretical and numerical results show that the latter procedure can improve the preconditioner updating strategy in [5].

The paper is organized as follows. In Section 2 we provide the basis for describing our preconditioner updating procedure along with some theoretical results which will be exploited to analyze it. In Section 3 we first recall the updating procedure based on the low-rank correction of the Schur complement and then we present the new updating step, aimed at recovering part of the information discarded by the low-rank approach. In Section 4 we show some numerical results illustrating the behavior of the overall updating approach. We give some conclusions in Section 5.

Henceforth we use the following notations. We denote by $\|\cdot\|$ the matrix 2-norm. For any symmetric matrix M , we denote by $\lambda(M)$ any eigenvalue of M , and by $\lambda_{\min}(M)$ and $\lambda_{\max}(M)$ the minimum and maximum of these eigenvalues; furthermore, we use $\text{diag}(M)$ to indicate the diagonal matrix with the same diagonal entries as M . Finally, for any complex number z , we denote by $\mathcal{R}(z)$ and $\mathcal{I}(z)$ the real and imaginary parts of z .

2 Preliminaries

The updating strategy presented in this paper builds a preconditioner for a generic matrix of the sequence (1),

$$\mathcal{A}_k = \begin{bmatrix} B_k & A^T \\ A & -\Theta_k \end{bmatrix},$$

exploiting information from a preconditioner built for a previous matrix of the sequence, which is denoted by

$$\mathcal{A}_{seed} = \begin{bmatrix} B_{seed} & A^T \\ A & -\Theta_{seed} \end{bmatrix}. \quad (3)$$

We assume that a CP \mathcal{P}_{seed} is available for \mathcal{A}_{seed} , having the following form:

$$\mathcal{P}_{seed} = \begin{bmatrix} H_{seed} & A^T \\ A & -\Theta_{seed} \end{bmatrix}, \quad (4)$$

where H_{seed} is the following approximation to B_{seed} :

$$H_{seed} = \text{diag}(B_{seed}).$$

The effectiveness of this preconditioner in the context of IP methods is widely recognized by the optimization community, see, e.g., [6, 10, 11, 21, 28].

In order to simplify the notations, henceforth we denote H_{seed} by H . The application of \mathcal{P}_{seed} requires its factorization. We consider the following block LDL^T factorization:

$$\mathcal{P}_{seed} = \begin{bmatrix} I_n & 0 \\ AH^{-1} & I_m \end{bmatrix} \begin{bmatrix} H & 0 \\ 0 & -S_{seed} \end{bmatrix} \begin{bmatrix} I_n & H^{-1}A^T \\ 0 & I_m \end{bmatrix}, \quad (5)$$

where I_j is the identity matrix of dimension j and S_{seed} is the negative Schur complement of H in \mathcal{P}_{seed} , i.e.,

$$S_{seed} = AH^{-1}A^T + \Theta_{seed}. \quad (6)$$

We also assume that a Cholesky factorization of S_{seed} has been computed:

$$S_{seed} = L_{seed}D_{seed}L_{seed}^T, \quad (7)$$

where L_{seed} is unit lower triangular.

Now we consider a generic system of the sequence and we drop the iteration index k from \mathcal{A}_k , B_k and Θ_k in order to simplify the notation. Then, as with the previous definition of \mathcal{P}_{seed} , the CP for

$$\mathcal{A} = \begin{bmatrix} B & A^T \\ A & -\Theta \end{bmatrix}, \quad (8)$$

is given by

$$\mathcal{P}_{ex} = \begin{bmatrix} G & A^T \\ A & -\Theta \end{bmatrix} = \begin{bmatrix} I_n & 0 \\ AG^{-1} & I_m \end{bmatrix} \begin{bmatrix} G & 0 \\ 0 & -S \end{bmatrix} \begin{bmatrix} I_n & G^{-1}A^T \\ 0 & I_m \end{bmatrix}, \quad (9)$$

where

$$G = \text{diag}(B), \quad S = AG^{-1}A^T + \Theta.$$

Concerning the eigenvalue distribution of $\mathcal{P}_{ex}^{-1}\mathcal{A}$, there exists an eigenvalue at 1 with multiplicity $2m - p$, with $p = \text{rank}(\Theta)$, and $n - m + p$ real positive eigenvalues such that the better G approximates B the more clustered around 1 they are [18]. On the other hand, computing a Cholesky factorization of S may account for a large part of the computational cost for solving the KKT system, and hence replacing S with a computationally cheaper approximation of it, S_{inex} , may be convenient [5, 21, 29, 31]. The resulting *inexact* CP has the following form:

$$\mathcal{P}_{inex} = \begin{bmatrix} I_n & 0 \\ AG^{-1} & I_m \end{bmatrix} \begin{bmatrix} G & 0 \\ 0 & -S_{inex} \end{bmatrix} \begin{bmatrix} I_n & G^{-1}A^T \\ 0 & I_m \end{bmatrix}. \quad (10)$$

Approximations of CPs may be also obtained by replacing the constraint matrix A in (9) with a sparse approximations of it [9].

In our work we focus on \mathcal{P}_{inex} . The spectral analysis of $\mathcal{P}_{inex}^{-1}\mathcal{A}$ has been addressed in [7, 8, 32] and further refined in [5]. Here we report some results from [5], which will be exploited to design and analyse our updating procedure. We note that, although the convergence of Krylov solvers for systems with coefficient matrix $\mathcal{P}_{inex}^{-1}\mathcal{A}$ is not fully characterized by the spectrum of the matrix, in many practical cases it depends on the distribution of the eigenvalues. Therefore we are interested in providing bounds on the eigenvalues of $\mathcal{P}_{inex}^{-1}\mathcal{A}$. The bounds given in the next theorem highlight the dependence of the spectrum of $\mathcal{P}_{inex}^{-1}\mathcal{A}$ on that of $S_{inex}^{-1}S$ (see [5, Theorem 2.1]).

Theorem 1 *Let \mathcal{A} and \mathcal{P}_{inex} be the matrices in (8) and (10), and let λ and $[x^T, y^T]^T$ be an eigenvalue of $\mathcal{P}_{inex}^{-1}\mathcal{A}$ and a corresponding eigenvector. Let*

$$X = G^{-\frac{1}{2}}BG^{-\frac{1}{2}} \quad (11)$$

and suppose that $2I_n - X$ is positive definite. Let

$$\bar{\lambda} = \lambda_{\max}(S_{inex}^{-1}S) \max\{2 - \lambda_{\min}(X), 1\}, \quad (12)$$

$$\underline{\lambda} = \lambda_{\min}(S_{inex}^{-1}S) \min\{2 - \lambda_{\max}(X), 1\}. \quad (13)$$

Then, $\mathcal{P}_{inex}^{-1}\mathcal{A}$ has at most $2m$ eigenvalues with nonzero imaginary part, counting conjugates. Furthermore, if λ has nonzero imaginary part, then

$$\frac{1}{2}(\lambda_{\min}(X) + \underline{\lambda}) \leq \mathcal{R}(\lambda) \leq \frac{1}{2}(\lambda_{\max}(X) + \bar{\lambda}); \quad (14)$$

otherwise,

$$\min\{\lambda_{\min}(X), \underline{\lambda}\} \leq \lambda \leq \max\{\lambda_{\max}(X), \bar{\lambda}\}, \quad \text{for } y \neq 0, \quad (15)$$

$$\lambda_{\min}(X) \leq \lambda \leq \lambda_{\max}(X), \quad \text{for } y = 0. \quad (16)$$

Finally, the imaginary part of λ satisfies

$$|\mathcal{I}(\lambda)| \leq \sqrt{\lambda_{\max}(S_{inex}^{-1}AG^{-1}A^T)}\|I_n - X\|. \quad (17)$$

We note that $2I_n - X$ can be made positive semidefinite by scaling X in (11) so that its eigenvalues are not greater than 2. Furthermore, if B is diagonal, then $X = I_n$ and all the eigenvalues of $\mathcal{P}_{inex}^{-1}\mathcal{A}$ are real. In this case $\mathcal{P}_{inex}^{-1}\mathcal{A}$ has at least n unit eigenvalues, with n associated independent eigenvectors of the form $[x^T, 0^T]^T$, and the remaining eigenvalues lie in the interval $[\lambda_{\min}(S_{inex}^{-1}S), \lambda_{\max}(S_{inex}^{-1}S)]$ (see also [21]).

Finally, from the analysis carried out in [32, section 5] it follows that the number of unit eigenvalues of $\mathcal{P}_{inex}^{-1}\mathcal{A}$ depends also on the number of zero eigenvalues of $S - S_{inex}$. In particular, when $S - S_{inex}$ has l zero eigenvalues, if $S^{-1/2}AA^TS^{-1/2}$ has no unit eigenvalues then $\mathcal{P}_{inex}^{-1}\mathcal{A}$ has l unit eigenvalues with geometric multiplicity l .

3 The preconditioner updating strategy

Our preconditioner updating strategy is based on building S_{inex} by a suitable update of the seed Schur complement S_{seed} . In [5] we presented an updated preconditioner of the form (10), where S_{inex} is a low-rank modification, S_{lr} , of S_{seed} . Here we make a step further, improving the approximation of S provided by S_{lr} , in the case where Θ is a nonzero matrix. Therefore, our approach for building S_{inex} consists of two steps: the first computes S_{lr} through the procedure discussed in [5]; the second employs updating techniques discussed in [3, 4] to form an approximate factorization of $S_{lr} + \Delta$, where Δ is a suitable diagonal positive semidefinite matrix containing information not included in the previous step. The latter factorization provides the matrix S_{inex} , henceforth denoted by S_{cu} because it is the result of the combination of two updates. The corresponding inexact preconditioner is

$$\mathcal{P}_{cu} = \begin{bmatrix} I_n & 0 \\ AG^{-1} & I_m \end{bmatrix} \begin{bmatrix} G & 0 \\ 0 & -S_{cu} \end{bmatrix} \begin{bmatrix} I_n & G^{-1}A^T \\ 0 & I_m \end{bmatrix}. \quad (18)$$

Next we provide a detailed description of our procedure and analyze the quality of the resulting preconditioner. First, we describe how to perform the low-rank update of \mathcal{P}_{seed} and summarize related theoretical results given in [5]. Second, we present the subsequent updating step and provide new bounds on the eigenvalues of the preconditioned matrix.

3.1 First step: building S_{lr}

Let $\mathcal{L} = \{i : \Theta_{ii} \neq 0\}$ and let m_1 be its cardinality. We recall that the set \mathcal{L} does not change throughout the sequence. Furthermore, let $\tilde{\Theta}_{seed}$ and $\tilde{\Theta}$ be the $m_1 \times m_1$ diagonal submatrices containing the nonzero diagonal entries of Θ_{seed} and Θ , respectively, and let \tilde{I}_m be the rectangular matrix consisting of the columns of I_m with indices in \mathcal{L} . Then, the Schur complements S_{seed} and S , corresponding to \mathcal{P}_{seed} and \mathcal{P}_{ex} , respectively, can be written as follows:

$$S_{seed} = \tilde{A}\tilde{H}^{-1}\tilde{A}^T, \quad S = \tilde{A}\tilde{G}^{-1}\tilde{A}^T,$$

where

$$\tilde{A} = \begin{bmatrix} A & \tilde{I}_m \end{bmatrix}, \quad \tilde{H}^{-1} = \begin{bmatrix} H^{-1} & 0 \\ 0 & \tilde{\Theta}_{seed} \end{bmatrix}, \quad \tilde{G}^{-1} = \begin{bmatrix} G^{-1} & 0 \\ 0 & \tilde{\Theta} \end{bmatrix}. \quad (19)$$

Trivially,

$$S = S_{seed} + \tilde{A} \left(\tilde{G}^{-1} - \tilde{H}^{-1} \right) \tilde{A}^T. \quad (20)$$

The updating procedure described in [5] consists in finding a low-rank correction of S_{seed} of the form

$$S_{lr} = S_{seed} + \tilde{A} \tilde{K} \tilde{A}^T = \tilde{A} \left(\tilde{H}^{-1} + \tilde{K} \right) \tilde{A}^T = \tilde{A} \tilde{J}^{-1} \tilde{A}^T, \quad (21)$$

where \tilde{K} , and hence \tilde{J} , is a suitable diagonal matrix. The matrix \tilde{J} (or, equivalently, \tilde{K}) is chosen with a double goal: tightening the bounds on the eigenvalues provided by Theorem 1 and limiting the cost for computing the Cholesky factorization of S_{lr} as a modification of the Cholesky factorization (7). A key role in achieving the first goal is played by a result in [1], reported next for completeness.

Lemma 2 *Let $U \in \mathbb{R}^{r \times s}$ be full rank and let $E, F \in \mathbb{R}^{s \times s}$ be symmetric positive definite. Then*

$$\lambda_{\min}(E^{-1}F) \leq \lambda((UEU^T)^{-1}UFU^T) \leq \lambda_{\max}(E^{-1}F).$$

For any diagonal positive definite matrix $C \in \mathbb{R}^{(n+m_1) \times (n+m_1)}$, let $\gamma(C)$ be the vector of dimension $n + m_1$ with entries $\gamma_i(C)$ equal to the diagonal entries of $C\tilde{G}^{-1}$ sorted in nondecreasing order, i.e.,

$$\min_{1 \leq i \leq n+m_1} \frac{C_{ii}}{\tilde{G}_{ii}} = \gamma_1(C) \leq \gamma_2(C) \leq \dots \leq \gamma_n(C) = \max_{1 \leq i \leq n+m_1} \frac{C_{ii}}{\tilde{G}_{ii}}. \quad (22)$$

Lemma 2 yields the bounds

$$\gamma_1(\tilde{J}) \leq \lambda(S_{lr}^{-1}S) \leq \gamma_{n+m_1}(\tilde{J}). \quad (23)$$

We now give some definitions useful to specify \tilde{J} . Let $l = (l_1, l_2, \dots, l_{n+m_1})^T$ be the vector of indices such that

$$\gamma_{l_i}(\tilde{H}) = \frac{\tilde{H}_{ii}}{\tilde{G}_{ii}}.$$

Given two real constants $\mu_\gamma \geq 1$ and $\nu_\gamma \in (0, 1]$, and two nonnegative integers q_1 and q_2 such that $q = q_1 + q_2 \leq n + m_1$, we define $\tilde{\Gamma}$ as

$$\tilde{\Gamma} = \tilde{\Gamma}_1 \cup \tilde{\Gamma}_2, \quad (24)$$

where

$$\begin{aligned} \tilde{\Gamma}_1 &= \{i : n + m_1 - q_1 + 1 \leq l_i \leq n + m_1 \text{ and } \gamma_{l_i}(\tilde{H}) > \mu_\gamma\}, \\ \tilde{\Gamma}_2 &= \{i : 1 \leq l_i \leq q_2 \text{ and } \gamma_{l_i}(\tilde{H}) < \nu_\gamma\}. \end{aligned} \quad (25)$$

Then we define the matrix \tilde{K} in (21) by setting

$$\tilde{K}_{ii} = \begin{cases} \tilde{G}_{ii}^{-1} - \tilde{H}_{ii}^{-1} & \text{if } i \in \tilde{\Gamma}, \\ 0 & \text{otherwise.} \end{cases} \quad (26)$$

As a consequence, the diagonal entries of \tilde{J} take the following form:

$$\tilde{J}_{ii} = \begin{cases} \tilde{G}_{ii} & \text{if } i \in \tilde{\Gamma}, \\ \tilde{H}_{ii} & \text{otherwise.} \end{cases} \quad (27)$$

We also note that \tilde{J}^{-1} can be written as follows:

$$\tilde{J}^{-1} = \begin{bmatrix} J^{-1} & 0 \\ 0 & \tilde{\Theta}_{lr} \end{bmatrix}, \quad (28)$$

where $J \in \mathbb{R}^{n \times n}$ accounts for the changes from H to G , and $\tilde{\Theta}_{lr} \in \mathbb{R}^{m_1 \times m_1}$ for those from $\tilde{\Theta}_{seed}$ to $\tilde{\Theta}$.

Suppose that the sets $\tilde{\Gamma}_1$ and $\tilde{\Gamma}_2$ in (25) have cardinality equal to $q_1^* \leq q_1$ and $q_2^* \leq q_2$, respectively. Then, from (21), $S_{lr} - S_{seed}$ is a correction of rank $q^* = q_1^* + q_2^*$ equal to the cardinality of $\tilde{\Gamma}$. Once S_{lr} is defined, the low-rank update preconditioner \mathcal{P}_{lr} is simply obtained by setting $S_{inex} = S_{lr}$ in (10), i.e.,

$$\mathcal{P}_{lr} = \begin{bmatrix} I_n & 0 \\ AG^{-1} & I_m \end{bmatrix} \begin{bmatrix} G & 0 \\ 0 & -S_{lr} \end{bmatrix} \begin{bmatrix} I_n & G^{-1}A^T \\ 0 & I_m \end{bmatrix}. \quad (29)$$

By combining Theorem 1 with (23) and (27), we get the following bounds on the eigenvalues of $\mathcal{P}_{lr}^{-1}\mathcal{A}$, (see [5, Corollary 3.3]).

Corollary 3 *Let \mathcal{A} , \mathcal{P}_{lr} and X be the matrices in (8), (29) and (11), and let λ be an eigenvalue of $\mathcal{P}_{lr}^{-1}\mathcal{A}$. Let $\gamma(\tilde{H})$ and $\gamma(\tilde{J})$ be defined as in (22), and let q_1^* and q_2^* be the cardinality of the sets $\tilde{\Gamma}_1$ and $\tilde{\Gamma}_2$ in (25), respectively. If $2I_n - X$ is positive definite, then λ satisfies either (14) or (15)–(16) with*

$$\bar{\lambda} \leq \gamma_{n+m_1}(\tilde{J}) \max\{2 - \lambda_{\min}(X), 1\}, \quad (30)$$

$$\lambda \geq \gamma_1(\tilde{J}) \min\{2 - \lambda_{\max}(X), 1\}, \quad (31)$$

where

$$\begin{aligned} \gamma_1(\tilde{J}) &= \min\{1, \min_{i \notin \tilde{\Gamma}} \gamma_{l_i}(\tilde{H})\} = \min\{1, \gamma_{q_2^*+1}(\tilde{H})\}, \\ \gamma_{n+m_1}(\tilde{J}) &= \max\{1, \max_{i \notin \tilde{\Gamma}} \gamma_{l_i}(\tilde{H})\} = \max\{1, \gamma_{n+m_1-q_1^*}(\tilde{H})\}. \end{aligned} \quad (32)$$

Furthermore,

$$|\mathcal{I}(\lambda)| \leq \sqrt{\gamma_{n+m_1}(\tilde{J})} \|I_n - X\|. \quad (33)$$

It is clear that the bounds provided by the previous corollary are expected to improve as the set \tilde{I} is enlarged, i.e., q is increased. Likewise, smaller (larger) values of μ_γ (ν_γ) may generally improve these bounds. However, the choice of the previous values must take into account the cost for computing the Cholesky factorization of S_{lr} by updating the factorization available for S_{seed} . From (7), (21) and (26) it follows that

$$S_{lr} = L_{seed} D_{seed} L_{seed}^T + \bar{A} \bar{K} \bar{A}^T,$$

where $\bar{A} \in \mathbb{R}^{m \times q^*}$ consists of the columns of \tilde{A} with indices in \tilde{I} , and $\bar{K} \in \mathbb{R}^{q^* \times q^*}$ is the diagonal matrix having on the diagonal the nonzero entries of \tilde{K} corresponding to those indices. Therefore, the Cholesky factorization

$$S_{lr} = L_{lr} D_{lr} L_{lr}^T, \quad (34)$$

is a rank- q^* modification of the factorization (7) and can be computed, e.g., through the update and downdate procedures in [17]. Note that an update is required if $\tilde{H}_{ii} > \tilde{G}_{ii}$ and a downdate is required if $\tilde{H}_{ii} < \tilde{G}_{ii}$.

The computational cost for building S_{lr} depends on the value of q , and for practical purposes the updating strategy is more convenient than computing the exact preconditioner as long as q is kept fairly small (see [5]). Similarly, values of μ_γ and ν_γ not too close to 1 are used in practice; thus indices that do not bring significant improvement in the eigenvalue bounds while increasing the cost for the updating strategy (see again [5]) are excluded from \tilde{I} . Since a large number of entries in $\tilde{G} - \tilde{H}$ must be discarded, we now propose a second step of our updating approach for recovering part of the discarded information.

3.2 Second step: updating S_{lr} to get S_{cu}

In order to recover information not included in S_{lr} , we observe that part of this information can be regarded as a positive semidefinite diagonal modification, Δ , of S_{lr} . Therefore, we can compute a low-cost approximate LDL^T factorization of the matrix $S_{lr} + \Delta$ by exploiting the procedures in [3, 4, 30] to update the factorization (34) of S_{lr} . Furthermore, by expressing $S_{lr} + \Delta$ as $\tilde{A} \tilde{J}_\Delta^{-1} \tilde{A}^T$, with \tilde{J}_Δ diagonal positive definite, we can exploit Lemma 2 to derive new eigenvalue bounds.

We consider the diagonal matrix $\tilde{\Delta} \in \mathbb{R}^{m_1 \times m_1}$ defined by

$$\tilde{\Delta}_{ii} = \begin{cases} \tilde{\Theta}_{ii} - (\tilde{\Theta}_{seed})_{ii} & \text{if } i + n \notin \tilde{I} \text{ and } \tilde{\Theta}_{ii} - (\tilde{\Theta}_{seed})_{ii} > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (35)$$

whose nonzero entries correspond to the ratios $\tilde{\Theta}_{ii}/(\tilde{\Theta}_{seed})_{ii}$ that are greater than 1 and have been discarded in the construction of S_{lr} . Then, recalling the definition of \tilde{I}_m in Section 3.1, we set $\Delta \in \mathbb{R}^{m \times m}$ as follows:

$$\Delta = \tilde{I}_m \tilde{\Delta} \tilde{I}_m^T, \quad (36)$$

and define the matrix

$$S_\Delta = S_{lr} + \Delta. \quad (37)$$

Using S_Δ instead of S_{lr} generally improves the quality of the preconditioner \mathcal{P}_{cu} , as shown next. Let

$$\tilde{\Gamma}_\Delta = \tilde{\Gamma} \cup \left\{ j : j = i + n, \text{ with } 1 \leq i \leq m_1 \text{ and } \gamma_{l_j}(\tilde{H}) = \frac{\tilde{\Theta}_{ii}}{(\tilde{\Theta}_{seed})_{ii}} > 1 \right\},$$

and let $\tilde{\Theta}_\Delta \in \mathbb{R}^{m_1 \times m_1}$ the diagonal matrix defined by

$$(\tilde{\Theta}_\Delta)_{ii} = \begin{cases} \tilde{\Theta}_{ii} & \text{if } i + n \in \tilde{\Gamma}_\Delta, \\ (\tilde{\Theta}_{seed})_{ii} & \text{otherwise.} \end{cases}$$

Then we have

$$S_\Delta = \tilde{A} \tilde{J}_\Delta^{-1} \tilde{A}^T, \quad \tilde{J}_\Delta^{-1} = \begin{bmatrix} J^{-1} & \\ & \tilde{\Theta}_\Delta \end{bmatrix} \tilde{A}^T,$$

where \tilde{A} and J are the matrices in (19) and (28), respectively. By reasoning as in Section 3.1, we get the following bounds on the eigenvalues of $S_\Delta^{-1}S$:

$$\hat{\gamma}_1 \leq \lambda(S_\Delta^{-1}S) \leq \hat{\gamma}_{n+m_1}, \quad (38)$$

where

$$\hat{\gamma}_1 = \gamma_1(\tilde{J}) = \min\{1, \gamma_{q_2^*+1}(\tilde{H})\}, \quad (39)$$

$$\hat{\gamma}_{n+m_1} = \max \left\{ 1, \max_{i \notin \tilde{\Gamma}_\Delta} \gamma_{l_i}(\tilde{H}) \right\}. \quad (40)$$

Trivially, $\hat{\gamma}_{n+m_1} \leq \gamma_{n+m_1}(\tilde{J})$, with $\gamma_{n+m_1}(\tilde{J})$ given in (32), and the benefit of using S_Δ depends on the size of $\hat{\gamma}_{n+m_1}$ with respect to $\gamma_{n+m_1}(\tilde{J})$.

On the other hand, using S_Δ in place of S_{lr} requires its factorization. In order to limit the computational cost of the preconditioner, we can compute an approximation to the Cholesky factorization of S_Δ ,

$$S_{cu} = L_{cu} D_{cu} L_{cu}^T \simeq S_\Delta, \quad (41)$$

by using the updating procedures mentioned at the beginning of this section. The procedure considered here updates the matrices L_{lr} and D_{lr} in (34) as follows (see [4] for details). Let W be the diagonal matrix with diagonal entries defined by

$$W_{ii} = \frac{(D_{lr})_{ii}}{(D_{lr})_{ii} + \Delta_{ii}}, \quad i = 1, \dots, m;$$

we define the matrices L_{cu} and D_{cu} as

$$\begin{aligned} D_{cu} &= D_{lr} + \Delta, \\ (L_{cu})_{jj} &= 1, \quad j = 1, \dots, m, \\ L_{cu}(j+1:m, j) &= W_{jj} L_{lr}(j+1:m, j), \quad j = 1, \dots, m-1, \end{aligned} \quad (42)$$

where, using the Matlab notation, $L_{lr}(j+1:n, j)$ denotes the strictly lower triangular part of the j -th column of L_{lr} . Note that the sparsity pattern of the factors of S_{lr} is preserved; furthermore, the cost of forming S_{cu} is low, since the cost of D_{cu} is negligible, and the computation of L_{cu} consists in scaling only the nonzero entries of the columns of L_{lr} corresponding to the nonzero entries of Δ . Obviously, if Δ is the zero matrix, the update is not performed and $S_{cu} = S_{lr}$, i.e., $\mathcal{P}_{cu} = \mathcal{P}_{lr}$. This limit case occurs if $\bar{\theta}_{ii} \leq (\bar{\theta}_{seed})_{ii}$ for $i = 1, \dots, m_1$, or if all the indices $i+n$ associated with the ratios $\bar{\theta}_{ii}/(\bar{\theta}_{seed})_{ii}$ that are greater than 1 belong to $\tilde{\Gamma}$.

We refer to [3, 4] for theoretical results and computational experiences motivating the previous updating approach. Here we report only some results concerning the eigenvalues of $S_{cu}^{-1}S$ (see [4, Theorems 2.2 and 2.4]), which are useful in the spectral analysis of $\mathcal{P}_{cu}^{-1}\mathcal{A}$.

Theorem 4 *Let S_Δ and S_{cu} be the matrices in (37) and (41). For all $\varepsilon > 0$ there exists $\eta > 0$ such that if $\|\Delta\| < \eta$, then*

$$|\lambda(S_{cu}^{-1}S_\Delta) - 1| < \varepsilon$$

for all the eigenvalues of $S_{cu}^{-1}S_\Delta$. Furthermore, if $S_\Delta - S_{cu}$ has rank $m-l$, then l eigenvalues of $S_{cu}^{-1}S_\Delta$ are equal to 1.

Theorem 5 *Let S_Δ and S_{cu} be the matrices in (37) and (41). For all $\varepsilon > 0$ there exists $\eta > 0$ such that if $\Delta_{ii} > \eta$ for r diagonal entries of Δ , then*

$$|\lambda(S_{cu}^{-1}S_\Delta) - 1| < \varepsilon$$

for r eigenvalues of $S_{cu}^{-1}S_\Delta$.

The next theorem provides bounds on the eigenvalues of $\mathcal{P}_{cu}^{-1}\mathcal{A}$ in terms of $\hat{\gamma}_1$ and $\hat{\gamma}_{n+m_1}$ and of the minimum and maximum eigenvalues of $S_{cu}^{-1}S_\Delta$.

Theorem 6 *Let \mathcal{A} , \mathcal{P}_{cu} , S_Δ , S_{cu} and X be the matrices in (8), (18), (37), (41) and (11), and let λ be an eigenvalue of $\mathcal{P}_{cu}^{-1}\mathcal{A}$. Let $\hat{\gamma}_1$ and $\hat{\gamma}_{n+m_1}$ be the scalars in (39) and (40). If $2I_n - X$ is positive definite, then λ satisfies either (14) or (15)–(16) with*

$$\bar{\lambda} \leq \hat{\gamma}_{n+m_1} \lambda_{\max}(S_{cu}^{-1}S_\Delta) \max\{2 - \lambda_{\min}(X), 1\}, \quad (43)$$

$$\underline{\lambda} \geq \hat{\gamma}_1 \lambda_{\min}(S_{cu}^{-1}S_\Delta) \min\{2 - \lambda_{\max}(X), 1\}. \quad (44)$$

Furthermore,

$$|\mathcal{I}(\lambda)| \leq \sqrt{\hat{\gamma}_{n+m_1} \lambda_{\max}(S_{cu}^{-1}S_\Delta)} \|I_n - X\|. \quad (45)$$

Proof The eigenvalue problem $S_{cu}^{-1}Sw = \lambda w$ is equivalent to

$$S_\Delta^{-\frac{1}{2}} S S_\Delta^{-\frac{1}{2}} \bar{w} = \lambda S_\Delta^{-\frac{1}{2}} S_{cu} S_\Delta^{-\frac{1}{2}} \bar{w},$$

with $\bar{w} = S_{\Delta}^{-\frac{1}{2}} w$. Hence, for any eigenvalue of $S_{cu}^{-1} S$ and any corresponding eigenvector w we have

$$\lambda(S_{cu}^{-1} S) = \frac{\bar{w}^T S_{\Delta}^{-\frac{1}{2}} S S_{\Delta}^{-\frac{1}{2}} \bar{w}}{\bar{w}^T S_{\Delta}^{-\frac{1}{2}} S_{cu} S_{\Delta}^{-\frac{1}{2}} \bar{w}}.$$

Then, by exploiting (38) we get

$$\lambda(S_{cu}^{-1} S) \leq \frac{\lambda_{\max}(S_{\Delta}^{-1} S)}{\lambda_{\min}(S_{\Delta}^{-1} S_{cu})} \leq \hat{\gamma}_{n+m_1} \lambda_{\max}(S_{cu}^{-1} S_{\Delta})$$

and

$$\lambda(S_{cu}^{-1} S) \geq \frac{\lambda_{\min}(S_{\Delta}^{-1} S)}{\lambda_{\max}(S_{\Delta}^{-1} S_{cu})} \geq \hat{\gamma}_1 \lambda_{\min}(S_{cu}^{-1} S_{\Delta}).$$

Inequalities (43) and (44) follow by recalling (12) and (13).

It remains to prove (45). The eigenvalue problem $S_{cu}^{-1} A G^{-1} A^T w = \lambda w$ is equivalent to

$$S_{\Delta}^{-\frac{1}{2}} A G^{-1} A^T S_{\Delta}^{-\frac{1}{2}} \bar{w} = \lambda S_{\Delta}^{-\frac{1}{2}} S_{cu} S_{\Delta}^{-\frac{1}{2}} \bar{w}.$$

Furthermore, since Θ is positive definite, for any vector $v \in \mathbb{R}^m$

$$v^T S_{\Delta}^{-\frac{1}{2}} A G^{-1} A^T S_{\Delta}^{-\frac{1}{2}} v \leq v^T S_{\Delta}^{-\frac{1}{2}} (A G^{-1} A^T + \Theta) S_{\Delta}^{-\frac{1}{2}} v,$$

and by matrix similarity

$$\lambda_{\max}(S_{\Delta}^{-1} A G^{-1} A^T) \leq \lambda_{\max}(S_{\Delta}^{-1} S).$$

By reasoning as in the first part of the proof and exploiting the previous inequality we get

$$\lambda(S_{cu}^{-1} A G^{-1} A^T) \leq \frac{\lambda_{\max}(S_{\Delta}^{-1} S)}{\lambda_{\min}(S_{\Delta}^{-1} S_{cu})} = \hat{\gamma}_{n+m_1} \lambda_{\max}(S_{cu}^{-1} S_{\Delta}),$$

and (45) follows from (17). \square

The following result is a straightforward consequence of Theorem 6.

Corollary 7 *Let \mathcal{A} , \mathcal{P}_{cu} , S_{Δ} , S_{cu} and X be the matrices in (8), (18), (37), (41) and (11), and let λ be an eigenvalue of $\mathcal{P}_{cu}^{-1} \mathcal{A}$. Let $\hat{\gamma}_1$ and $\hat{\gamma}_{n+m_1}$ be the scalars in (39) and (40). If $2I_n - X$ is positive definite and, for all the eigenvalues of $S_{cu}^{-1} S_{\Delta}$,*

$$|\lambda(S_{cu}^{-1} S_{\Delta}) - 1| \leq \varepsilon$$

for some $\varepsilon > 0$, then λ satisfies either (14) or (15)–(16) with

$$\begin{aligned} \bar{\lambda} &\leq (1 + \varepsilon) \hat{\gamma}_{n+m_1} \max\{2 - \lambda_{\min}(X), 1\}, \\ \underline{\lambda} &\geq (1 - \varepsilon) \hat{\gamma}_1 \min\{2 - \lambda_{\max}(X), 1\}. \end{aligned}$$

Furthermore,

$$|\mathcal{I}(\lambda)| \leq \sqrt{(1 + \varepsilon) \hat{\gamma}_{n+m_1}} \|I_n - X\|.$$

The previous results show that the bounds (30) and (33) can be effectively improved if $\lambda_{\max}(S_{cu}^{-1}S_{\Delta})$ is not too far from 1. Furthermore, $\lambda_{\min}(S_{cu}^{-1}S_{\Delta})$ must not be too far from 1, in order to avoid a significant deterioration of the bound (31). By Theorems 4 and 5, this happens when the entries of Δ are either all sufficiently small or all sufficiently large. In general, we cannot expect that all the eigenvalues of $S_{cu}^{-1}S_{\Delta}$ are close to 1; nevertheless, the ability of S_{cu} of clustering around 1 some eigenvalues of S_{Δ} when Δ has large entries, provides a way to tighten the bounds on the spectrum of $\mathcal{P}_{cu}^{-1}\mathcal{A}$, as confirmed by numerical experiments.

We conclude this section by summarizing in Algorithm 1 the main steps of the overall updating procedure.

Algorithm 1: updating the constraint preconditioner

- Given $\mu_{\gamma} \geq 1$, $\nu_{\gamma} \in (0, 1]$ and the nonnegative integers q_1 and q_2 ,
1. form \tilde{H} defined in (19) and compute $\gamma_i(\tilde{H})$, $i = 1, \dots, n + m_1$, according to (22);
 2. set \tilde{T} as in (24), q^* as the cardinality of \tilde{T} , and \tilde{J} as in (27);
 3. compute the factorization in (34), $S_{lr} = L_{lr}D_{lr}L_{lr}^T$, by a rank- q^* correction of S_{seed} in (7);
 4. compute the matrix Δ given in (36);
 5. if $\Delta \neq 0$
 - compute the approximate factorization $L_{cu}D_{cu}L_{cu}^T$ of $L_{lr}D_{lr}L_{lr}^T + \Delta$ (see (42));
 - else
 - set $L_{cu} = L_{lr}$, $D_{cu} = D_{lr}$;
 6. set

$$\mathcal{P}_{cu} = \begin{bmatrix} I_n & 0 \\ AG^{-1} & L_{cu} \end{bmatrix} \begin{bmatrix} G & 0 \\ 0 & -D_{cu} \end{bmatrix} \begin{bmatrix} I_n & G^{-1}A^T \\ 0 & L_{cu}^T \end{bmatrix}.$$

4 Numerical results

We provide some illustrative examples of the behavior of the preconditioner \mathcal{P}_{cu} built with Algorithm 1, compared with the exact CP preconditioner \mathcal{P}_{ex} in (9) and the updated preconditioner \mathcal{P}_{lr} in (29).

We consider five sequences of KKT systems that arise in the solution by an IP method of convex quadratic programming problems with linear inequality constraints. These problems have been obtained by modifying the CUTEst [27] problems CVXQP1, CVXQP3 and MOSARQP1. The modifications have been made because large or variable-size CUTEst convex quadratic programming problems with inequality constraints have Schur complements that are very inexpensive to factorize and thus are useless for our experiments. More precisely, CVXQP1 and CVXQP3, which have non-banded Schur complements, have been modified so that their original equality constraint $Ax = b$ became $Ax \geq b$ (the modified problems have been identified by appending “-M” to the

original names). Furthermore, in order to increase the density of their Schur complements, four nonzero entries per row have been added to the matrix A of CVXQP1-M, and one nonzero entry per row to the matrix A of CVXQP3-M (the denser problems have been identified by appending “-D” and “-D2” to their names, respectively, according to the notation used in [5]). Finally, nonzeros in the positions (i, n) , with $\text{mod}(i, 10) = 1$, have been added to the matrix A of MOSARQP1, obtaining problem MOSARQP1-D (note that the original problem has a narrow-banded Schur complement). The dimensions n and m of the five problems and the number $\text{nnz}(S)$ of nonzero entries of their Schur complements are reported in the second column of Table 1.

The sequences of KKT systems have been obtained by running the Fortran 90 PRQP code, which implements an infeasible inexact potential reduction IP method [11, 13, 15], and extracting the KKT matrices arising at each IP iteration and the corresponding right-hand sides. Successively, these sequences have been solved offline, applying \mathcal{P}_{ex} , \mathcal{P}_{lr} and \mathcal{P}_{cu} . The starting point for the IP procedure in PRQP has been built with the STP2 strategy described in [16] and the tolerances on the relative duality gap and the relative infeasibilities have been set to 10^{-6} and 10^{-7} , respectively. Within PRQP the KKT systems have been solved by the SQMR method coupled with the exact CP, using an adaptive tolerance in the stopping criterion, which relates the accuracy in the solution of the KKT system to the quality of the current IP iterate, in the spirit of inexact IP methods [2, 12]. Specifically, in PRQP the SQMR iterations were stopped as soon as the norm of the unpreconditioned residual was below a tolerance of the form $\tau = \min\{\max\{\tau_1, 10^{-8}\}, 10^{-2}\|r_0\|\}$, where τ_1 depends on the duality gap value at the current IP iteration and r_0 is the initial unpreconditioned residual (see [12] for more details). The values of τ corresponding to all the systems of each sequence were saved to be used in our experiments.

The numerical experiments have been performed in the Matlab environment, exploiting C code for the construction of \mathcal{P}_{lr} and \mathcal{P}_{cu} . More precisely, the CHOLMOD package [14] has been called, through its MEX interface, to compute the sparse LDL^T factorization of S_{seed} and the low-rank updates and downdates required to build \mathcal{P}_{lr} . The updates (42) needed to form \mathcal{P}_{cu} have been implemented in C as Matlab MEX-files too. The preconditioners have been built using $q_1 = q_2 = q/2 = 25$, in order to keep low the overhead of the updating/downdating phase; furthermore, $\mu_\gamma = 10$ and $\nu_\gamma = 0.1$ have been used to select the indices to be put in \tilde{I} . These choices agree with the results of the experiments with \mathcal{P}_{lr} discussed in [5]. When less than q_1 ratios $\gamma_{i_i}(\tilde{H}) \geq \mu_\gamma$ (or less than q_2 ratios $\gamma_{i_i}(\tilde{H}) \leq \nu_\gamma$) were available, the number of ratios $\gamma_{i_i}(\tilde{H}) \leq \nu_\gamma$ (or $\gamma_{i_i}(\tilde{H}) \geq \mu_\gamma$) was increased to have a total number of ratios in \tilde{I} as close as possible to q . We note that we have not applied any scaling to the matrix X in (11) to ensure positive definiteness of $2I_n - X$, although this is assumed in our theory. Nevertheless, the results generally appear to be in agreement with the theory (see also [5]). The linear systems have been solved using a Matlab implementation of the SQMR method without

look-ahead [22], stopping the iterations when the norm of the residual became smaller than the associated tolerance, as in the solution of the KKT systems within the IP code. A maximum number of 500 iterations has been considered too.

Following [5], the preconditioners have been “refreshed” as explained next. When, for a system of the sequence, the time for computing \mathcal{P}_{lr} or \mathcal{P}_{cu} and solving the preconditioned linear system exceeded 90% of the time for building the last exact preconditioner and solving the corresponding system, the next system of the sequence was solved using the exact CP in (9). Furthermore, a maximum number of consecutive updates, k_{\max} , was also considered; here $k_{\max} = 4$. This strategy aims at avoiding deterioration of the quality of the updated preconditioner and hence excessive increase of the number of SQMR iterations, which may make the updating strategy useless. We note that, by using a refreshing criterion based on the execution time, both the problem and the computing environment are taken into account to get an efficient preconditioning strategy.

The tests have been performed on a six-core Xeon processor with clock frequency of 2.4 GHz, 24 GB of RAM and 12 MB of cache memory, running Ubuntu/Linux 12.04.5 (kernel version 3.2.0.83_generic). CHOLMOD and the C code for the updates (42) have been compiled with gcc 4.6.3, and Matlab R2015a (v. 8.5, 64-bit) has been employed. The `tic` and `toc` Matlab commands have been used to measure the execution times.

We start the description of our numerical results comparing \mathcal{P}_{lr} and \mathcal{P}_{cu} . For the problems from the CVXQP family, we observe that the number of elements used for the second step of the update, namely $nnz(\Delta)$ with Δ defined in (35)–(36), is zero or negligible (at most twenty) in the first systems of the sequence (first nine systems for CVXQP1-M-D, and first ten systems for CVXQP1-M, CVXQP3-M and CVXQP3-M-D2). Since no significant benefit can be obtained from the update strategy (42) with such a small number of elements, we exclude these first systems from the following analysis.

The results obtained with \mathcal{P}_{lr} and \mathcal{P}_{cu} are shown in Table 1. We report the range $IPits$ of IP iterations considered for each sequence and, for each preconditioner, the total number nit of SQMR iterations performed, the number $nref$ of times the preconditioner has been refreshed, the total time T_{prec} for building the preconditioner, the total time T_{solve} for solving the preconditioned system, and the sum T_{tot} of the two times. The times are expressed in seconds. We see that all the runs with \mathcal{P}_{cu} are faster than those with \mathcal{P}_{lr} . The gain, in terms of total execution time, varies between 6% and 23%. Furthermore, for two sequences, the savings obtained in the number of SQMR iterations reduce the number of preconditioner refreshes with respect to the use of \mathcal{P}_{lr} .

Details on the solution of the sequences of KKT systems by using \mathcal{P}_{lr} and \mathcal{P}_{cu} are shown in Tables 2, 3 and 4 for CVXQP1-M, CVXQP3-M-D2 and MOSARQP1-D, respectively. The IP iteration number corresponding to each system of a sequence is indicated by $IP\#$. For each system and for both preconditioners, we report the number nit of SQMR iterations, the time T_{prec} for building the preconditioner for that system, the time T_{solve} for solving the

Table 1 Comparison of \mathcal{P}_h and \mathcal{P}_{cu} : total number of SQMR iterations, refreshes of the preconditioners and execution times.

<i>Problem</i>	n m $mz(S)$	IP_{its}	\mathcal{P}_h ($q = 50$)					\mathcal{P}_{cu}				
			nit	$nref$	T_{prec}	T_{solve}	T_{tot}	nit	$nref$	T_{prec}	T_{solve}	T_{tot}
CVXQP1-M	20000 10000 67976	11-26	920	5	1.84e+1	4.49e+1	6.33e+1	597	5	1.92e+1	2.98e+1	4.90e+1
CVXQP1-M-D	20000 10000 240494	10-28	968	9	4.12e+1	1.16e+2	1.57e+2	750	6	4.03e+1	9.13e+1	1.32e+2
CVXQP3-M	20000 15000 155942	11-30	1066	5	1.03e+2	2.21e+2	3.24e+2	943	5	1.08e+2	1.96e+2	3.04e+2
CVXQP3-M-D2	20000 15000 224396	11-25	632	6	7.36e+1	1.70e+2	2.44e+2	583	5	7.02e+1	1.55e+2	2.25e+2
MOSARQP1-D	22500 20000 4259250	1-24	318	7	3.12e+1	5.37e+1	8.49e+1	227	7	3.80e+01	3.91e+01	7.71e+1

preconditioned system, and $T_{sum} = T_{prec} + T_{solve}$. For \mathcal{P}_{lr} we also display the size q^* of the low-rank update performed, i.e., the cardinality of the set \tilde{I} in (24), while for \mathcal{P}_{cu} we display both q^* and the number of elements used for the second step of the update, i.e., $nnz(\Delta)$. The data corresponding to refreshes are in bold and the times are expressed in seconds.

We observe that, in general, forming \mathcal{P}_{cu} is inexpensive although $nnz(\Delta)$ is large. Using \mathcal{P}_{cu} may significantly reduce the number of SQMR iterations, especially in the last runs before a refresh. In fact, the performance of both \mathcal{P}_{lr} and \mathcal{P}_{cu} tends to deteriorate progressively after a refresh, but the use of the diagonal modification in \mathcal{P}_{cu} may considerably improve the effectiveness of \mathcal{P}_{cu} with respect to \mathcal{P}_{lr} . This behavior is clearly illustrated, e.g., by the data of the 15th, 19th and 22nd system of CVXQP1-M (Table 2); a similar behavior can be recognized in Tables 3 and 4. Taking into account this numerical evidence and the fact that forming \mathcal{P}_{cu} is simple and inexpensive, the second phase of the update appears worthy to be coupled with the low-rank correction in the update of CP preconditioners for KKT systems.

To provide further insight into the behavior of \mathcal{P}_{cu} , in Table 5 we report the maximum and minimum eigenvalues of $S_{lr}^{-1}S$ and $S_{cu}^{-1}S$ from the 12th to the 25th iteration of CVXQP1-M. Note that at the 25th iteration S_{lr} is the matrix which S_{cu} is built from, and not the matrix associated with the preconditioner \mathcal{P}_{lr} considered in Table 2 (with \mathcal{P}_{lr} and \mathcal{P}_{cu} the refresh occurs at different IP iterations). We see that the updating strategy used to build S_{cu} practically does not change the smallest eigenvalue, while it reduces the largest one, according to the decrease observed in the number of iterations (see Table 2).

We conclude this section by comparing the preconditioners \mathcal{P}_{cu} and \mathcal{P}_{ex} . In Table 6 we summarize the total number of SQMR iterations and execution times for \mathcal{P}_{ex} ; the same data for \mathcal{P}_{cu} , available in Table 1, are repeated for the sake of readability. The reduction of the total time obtained with \mathcal{P}_{cu} over \mathcal{P}_{ex} varies between 6% and 25% although, as expected, the number of SQMR iterations performed is larger than with \mathcal{P}_{ex} . It is noteworthy that \mathcal{P}_{cu} speeds up the solution of the sequences associated with CVXQP1-M, CVXQP1-M-D and MOSARQP1, where the use of \mathcal{P}_{lr} is not beneficial with respect to \mathcal{P}_{ex} . We also report briefly on experiments with inexact CPs obtained by approximating the exact Schur complement S with an inexact Cholesky factorization of it. These experiments have been performed to further assess the reliability of our updating approach. Unmodified and modified incomplete Cholesky factorizations of S have been computed using the Matlab function `ichol` and drop tolerances ranging from 10^{-1} to 10^{-5} . However, the resulting inexact CPs seem to lack robustness on our sequences, as SQMR failed in the solution of at least one system of each sequence.

Table 2 CVXQP1-M: SQMR iterations and execution times obtained with P_{lr} and P_{cu} (the data concerning the application of the exact CP are in bold). The number q^* of low-rank updates and the number $mz(\Delta)$ of elements used for building P_{cu} are also reported.

$IP\#$	$P_{lr} (q = 50)$					P_{cu}				
	nit	q^*	T_{prec}	T_{solve}	T_{sum}	nit	$q^*/mz(\Delta)$	T_{prec}	T_{solve}	T_{sum}
11	9		2.70e+0	5.04e-1	1.92e+1	9		2.70e+0	5.04e-1	3.20e+0
12	13	1	1.02e-1	6.85e-1	3.57e+1	10	1/1696	1.76e-1	5.48e-1	7.24e-1
13	19	14	3.39e-1	9.63e-1	3.55e+1	14	14/1841	4.12e-1	7.44e-1	1.16e+0
14	37	50	8.54e-1	1.85e+0	3.51e+1	27	50/1949	9.28e-1	1.37e+0	2.29e+0
15	106	50	7.52e-1	5.15e+0	3.42e+1	46	50/2084	8.27e-1	2.29e+0	3.12e+0
16	16		2.62e+0	8.39e-1	3.33e+1	16		2.62e+0	8.39e-1	3.46e+0
17	22	5	1.16e-1	1.10e+0	3.07e+1	20	5/2495	1.91e-1	1.02e+0	1.21e+0
18	41	50	6.47e-1	2.02e+0	3.05e+1	36	50/2486	7.22e-1	1.81e+0	2.53e+0
19	89	50	6.38e-1	4.38e+0	2.98e+1	57	50/2491	7.13e-1	2.83e+0	3.55e+0
20	20		2.68e+0	1.03e+0	2.91e+1	20		2.68e+0	1.03e+0	3.71e+0
21	37	4	2.01e-1	1.83e+0	2.64e+1	28	4/2537	2.78e-1	1.49e+0	1.77e+0
22	297	50	6.64e-1	1.42e+1	2.61e+1	114	50/2517	7.39e-1	5.51e+0	6.25e+0
23	37		2.63e+0	1.87e+0	2.54e+1	37		2.63e+0	1.87e+0	4.50e+0
24	86	14	2.51e-1	4.05e+0	2.28e+1	47	14/2537	3.28e-1	2.35e+0	2.68e+0
25	37		2.60e+0	1.86e+0	2.24e+1	79	50/2514	6.67e-1	3.76e+0	4.42e+0
26	54	47	5.61e-1	2.62e+0	2.18e+1	37		2.57e+0	1.85e+0	4.43e+0
	920		1.84e+1	4.49e+1	6.33e+1	597		1.92e+1	2.98e+1	4.90e+1

Table 3 CVXQP3-M-D2: SQMR iterations and execution times obtained with \mathcal{P}_r and \mathcal{P}_{cu} (the data concerning the application of the exact CP are in bold). The number q^* of low-rank updates and the number $nnz(\Delta)$ of elements used for building \mathcal{P}_{cu} are also reported.

IP#	$\mathcal{P}_{lr} \ (q = 50)$					\mathcal{P}_{cu}				
	nit	q^*	T_{prec}	T_{solve}	T_{sum}	nit	$q^*/nnz(\Delta)$	T_{prec}	T_{solve}	T_{sum}
11	9	9.58e+0	2.59e+0	1.22e+1	1.22e+1	9	9.58e+0	2.59e+0	1.22e+1	1.22e+1
12	20	6	6.59e-1	5.33e+0	5.99e+0	12	6/2227	1.10e+0	3.34e+0	4.44e+0
13	33	50	2.72e+0	9.07e+0	1.18e+1	19	50/2752	3.18e+0	5.36e+0	8.54e+0
14	15	9.56e+0	4.11e+0	1.37e+1	1.37e+1	41	50/3282	3.68e+0	1.10e+1	1.47e+1
15	22	1	5.70e-1	6.00e+0	6.57e+0	17	9.54e+0	4.47e+0	1.40e+1	1.40e+1
16	47	50	2.73e+0	1.25e+1	1.52e+1	28	2/4897	8.79e-1	7.51e+0	8.39e+0
17	24	9.44e+0	6.31e+0	1.58e+1	1.58e+1	51	5/4970	3.11e+0	1.41e+1	1.73e+1
18	40	7	4.58e-1	1.03e+1	1.07e+1	28	9.55e+0	7.26e+0	1.68e+1	1.68e+1
19	68	50	2.78e+0	1.85e+1	2.12e+1	43	3/5152	9.43e-1	1.17e+1	1.26e+1
20	32	9.62e+0	8.18e+0	1.78e+1	1.78e+1	58	50/5156	3.62e+0	1.64e+1	2.00e+1
21	46	2	5.38e-1	1.25e+1	1.30e+1	36	9.60e+0	9.72e+0	1.93e+1	1.93e+1
22	83	50	3.93e+0	2.32e+1	2.71e+1	56	4/5210	1.20e+0	1.37e+1	1.49e+1
23	41	9.49e+0	1.13e+1	2.08e+1	2.08e+1	61	30/5204	3.40e+0	1.47e+1	1.81e+1
24	82	18	1.97e+0	2.18e+1	2.37e+1	43	9.77e+0	1.15e+1	2.13e+1	2.13e+1
25	70	9.58e+0	1.86e+1	2.82e+1	2.82e+1	81	1/5207	1.02e+0	2.19e+1	2.29e+1
	632	7.36e+1	1.70e+2	2.44e+2	2.44e+2	583		7.02e+1	1.55e+2	2.25e+2

Table 4 MOSARQP1-D: SQMR iterations and execution times obtained with \mathcal{P}_{lr} and \mathcal{P}_{cu} (the data concerning the application of the exact CP are in bold). The number q^* of low-rank updates and the number $\text{nnz}(\Delta)$ of elements used for building \mathcal{P}_{cu} are also reported.

$IP\#$	$\mathcal{P}_{lr} (q = 50)$					\mathcal{P}_{cu}				
	nit	q^*	T_{prec}	T_{solve}	T_{sum}	nit	$q^*/\text{nnz}(\Delta)$	T_{prec}	T_{solve}	T_{sum}
1	1		2.71e+0	2.98e-1	3.01e+0	1		2.71e+0	2.98e-1	3.01e+0
2	3	1	4.28e-1	6.01e-1	1.03e+0	4	1/2068	6.45e-1	7.41e-1	1.39e+0
3	7	2	4.24e-1	1.28e+0	1.70e+0	5	2/2071	6.41e-1	9.78e-1	1.62e+0
4	13	50	4.80e-1	2.23e+0	2.71e+0	8	50/2076	6.80e-1	1.31e+0	1.99e+0
5	2		2.88e+0	4.62e-1	3.34e+0	10	50/2160	6.46e-1	1.71e+0	2.36e+0
6	10	2	2.20e-1	1.72e+0	1.94e+0	3		2.88e+0	5.93e-1	3.47e+0
7	11	50	4.27e-1	1.89e+0	2.32e+0	4	1/12158	6.64e-1	7.24e-1	1.39e+0
8	28	50	1.39e+0	4.57e+0	5.96e+0	6	50/11075	2.82e+0	1.11e+0	3.93e+0
9	4		2.77e+0	7.79e-1	3.55e+0	4		2.77e+0	7.79e-1	3.55e+0
10	8	1	4.27e-1	1.39e+0	1.82e+0	6	1/9866	6.38e-1	1.09e+0	1.73e+0
11	11	2	4.25e-1	1.89e+0	2.32e+0	11	2/9250	4.25e-1	1.89e+0	2.32e+0
12	20	4	4.23e-1	3.30e+0	3.72e+0	11	4/8675	6.55e-1	1.90e+0	2.56e+0
13	4		2.77e+0	7.94e-1	3.57e+0	14	50/8412	2.11e+0	2.39e+0	4.49e+0
14	5	1	2.20e-1	9.34e-1	1.15e+0	4		2.82e+0	7.96e-1	3.62e+0
15	9	2	2.19e-1	1.54e+0	1.76e+0	7	1/10130	4.63e-1	1.25e+0	1.71e+0
16	17	2	2.18e-1	2.82e+0	3.04e+0	11	2/10129	4.37e-1	1.87e+0	2.31e+0
17	43	50	1.97e+0	6.70e+0	8.67e+0	21	50/10104	2.21e+0	3.47e+0	5.68e+0
18	5		2.77e+0	9.54e-1	3.72e+0	5		2.77e+0	9.54e-1	3.72e+0
19	15	1	4.25e-1	2.59e+0	3.02e+0	13	1/10130	6.58e-1	2.23e+0	2.89e+0
20	47	50	3.00e+0	7.76e+0	1.08e+01	31	50/10130	3.29e+0	5.05e+0	8.34e+0
21	6		2.81e+0	1.10e+0	3.91e+0	6		2.81e+0	1.10e+0	3.91e+0
22	21	7	5.89e-1	3.43e+0	4.02e+0	17	7/10130	8.18e-1	2.70e+0	3.51e+0
23	6		2.76e+0	1.11e+0	3.87e+0	6		2.76e+0	1.11e+0	3.87e+0
24	22	10	4.29e-1	3.58e+0	4.01e+0	19	10/10130	6.63e-1	3.08e+0	3.74e+0
	318		3.12e+1	5.37e+1	8.49e+1	227		3.80e+1	3.91e+1	7.71e+1

Table 5 CVXQP1-M: minimum and maximum eigenvalues of $S_{lr}^{-1}S$ and $S_{cu}^{-1}S$.

CVXQP1-M				
$IP\#$	$\lambda_{\min}(S_{lr}^{-1}S)$	$\lambda_{\max}(S_{lr}^{-1}S)$	$\lambda_{\min}(S_{cu}^{-1}S)$	$\lambda_{\max}(S_{cu}^{-1}S)$
12	4.15e-1	2.34e+0	4.16e-1	1.39e+0
13	2.02e-1	6.79e+0	2.03e-1	2.07e+0
14	8.63e-2	8.83e+0	8.67e-2	2.64e+0
15	3.77e-2	2.42e+1	3.81e-2	4.27e+0
17	2.85e-1	4.75e+0	2.85e-1	2.88e+0
18	1.35e-1	9.56e+0	1.35e-1	5.52e+0
19	6.47e-2	2.33e+1	6.47e-2	9.84e+0
21	1.97e-1	8.69e+0	1.97e-1	3.59e+0
22	2.72e-2	4.70e+1	2.72e-2	7.99e+0
24	1.02e-1	9.58e+0	1.02e-1	2.96e+0
25	1.44e-2	6.23e+1	1.44e-2	1.13e+1

5 Conclusions

We have presented a procedure for updating CPs for sequences of regularized KKT systems with the nonzero (2,2) block having a fixed sparsity pattern. This procedure combines a recently proposed technique for updating a block LDL^T factorization of a given seed CP, relying on a low-rank correction of its Schur complement, with a further update of the Schur complement. The latter update is performed to introduce into the preconditioner some information that has been discarded in the low-rank correction. It is based on a low-cost technique for approximating the Cholesky factorization when the matrix undergoes a diagonal positive semidefinite modification.

Theoretical results show that the procedure proposed here provides the possibility of tightening the bounds on the eigenvalues of the preconditioned matrix with respect to the procedure based on the low-rank correction alone, and that its effectiveness depends on the accuracy of the approximate Cholesky factorization computed in the second updating step. Numerical experiments confirm this behavior and show that the new approach can enhance the low-rank strategy.

Acknowledgements We are grateful to William Hager for helping us to understand some details in the implementation of the sparse Cholesky factorization available from the CHOLMOD package. We also appreciate the anonymous referees for their insightful comments, which led us to improve our work.

Table 6 Comparison of \mathcal{P}_{ex} and \mathcal{P}_{cu} : total number of SQMR iterations, refreshes for \mathcal{P}_{cu} and execution times.

<i>Problem</i>	n m $mz(S)$	I_{Pits}	P_{ex}				P_{cu}			
			n_{it}	T_{prec}	T_{solve}	T_{tot}	n_{it}	n_{ref}	T_{prec}	T_{solve} T_{tot}
CVXQP1-M	20000	11-26	351	4.20e+1	1.83e+1	6.03e+1	597	5	1.92e+1	2.98e+1 4.90e+1
	10000									
	67976									
CVXQP1-M-D	20000	10-28	670	6.36e+1	8.20e+1	1.40e+2	750	6	4.03e+1	9.13e+1 1.32e+2
	10000									
	240494									
CVXQP3-M	20000	11-30	520	2.92e+2	1.13e+2	4.05e+2	943	5	1.08e+2	1.96e+2 3.04e+2
	15000									
	155942									
CVXQP3-M-D2	20000	11-25	424	1.44e+2	1.14e+2	2.57e+2	583	5	7.02e+1	1.55e+2 2.25e+2
	15000									
	224396									
MOSARQP1-D	22500	1-24	92	6.69e+1	1.81e+1	8.50e+1	227	7	3.80e+01	3.91e+01 7.71e+1
	20000									
	4259250									

References

1. V. Baryamureeba, T. Steihaug, Y. Zhang, *Properties of a Class of Preconditioners for Weighted Least Squares Problems*, Technical Report No. 170, Department of Informatics, University of Bergen, and Technical Report No. TR99-16, Department of Computational and Applied Mathematics, Rice University, Houston, 1999.
2. S. Bellavia, *Inexact Interior-Point Method*, Journal of Optimization Theory and Applications, 96 (1998), pp. 109–121.
3. S. Bellavia, V. De Simone, D. di Serafino, B. Morini, *Efficient preconditioner updates for shifted linear systems*, SIAM Journal on Scientific Computing, 33 (2011), pp. 1785–1809.
4. S. Bellavia, V. De Simone, D. di Serafino, B. Morini, *A preconditioning framework for sequences of diagonally modified linear systems arising in optimization*, SIAM Journal on Numerical Analysis, 50 (2012), pp. 3280–3302.
5. S. Bellavia, V. De Simone, D. di Serafino, B. Morini, *Updating constraint preconditioners for KKT systems in quadratic programming via low-rank corrections*, SIAM Journal on Optimization, 25 (2015), pp. 1787–1808.
6. M. Benzi, G. H. Golub, J. Liesen, *Numerical solution of saddle point problems*, Acta Numerica, 14 (2005), pp. 1–137.
7. M. Benzi, V. Simoncini, *On the eigenvalues of a class of saddle point matrices*, Numerische Mathematik, 103 (2006), pp. 173–196.
8. L. Bergamaschi, *Eigenvalue distribution of constraint-preconditioned symmetric saddle point matrices*, Numerical Linear Algebra with Applications, 19 (2012), pp. 754–772.
9. L. Bergamaschi, J. Gondzio, M. Venturin, G. Zilli, *Inexact constraint preconditioners for linear systems arising in interior point methods*, Computational Optimization and Applications, 36 (2007), pp. 137–147.
10. L. Bergamaschi, J. Gondzio, G. Zilli, *Preconditioning Indefinite Systems in Interior Point Methods for Optimization*, Computational Optimization and Applications, 28 (2004), pp. 149–171.
11. S. Cafieri, M. D’Apuzzo, V. De Simone, D. di Serafino, *On the iterative solution of KKT systems in potential reduction software for large-scale quadratic problems*, Computational Optimization and Applications, 38 (2007), pp. 27–45.
12. S. Cafieri, M. D’Apuzzo, V. De Simone, and D. di Serafino, *Stopping criteria for inner iterations in inexact potential reduction methods: a computational study*, Computational Optimization and Applications, 36 (2007), pp. 165–193.
13. S. Cafieri, M. D’Apuzzo, V. De Simone, D. di Serafino, G. Toraldo, *Convergence analysis of an inexact potential reduction method for convex quadratic programming*, Journal of Optimization Theory and Applications, 135 (2007), pp. 355–366.
14. Y. Chen, T.A. Davis, W.W. Hager, S. Rajamanickam, *Algorithm 887: CHOLMOD, Supernodal Sparse Cholesky Factorization and Update/Downdate*, ACM Transactions on Mathematical Software, 35 (2008), pp. 22:1–22:14.
15. M. D’Apuzzo, V. De Simone, D. di Serafino, *On mutual impact of numerical linear algebra and large-scale optimization with focus on interior point methods*, Computational Optimization and Applications, 45 (2010), pp. 283–310.
16. M. D’Apuzzo, V. De Simone, D. di Serafino, *Starting-point strategies for an infeasible potential reduction method*, Optimization Letters, 4 (2010), pp. 131–146.
17. T.A. Davis, W.W. Hager, *Dynamic supernodes in sparse Cholesky update/downdate and triangular solves*, ACM Transactions on Mathematical Software, 35 (2009), pp. 27:1–27:23.
18. H.S. Dollar, *Constraint-style preconditioners for regularized saddle point problems*, SIAM Journal on Matrix Analysis and Applications, 29 (2007), pp. 672–684.
19. H.S. Dollar, N.I.M. Gould, W.H.A. Schilders, A.J. Wathen, *Using constraint preconditioners with regularized saddle-point systems*, Computational Optimization and Applications, 36 (2007), pp. 249–270.
20. H.S. Dollar, A.J. Wathen, *Approximate factorization constraint preconditioners for saddle-point matrices*, SIAM Journal on Scientific Computing, 27 (2006), pp. 1555–1572.
21. C. Durazzi, V. Ruggiero, *Indefinitely preconditioned conjugate gradient method for large sparse equality and inequality constrained quadratic problems*, Numerical Linear Algebra with Applications, 10 (2003), pp. 673–688.

22. R. Freund, N. Nachtigal, *Software for simplified Lanczos and QMR algorithms*, Applied Numerical Mathematics, 19 (1995), pp. 319–341.
23. M.P. Friedlander, D. Orban, *A primal-dual regularized interior-point method for convex quadratic programs*, Mathematical Programming Computation, 4 (2012), pp. 71–107.
24. A. Forsgren P.E. Gill, J.D. Griffin, *Iterative solution of augmented systems arising in interior methods*, SIAM Journal on Optimization, 18 (2007), pp. 666–690.
25. J. Gondzio, *Interior point methods 25 years later*, European Journal of Operational Research, 218 (2012), pp. 587–601.
26. J. Gondzio, *Matrix-free interior point method*, Computational Optimization and Applications, 51 (2012), pp. 457–480.
27. N.I.M. Gould, D. Orban, Ph.L. Toint, *CUTEst: a Constrained and Unconstrained Testing Environment with safe threads*, Technical Report RAL-TR-2013-005, STFC Rutherford Appleton Laboratory, Chilton, Oxfordshire, UK, 2013.
28. C. Keller, N.I.M. Gould, A.J. Wathen, *Constraint preconditioning for indefinite linear systems*, SIAM Journal on Matrix Analysis and Applications, 21 (2000), pp. 1300–1317.
29. L. Lukšan, J. Vlček, *Indefinitely Preconditioned Inexact Newton Method for Large Sparse Equality Constrained Non-linear Programming Problems* Numerical Linear Algebra with Applications, 5 (1998), pp. 219–247.
30. G. Meurant, *On the incomplete Cholesky decomposition of a class of perturbed matrices*, SIAM Journal on Scientific Computing, 23 (2001), pp. 419–429.
31. I. Perugia, V. Simoncini, *Block-diagonal and indefinite symmetric preconditioners for mixed finite element formulations*, Numerical Linear Algebra with Applications, 7 (2000), pp. 585–616.
32. D. Sesana, V. Simoncini, *Spectral analysis of inexact constraint preconditioning for symmetric saddle point matrices*, Linear Algebra and its Applications, 438 (2013), pp. 2683–2700.
33. W. Wang, D.P. O’Leary, *Adaptive use of iterative methods in predictor-corrector interior point methods for linear programming*, Numerical Algorithms, 25 (2000), pp. 387–406.
34. S.J. Wright, *Primal-Dual Interior-Point Methods*, SIAM, Philadelphia, 1997.