

UPDATING CONSTRAINT PRECONDITIONERS FOR KKT SYSTEMS IN QUADRATIC PROGRAMMING VIA LOW-RANK CORRECTIONS*

STEFANIA BELLAVIA[†], VALENTINA DE SIMONE[‡], DANIELA DI SERAFINO^{‡§}, AND
BENEDETTA MORINI[†]

Revised version – July 7, 2015

Abstract. This work focuses on the iterative solution of sequences of KKT linear systems arising in interior point methods applied to large convex quadratic programming problems. This task is the computational core of the interior point procedure and an efficient preconditioning strategy is crucial for the efficiency of the overall method. Constraint preconditioners are very effective in this context; nevertheless, their computation may be very expensive for large-scale problems and resorting to approximations of them may be convenient. Here we propose a procedure for building inexact constraint preconditioners by updating a *seed* constraint preconditioner computed for a KKT matrix at a previous interior point iteration. These updates are obtained through low-rank corrections of the Schur complement of the (1,1) block of the seed preconditioner. The updated preconditioners are analyzed both theoretically and computationally. The results obtained show that our updating procedure, coupled with an adaptive strategy for determining whether to reinitialize or update the preconditioner, can enhance the performance of interior point methods on large problems.

Key words. KKT systems, constraint preconditioners, matrix updates, convex quadratic programming, interior point methods.

AMS subject classifications. 65F08, 65F10, 90C20, 90C51.

1. Introduction. Second-order methods for constrained and unconstrained optimization require, at each iteration, the solution of a system of linear equations. For large-scale problems it is common to solve each system by an iterative method coupled with a suitable preconditioner. Since the computation of a preconditioner for each linear system can be very expensive, in recent years there has been a growing interest in reducing the cost for preconditioning by sharing some computational effort through subsequent linear systems.

Updating preconditioner frameworks for sequences of linear systems have a common feature: based on information generated during the solution of a linear system in the sequence, a preconditioner for a subsequent system is generated. An efficient updating procedure is expected to build a preconditioner which is less effective in terms of linear iterations than the one computed from scratch, but more convenient in terms of cost for the overall linear algebra phase. The approaches existing in literature can be broadly classified as: limited-memory quasi-Newton preconditioners for symmetric positive definite and nonsymmetric matrices (see, e.g., [10, 36, 41]), recycled Krylov information preconditioners for symmetric and nonsymmetric matrices (see, e.g., [18, 30, 32, 38]), updates of factorized preconditioners for symmetric positive definite and nonsymmetric matrices (see, e.g., [3, 4, 5, 6, 7, 17, 25, 40]).

*Work partially supported by INdAM-GNCS, under the 2013 Projects *Strategie risolutive per sistemi lineari di tipo KKT con uso di informazioni strutturali* and *Metodi numerici e software per l'ottimizzazione su larga scala con applicazioni all'immagine processing*.

[†]Dipartimento di Ingegneria Industriale, Università degli Studi di Firenze, viale Morgagni 40, 50134 Firenze, Italy, stefania.bellavia@unifi.it, benedetta.morini@unifi.it.

[‡]Dipartimento di Matematica e Fisica, Seconda Università degli Studi di Napoli, viale A. Lincoln 5, 81100 Caserta, Italy, valentina.desimone@unina2.it, daniela.diserafino@unina2.it.

[§]Istituto di Calcolo e Reti ad Alte Prestazioni, CNR, via P. Castellino 111, 80131 Napoli, Italy.

In this paper we study the problem of preconditioning sequences of KKT systems arising in the solution of the convex quadratic programming (QP) problem

$$(1.1) \quad \begin{aligned} & \text{minimize} && \frac{1}{2}x^T Qx + c^T x, \\ & \text{subject to} && A_1x - s = b_1, \quad A_2x = b_2, \quad x + v = u, \quad (x, s, v) \geq 0 \end{aligned}$$

by Interior Point (IP) methods [34, 46]. Here $Q \in \mathbb{R}^{n \times n}$ is symmetric positive semidefinite, and $A_1 \in \mathbb{R}^{m_1 \times n}$, $A_2 \in \mathbb{R}^{m_2 \times n}$, with $m = m_1 + m_2 \leq n$. Note that s and v are slack variables, used to transform the inequality constraints $A_1x \geq b_1$ and $x \leq u$ into equality constraints.

The updating strategy proposed here concerns constraint preconditioners (CPs) [11, 13, 19, 22, 31, 37, 42]. Using the factorization of a *seed* CP computed for some KKT matrix of the sequence, the factorization of an approximate CP is built for subsequent systems. The resulting preconditioner is a special case of *inexact* CP [12, 26, 39, 42, 44] and is intended to reduce the computational cost while preserving effectiveness. To the best of our knowledge, this is the first attempt to study, both theoretically and numerically, a preconditioner updating technique for sequences of KKT systems arising from (1.1). A computational analysis of the reuse of CPs in consecutive IP steps, which can be regarded as a limit case of preconditioner updating, has been presented in [15].

Concerning preconditioner updates for sequences of systems arising from IP methods, we are aware of the work in [1, 45], where linear programming problems are considered. In this case, the KKT linear systems are reduced to the so-called normal equation form; some systems of the sequence are solved by the Cholesky factorization, while the remaining ones are solved by the Conjugate Gradient method preconditioned with a low-rank correction of the last computed Cholesky factor.

Motivated by [1, 45], in this work we adapt the low-rank corrections given therein to our problem. In our approach the updated preconditioner is an inexact CP where the Schur complement of the (1,1) block is replaced by a low-rank modification of the corresponding Schur complement in the seed preconditioner. The validity of the proposed procedure is supported by a spectral analysis of the preconditioned matrix and by numerical results illustrating its performance.

The paper is organized as follows. In Section 2 we provide preliminaries on CPs and new spectral analysis results for general inexact CPs. In Section 3 we present our updating procedure and specialize the spectral analysis conducted in the previous section to the updated preconditioners. In Section 4 we discuss implementation issues of the updating procedure and present numerical results obtained by solving sequences of linear systems arising in the solution of convex QP problems. These results show that our updating technique is able to reduce the computational cost for solving the overall sequence whenever the updating strategy is performed in conjunction with adaptive strategies for determining whether to recompute the seed preconditioner or update the current one.

In the following, $\|\cdot\|$ denotes the vector or matrix 2-norm and, for any symmetric matrix A , $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ denote its minimum and maximum eigenvalues. Finally, for any complex number α , $\mathcal{R}(\alpha)$ and $\mathcal{I}(\alpha)$ denote the real and imaginary parts of α , respectively.

2. Inexact Constraints Preconditioners. In this section we discuss the features of the KKT matrices arising in the solution of problem (1.1) by IP methods and present spectral properties of both CPs and inexact CPs. This analysis will be used

to develop our updating strategy. Throughout the paper we assume that the matrix $A = \begin{bmatrix} A_1 \\ A_2 \end{bmatrix} \in \mathbb{R}^{m \times n}$ is full rank.

The application of an IP method to problem (1.1) gives rise to a sequence of symmetric indefinite matrices differing by a diagonal, possibly indefinite, matrix. In fact, at the k th IP iteration, the KKT matrix takes the form

$$\mathcal{A}_k = \begin{bmatrix} Q + \Theta_k^{(1)} & A^T \\ A & -\Theta_k^{(2)} \end{bmatrix},$$

where $\Theta_k^{(1)} \in \mathbb{R}^{n \times n}$ is diagonal positive definite and $\Theta_k^{(2)} \in \mathbb{R}^{m \times m}$ is diagonal positive semidefinite. In particular,

$$\begin{aligned} \Theta_k^{(1)} &= X_k^{-1}W_k + V_k^{-1}T_k, \\ \Theta_k^{(2)} &= \begin{bmatrix} Y_k^{-1}S_k & 0 \\ 0 & 0 \end{bmatrix}, \end{aligned}$$

where (x_k, w_k) , (s_k, y_k) , and (v_k, t_k) are the pairs of complementary variables of problem (1.1) evaluated at the current iteration, and X_k, W_k, S_k, Y_k, V_k and T_k are the corresponding diagonal matrices according to the standard IP notation. If the QP problem has no linear inequality constraints then $\Theta_k^{(2)}$ is the zero matrix; otherwise $\Theta_k^{(2)}$ admits positive diagonal entries corresponding to slack variables for linear inequality constraints.

To simplify the notation, in the rest of the paper we drop the iteration index k from \mathcal{A}_k , $\Theta_k^{(1)}$ and $\Theta_k^{(2)}$. Hence,

$$(2.1) \quad \mathcal{A} = \mathcal{A}_k$$

and the CP for \mathcal{A} is given by

$$(2.2) \quad \mathcal{P}_{ex} = \begin{bmatrix} G & A^T \\ A & -\Theta^{(2)} \end{bmatrix},$$

where G is an approximation to $Q + \Theta^{(1)}$. Clearly, the application of \mathcal{P}_{ex} requires its factorization.

In order to define \mathcal{P}_{ex} , one possibility is to specify G and then explicitly factorize the preconditioner; alternatively, implicit factorizations can be used, where \mathcal{P}_{ex} is derived from specially chosen factors M and C in $\mathcal{P}_{ex} = MCM^T$ (see, e.g., [23, 24]). We follow the former approach and choose G as the diagonal matrix with the same diagonal entries as $Q + \Theta^{(1)}$, i.e.,

$$(2.3) \quad G = \text{diag}(Q + \Theta^{(1)}).$$

The resulting matrix $\mathcal{P}_{ex}^{-1}\mathcal{A}$ has an eigenvalue at 1 with multiplicity $2m - p$, with $p = \text{rank}(\Theta^{(2)})$, and $n - m + p$ real positive eigenvalues such that the better G approximates $Q + \Theta^{(1)}$ the more clustered around 1 they are [22, 37].

The factorization of \mathcal{P}_{ex} can be obtained by computing a Cholesky-like factorization of the negative Schur complement, S , of G in \mathcal{A} ,

$$(2.4) \quad S = AG^{-1}A^T + \Theta^{(2)},$$

and using the block decomposition

$$(2.5) \quad \mathcal{P}_{ex} = \begin{bmatrix} I_n & 0 \\ AG^{-1} & I_m \end{bmatrix} \begin{bmatrix} G & 0 \\ 0 & -S \end{bmatrix} \begin{bmatrix} I_n & G^{-1}A^T \\ 0 & I_m \end{bmatrix},$$

where I_r is the identity matrix of dimension r , for any integer r .

In problems where a large part of the computational cost for solving the linear system depends on the computation of a Cholesky-like factorization of S , this matrix may be replaced by a computationally cheaper approximation of it [26, 39, 42]. This yields the following *inexact* CP:

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

where S_{inex} is the approximation of S .

2.1. Analysis of the eigenvalues. Due to the indefiniteness of $\mathcal{P}_{inex}^{-1}\mathcal{A}$ and the specific form of \mathcal{P}_{inex} , the preconditioned linear system must be solved by either nonsymmetric solvers, such as GMRES [43] and QMR [27], or by the simplified version of the QMR method [28]; such a simplified solver, named SQMR, is capable of exploiting the symmetry of \mathcal{A} and \mathcal{P}_{inex} and is used in our numerical experiments.

Although the convergence of many Krylov solvers is not fully characterized by the spectral properties of the coefficient matrix, in many practical cases it depends to a large extent on the eigenvalue distribution. For this reason, exploiting the spectral analysis of $\mathcal{P}_{inex}^{-1}\mathcal{A}$ made in the general context of saddle point problems [8, 9, 44], we provide bounds on the eigenvalues which will guide our preconditioner updates.

In order to carry out our analysis we write the eigenvalue problem for the matrix $\mathcal{P}_{inex}^{-1}\mathcal{A}$ in the form

$$(2.6) \quad \begin{bmatrix} Q + \Theta^{(1)} & A^T \\ A & -\Theta^{(2)} \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \lambda \mathcal{P}_{inex} \begin{bmatrix} x \\ y \end{bmatrix}.$$

Starting from results in [8], in the next theorem we give bounds on the eigenvalues which highlight the dependence of the spectrum of $\mathcal{P}_{inex}^{-1}\mathcal{A}$ on that of $S_{inex}^{-1}S$.

THEOREM 2.1. *Let λ and $[x^T, y^T]^T$ be an eigenvalue of problem (2.6) and a corresponding eigenvector. Let*

$$(2.7) \quad X = G^{-\frac{1}{2}}(Q + \Theta^{(1)})G^{-\frac{1}{2}}$$

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

$$(2.8) \quad \bar{\lambda} = \lambda_{\max}(S_{inex}^{-1}S) \max\{2 - \lambda_{\min}(X), 1\},$$

$$(2.9) \quad \underline{\lambda} = \lambda_{\min}(S_{inex}^{-1}S) \min\{2 - \lambda_{\max}(X), 1\},$$

if $\Theta^{(2)} \neq 0$, and

$$(2.10) \quad \bar{\lambda} = \lambda_{\max}(S_{inex}^{-1}S) (2 - \lambda_{\min}(X)),$$

$$(2.11) \quad \underline{\lambda} = \lambda_{\min}(S_{inex}^{-1}S) (2 - \lambda_{\max}(X)),$$

otherwise.

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

$$(2.12) \quad \frac{1}{2}(\lambda_{\min}(X) + \underline{\lambda}) \leq \mathcal{R}(\lambda) \leq \frac{1}{2}(\lambda_{\max}(X) + \bar{\lambda}),$$

otherwise

$$(2.13) \quad \min\{\lambda_{\min}(X), \underline{\lambda}\} \leq \lambda \leq \max\{\lambda_{\max}(X), \bar{\lambda}\}, \quad \text{for } y \neq 0,$$

$$(2.14) \quad \lambda_{\min}(X) \leq \lambda \leq \lambda_{\max}(X), \quad \text{for } y = 0.$$

Finally, the imaginary part of λ satisfies

$$(2.15) \quad |\mathcal{I}(\lambda)| \leq \sqrt{\lambda_{\max}(S_{\text{inex}}^{-1}AG^{-1}A^T)}\|I_n - X\|.$$

Proof. Let

$$S_{\text{inex}} = RR^T,$$

be the Cholesky factorization of S_{inex} with R lower triangular. With a little algebra, the eigenvalue problem (2.6) becomes

$$\begin{bmatrix} X & Y^T \\ Y & -Z \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \lambda \begin{bmatrix} I_n & 0 \\ 0 & -I_m \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix},$$

where X is the matrix in (2.7) and

$$(2.16) \quad \begin{aligned} Y &= R^{-1}AG^{-\frac{1}{2}}(I_n - X), \\ Z &= R^{-1}AG^{-\frac{1}{2}}(2I_n - X)G^{-\frac{1}{2}}A^TR^{-T} + R^{-1}\Theta^{(2)}R^{-T}, \\ \begin{bmatrix} u \\ v \end{bmatrix} &= \begin{bmatrix} G^{\frac{1}{2}} & G^{-1/2}A^T \\ 0 & R^T \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \end{aligned}$$

(see [8]).

From [8, Proposition 2.3] it follows that $\mathcal{P}_{\text{inex}}^{-1}\mathcal{A}$ has at most $2m$ eigenvalues with nonzero imaginary part, counting conjugates. Furthermore, [8, Proposition 2.12] shows that, when Y is full rank and Z is positive semidefinite, if λ has nonzero imaginary part, then

$$(2.17) \quad \frac{1}{2}(\lambda_{\min}(X) + \lambda_{\min}(Z)) \leq \mathcal{R}(\lambda) \leq \frac{1}{2}(\lambda_{\max}(X) + \lambda_{\max}(Z)),$$

otherwise

$$(2.18) \quad \min\{\lambda_{\min}(X), \lambda_{\min}(Z)\} \leq \lambda \leq \max\{\lambda_{\max}(X), \lambda_{\max}(Z)\}, \quad \text{for } v \neq 0,$$

$$(2.19) \quad \lambda_{\min}(X) \leq \lambda \leq \lambda_{\max}(X), \quad \text{for } v = 0.$$

In addition, the imaginary part of λ satisfies

$$(2.20) \quad |\mathcal{I}(\lambda)| \leq \|Y\|.$$

The same results hold if Y is rank deficient with Z positive definite on $\ker(Y^T)$.

By assumptions, Z is positive definite and hence inequalities (2.17)–(2.20) hold. Furthermore, since $v = R^Ty$ and R is nonsingular, (2.14) is equivalent to (2.19). It remains to prove (2.12), (2.13) and (2.15); we proceed by bounding $\lambda_{\min}(Z)$, $\lambda_{\max}(Z)$ and $\|Y\|$.

We consider first the case $\Theta^{(2)} \neq 0$. In order to provide an upper bound on $\lambda_{\max}(Z)$, we note that

$$Z = R^{-1} \begin{bmatrix} AG^{-\frac{1}{2}} & (\Theta^{(2)})^{\frac{1}{2}} \end{bmatrix} \begin{bmatrix} 2I_n - X & 0 \\ 0 & I_m \end{bmatrix} \begin{bmatrix} G^{-\frac{1}{2}}A^T \\ (\Theta^{(2)})^{\frac{1}{2}} \end{bmatrix} R^{-T}.$$

Let VU be the rank-retaining factorization of $\begin{bmatrix} G^{-\frac{1}{2}}A^T \\ (\Theta^{(2)})^{\frac{1}{2}} \end{bmatrix}$, where $V \in \mathbb{R}^{(n+m) \times m}$ has orthogonal columns and $U \in \mathbb{R}^{m \times m}$ is upper triangular and nonsingular. Then, $S = U^T U$ and

$$(2.21) \quad Z = R^{-1}U^T V^T \begin{bmatrix} 2I_n - X & 0 \\ 0 & I_m \end{bmatrix} VUR^{-T}.$$

Letting $N = R^{-1}U^T$, we have

$$\lambda_{\max}(Z) = \|Z\| \leq \|N\|^2 \max\{2 - \lambda_{\min}(X), 1\}.$$

Since $N^T N$ is similar to $S_{inex}^{-1} S$, the upper bounds in (2.12) and (2.13), with $\bar{\lambda}$ given in (2.8), follow from the upper bounds in (2.17) and (2.18).

In order to provide a lower bound on $\lambda_{\min}(Z)$, we observe that

$$(2.22) \quad \frac{1}{\lambda_{\min}(Z)} = \|Z^{-1}\| \leq \|N^{-1}\|^2 \left\| \left(V^T \begin{bmatrix} 2I_n - X & 0 \\ 0 & I_m \end{bmatrix} V \right)^{-1} \right\|.$$

By noting that $N^{-T} N^{-1}$ is similar to $S^{-1} S_{inex}$ and hence

$$\|N^{-1}\|^2 = \frac{1}{\lambda_{\min}(S_{inex}^{-1} S)},$$

and by applying the Courant-Fischer minimax characterization (see, e.g., [33, Theorem 8.1.2]) to the rightmost term in (2.22), we have that the lower bounds in (2.12) and (2.13), with $\underline{\lambda}$ given in (2.9), follow from the lower bounds in (2.17) and (2.18).

When $\Theta^{(2)} = 0$, inequalities (2.12)–(2.13) can be derived by assuming that VU is the rank-retaining factorization of $G^{-\frac{1}{2}}A^T$, with $V \in \mathbb{R}^{n \times m}$ having orthogonal columns and $U \in \mathbb{R}^{m \times m}$ upper triangular and nonsingular. In this case equality (2.21) becomes

$$Z = R^{-1}U^T V^T (2I_n - X) VUR^{-T}$$

and the thesis follows by reasoning as above.

Finally, from (2.16) it follows that

$$(2.23) \quad \|Y\| \leq \|R^{-1}AG^{-\frac{1}{2}}\| \|I_n - X\| = \sqrt{\lambda_{\max}(R^{-1}AG^{-1}A^T R^{-T})} \|I_n - X\|.$$

Inequality (2.15) follows from (2.20) and (2.23) by noting that $R^{-1}AG^{-1}A^T R^{-T}$ and $S_{inex}^{-1}AG^{-1}A^T$ are similar. \square

REMARK 2.1. *Note that the assumption on $2I_n - X$ in the theorem can be fulfilled by a proper scaling of X enforcing its eigenvalues to be smaller than 2. Furthermore, if Q is diagonal, then $X = I_n$ and from (2.15) it follows that 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 [26]).*

3. Building an inexact constraint preconditioner by updates. In this section we design a strategy for updating the CP built for some *seed* matrix of the KKT sequence. The update is based on low-rank corrections and generates inexact constraint preconditioners for subsequent systems.

Let us consider a KKT matrix \mathcal{A}_{seed} generated at some iteration r of the IP procedure,

$$(3.1) \quad \mathcal{A}_{seed} = \begin{bmatrix} Q + \Theta_{seed}^{(1)} & A^T \\ A & -\Theta_{seed}^{(2)} \end{bmatrix},$$

where $\Theta_{seed}^{(1)} \in \mathbb{R}^{n \times n}$ is diagonal positive definite and $\Theta_{seed}^{(2)} \in \mathbb{R}^{m \times m}$ is diagonal positive semidefinite. The corresponding *seed* CP has the form

$$(3.2) \quad \mathcal{P}_{seed} = \begin{bmatrix} H & A^T \\ A & -\Theta_{seed}^{(2)} \end{bmatrix},$$

where $H = \text{diag}(Q + \Theta_{seed}^{(1)})$. Assume that a block factorization of \mathcal{P}_{seed} has been obtained by computing the Cholesky-like factorization of the negative Schur complement, S_{seed} , of H in \mathcal{P}_{seed} ,

$$(3.3) \quad S_{seed} = AH^{-1}A^T + \Theta_{seed}^{(2)} = LDL^T,$$

where L is unit lower triangular and D is diagonal positive definite.

Let \mathcal{A} in (2.1) be a subsequent matrix of the KKT sequence, \mathcal{P}_{ex} in (2.5) the corresponding CP, with G defined as in (2.3), and S the Schur complement (2.4). We replace S with a suitable update of S_{seed} , named S_{upd} , obtaining the inexact preconditioner

$$(3.4) \quad \mathcal{P}_{upd} = \begin{bmatrix} I_n & 0 \\ AG^{-1} & I_m \end{bmatrix} \begin{bmatrix} G & 0 \\ 0 & -S_{upd} \end{bmatrix} \begin{bmatrix} I_n & G^{-1}A^T \\ 0 & I_m \end{bmatrix}.$$

In the remainder of this section we show that specific choices of S_{upd} provide easily computable bounds on the eigenvalues of $S_{upd}^{-1}S$; we use these bounds, together with the spectral analysis in Section 2, for constructing practical inexact preconditioners by updating techniques. For ease of presentation, we consider the cases $\Theta^{(2)} = 0$ and $\Theta^{(2)} \neq 0$ separately.

3.1. Updated preconditioners for $\Theta^{(2)} = 0$. We consider a low-rank update/downdate of $S_{seed} = AH^{-1}A^T = LDL^T$ that generates a matrix S_{upd} of the form

$$(3.5) \quad S_{upd} = AJ^{-1}A^T = L_{upd}D_{upd}L_{upd}^T,$$

where J is a suitably chosen diagonal positive definite matrix, L_{upd} is unit lower triangular and D_{upd} is diagonal positive definite.

A guideline for choosing J is provided by Theorem 2.1 and by a result in [1], reported next for completeness.

LEMMA 3.1. *Let $B \in \mathbb{R}^{m \times n}$ be full rank and let $E, F \in \mathbb{R}^{n \times n}$ be symmetric and positive definite. Then, any eigenvalue λ of $(BEB^T)^{-1}BFB^T$ satisfies*

$$(3.6) \quad \lambda_{\min}(E^{-1}F) \leq \lambda((BEB^T)^{-1}BFB^T) \leq \lambda_{\max}(E^{-1}F).$$

For any diagonal matrix $W \in \mathbb{R}^n$, let $\gamma(W) = (\gamma_1(W), \dots, \gamma_n(W))^T$ be the vector with entries given by the diagonal entries of WG^{-1} sorted in nondecreasing order, i.e.,

$$(3.7) \quad \min_{1 \leq i \leq n} \frac{W_{ii}}{G_{ii}} = \gamma_1(W) \leq \gamma_2(W) \leq \dots \leq \gamma_n(W) = \max_{1 \leq i \leq n} \frac{W_{ii}}{G_{ii}}.$$

By Lemma 3.1, the eigenvalues of $S_{upd}^{-1}S$ can be bounded by using the diagonal entries of JG^{-1} , i.e.,

$$(3.8) \quad \gamma_1(J) \leq \lambda(S_{upd}^{-1}S) \leq \gamma_n(J);$$

then, Theorem 2.1 implies the following result.

COROLLARY 3.2. *Let \mathcal{A} , \mathcal{P}_{upd} and S_{upd} be the matrices in (2.1), (3.4) and (3.5), respectively, and λ an eigenvalue of $\mathcal{P}_{upd}^{-1}\mathcal{A}$. Let $\gamma_1(J)$ and $\gamma_n(J)$ be defined according to (3.7) and let X be the matrix in (2.7). If $2I_n - X$ is positive definite, then (2.12)–(2.14) hold with*

$$(3.9) \quad \bar{\lambda} \leq \gamma_n(J) (2 - \lambda_{\min}(X)),$$

$$(3.10) \quad \underline{\lambda} \geq \gamma_1(J) (2 - \lambda_{\max}(X)).$$

Furthermore,

$$(3.11) \quad |\mathcal{I}(\lambda)| \leq \sqrt{\gamma_n(J)} \|I_n - X\|.$$

The key issue is to define J such that a good tradeoff can be achieved between the effectiveness of the bounds provided by Corollary 3.2 and the cost for building the factors L_{upd} and D_{upd} in (3.5). On the basis of this consideration, following [1] we define S_{upd} as a low-rank correction of S_{seed} of the form

$$S_{upd} = AH^{-1}A^T + \bar{A}\bar{K}\bar{A}^T = A(H^{-1} + K)A^T,$$

where $K \in \mathbb{R}^{n \times n}$ is a diagonal matrix with only $q < n$ nonzero entries on the diagonal, $\bar{K} \in \mathbb{R}^{q \times q}$ is the principal submatrix of K containing these nonzero entries, $\bar{A} \in \mathbb{R}^{m \times q}$ is made of the corresponding q columns of A , and $J^{-1} = H^{-1} + K$ accounts for major changes from H to G . In order to choose K , we consider the vector $\gamma(H)$, defined according to (3.7), and the vector of integers $l = (l_1, \dots, l_n)^T$ such that l_i gives the position of H_{ii}/G_{ii} in $\gamma(H)$, i.e.,

$$\gamma_{l_i}(H) = \frac{H_{ii}}{G_{ii}},$$

and define the set

$$(3.12) \quad \Gamma = \{i : n - q_1 + 1 \leq l_i \leq n \text{ and } \gamma_{l_i}(H) > \mu_\gamma\} \\ \cup \{i : 1 \leq l_i \leq q_2 \text{ and } \gamma_{l_i}(H) < \nu_\gamma\},$$

where $\mu_\gamma \geq 1$, $\nu_\gamma \in (0, 1]$, and q_1 and q_2 are nonnegative integers such that $q_1 + q_2 = q$. Then we define the diagonal matrix K by setting

$$(3.13) \quad K_{ii} = \begin{cases} G_{ii}^{-1} - H_{ii}^{-1}, & \text{if } i \in \Gamma, \\ 0, & \text{otherwise,} \end{cases}$$

Algorithm 3.1 (Building S_{upd} via low-rank update)

$[L_{upd}, D_{upd}] = \mathbf{lr_update}(H, G, A, L, D, q_1, q_2, \mu_\gamma, \nu_\gamma)$

- 1: Build the vector $\gamma(H)$, containing the diagonal entries of HG^{-1} sorted in nondecreasing order, and the vector l , containing the positions of the diagonal entries of HG^{-1} in $\gamma(H)$.
- 2: Build the set Γ in (3.12) and the vector $C_\Gamma \in \mathbb{R}^q$ that contains the indices $i \in \Gamma$ in increasing order.
- 3: Build the matrix \bar{K} by defining its diagonal entries as

$$\bar{K}_{ss} = G_{jj}^{-1} - H_{jj}^{-1}, \quad j = C_\Gamma(s), \quad s = 1, \dots, q,$$

and set $\bar{A} = A(:, C_\Gamma)$.

- 4: Compute the factorization $L_{upd}D_{upd}L_{upd}^T$ of $S_{upd} = S_{seed} + \bar{A}\bar{K}\bar{A}^T$ by updating/downdating the factors L and D .
-

and choose \bar{K} as the principal submatrix of K having as diagonal entries the values K_{ii} with index $i \in \Gamma$. Finally, we define \bar{A} as the matrix consisting of the columns of A corresponding to the indices in Γ , ordered as they are in A .

The matrix S_{upd} is positive definite and $S_{upd} - S_{seed}$ is a low-rank matrix if the cardinality of Γ is small, i.e., $q \ll n$. In this case the Cholesky-like factorization of S_{upd} can be conveniently computed by updating or downdating the factorization LDL^T of S_{seed} . Specifically, an update must be performed if $H_{ii} > G_{ii}$, and a downdate if $H_{ii} < G_{ii}$. This task can be accomplished by either using efficient procedures for updating and downdating the Cholesky factorization [21], or by the Sherman-Morrison-Woodbury formula (see, e.g., [33, Section 2.1.3]). Clearly, once S_{upd} has been factorized, the factorization of \mathcal{P}_{upd} is readily available.

The procedure described so far is summarized in Algorithm 3.1 and is called *lr_update*. It takes in input the matrices H , G and A , the factors L and D of S_{seed} , and the scalars q_1 , q_2 , μ_γ , and ν_γ , and returns in output the factors L_{upd} and D_{upd} of S_{upd} . For simplicity, we assume that the first set on the right-hand side of (3.12) contains q_1 elements and the second set contains q_2 elements. Note that we borrow the Matlab notation.

The previous choice of K implies that the matrix J has the following diagonal entries:

$$J_{ii} = \begin{cases} G_{ii}, & \text{if } i \in \Gamma, \\ H_{ii}, & \text{otherwise,} \end{cases}$$

and hence

$$(3.14) \quad \gamma_1(J) = \min\{1, \min_{i \notin \Gamma} \gamma_{l_i}(H)\} = \min\{1, \gamma_{q_2+1}(H)\},$$

$$(3.15) \quad \gamma_n(J) = \max\{1, \max_{i \notin \Gamma} \gamma_{l_i}(H)\} = \max\{1, \gamma_{n-q_1}(H)\}.$$

Then, from Corollary 3.2 it follows that, among the possible choices of J , the one considered here is expected to provide good eigenvalue bounds as long as $\gamma_{q_2}(H)$ and $\gamma_{n-q_1+1}(H)$ are well separated from $\gamma_{q_2+1}(H)$ and $\gamma_{n-q_1}(H)$, respectively. We




FIG. 3.1. Spectra of the matrices \mathcal{A} , $\mathcal{P}_{ex}^{-1}\mathcal{A}$, and $\mathcal{P}_{upd}^{-1}\mathcal{A}$ with $q_1 = q_2 = 25$, at the 10th iteration of an IP method applied to problem CVXQP1 ($n = 1000$, $m = 500$). The updated preconditioner is built from the Schur complement of the exact preconditioner at the 6th IP iteration.

observe that by setting $\mu_\gamma = \nu_\gamma = 1$ we allow to include in Γ indices corresponding to values $\gamma_{l_i}(H)$ close to 1, which may not change much the bounds provided by (3.11) and by (2.12)–(2.14) with (3.9)–(3.10). Therefore, the use of larger (smaller) values for μ_γ (ν_γ) may be effective anyway, while saving computational cost (see Section 4 for further details). Furthermore, a consequence of this low-rank correction strategy is that $S - S_{upd}$ has q zero eigenvalues; thus $\mathcal{P}_{inex}^{-1}\mathcal{A}_k$ has $2q$ unit eigenvalues with geometric multiplicity q [44, Theorem 3.3].

We note that in the limit case $q = 0$ the set Γ is empty; hence $S_{upd} = S_{seed}$ and

$$\gamma_{l_i}(H) = \frac{Q_{ii} + (\Theta_{seed}^{(1)})_{ii}}{Q_{ii} + \Theta_{ii}^{(1)}}.$$

The element $\gamma_{l_i}(H)$ is expected to be close to 1 if $(\Theta^{(1)})_{ii} - (\Theta_{seed}^{(1)})_{ii}$ is small, while it may significantly differ from 1 if $\Theta_{ii}^{(1)}$ tends to zero or infinity, as it happens when the IP iterate approaches an optimal solution where strict complementarity holds.

We conclude this section showing the spectra of the matrices \mathcal{A} , $\mathcal{P}_{ex}^{-1}\mathcal{A}$ and $\mathcal{P}_{upd}^{-1}\mathcal{A}$, where \mathcal{A} has been obtained by applying an IP solver to problem CVXQP1 from the CUTEst collection [35], with dimensions $n = 1000$ and $m = 500$ (see Figure 3.1). The IP solver, described in Section 4, has been run using \mathcal{P}_{ex} , and \mathcal{A} is the KKT matrix at

the 10th IP iteration. The updated preconditioner \mathcal{P}_{upd} has been built by updating the matrix \mathcal{A}_{seed} obtained at the 6th IP iteration, using $q_1 = q_2 = 25$. In this case $\gamma_{q_2+1} = 1.85e-3$ and $\gamma_{n-q_1} = 2.73e+0$. We see that, unlike \mathcal{P}_{ex} , \mathcal{P}_{upd} moves some eigenvalues from the real to the complex field. Nevertheless, \mathcal{P}_{upd} tends to cluster the eigenvalues of \mathcal{A} around 1, and γ_{q_2+1} and γ_{n-q_1} provide approximate bounds on the real and imaginary parts of the eigenvalues of the preconditioned matrix, according to (3.11)–(3.10) and (3.14)–(3.15). Of course, this clustering is less effective than the one performed by \mathcal{P}_{ex} , but it is useful in several cases, as shown in Section 4.

3.2. Updated preconditioners for $\Theta^{(2)} \neq 0$. The updating strategy described in the previous section can be generalized to the case $\Theta^{(2)} \neq 0$. To this end, we note that the sparsity pattern of $\Theta^{(2)}$ does not change throughout the IP iterations and the set $\mathcal{L} = \{i : \Theta_{ii}^{(2)} \neq 0\}$ has cardinality equal to the number m_1 of linear inequality constraints in the QP problem. Let $\tilde{\Theta}_{seed}^{(2)}$ and $\tilde{\Theta}^{(2)}$ be the $m_1 \times m_1$ diagonal submatrices containing the nonzero diagonal entries of $\Theta_{seed}^{(2)}$ and $\Theta^{(2)}$, respectively, and let \tilde{I}_m be the rectangular matrix consisting of the columns of I_m with indices in \mathcal{L} . Then, we have

$$\begin{aligned} S_{seed} &= AH^{-1}A^T + \Theta_{seed}^{(2)} = \tilde{A}\tilde{H}^{-1}\tilde{A}^T, \\ S &= AG^{-1}A^T + \Theta^{(2)} = \tilde{A}\tilde{G}^{-1}\tilde{A}^T, \end{aligned}$$

where

$$\tilde{A} = [A \quad \tilde{I}_m], \quad \tilde{H}^{-1} = \begin{bmatrix} H^{-1} & 0 \\ 0 & \tilde{\Theta}_{seed}^{(2)} \end{bmatrix}, \quad \tilde{G}^{-1} = \begin{bmatrix} G^{-1} & 0 \\ 0 & \tilde{\Theta}^{(2)} \end{bmatrix}.$$

Analogously, letting

$$(3.16) \quad S_{upd} = AJ^{-1}A^T + \Theta_{upd}^{(2)},$$

we have

$$S_{upd} = \tilde{A}\tilde{J}^{-1}\tilde{A}^T, \quad \text{where} \quad \tilde{J}^{-1} = \begin{bmatrix} J^{-1} & 0 \\ 0 & \tilde{\Theta}_{upd}^{(2)} \end{bmatrix},$$

and $\tilde{\Theta}_{upd}^{(2)}$ is the $m_1 \times m_1$ diagonal submatrix of $\Theta_{upd}^{(2)}$ containing its nonzero diagonal entries. Thus, we can choose \tilde{J} using the same arguments as in the previous section. With a little abuse of notation, let $\gamma(\tilde{J}) = (\gamma_1(\tilde{J}), \dots, \gamma_{n+m_1}(\tilde{J}))$ be the vector with elements equal to the diagonal entries of $\tilde{J}\tilde{G}^{-1}$ sorted in nondecreasing order. Then, by Lemma 3.1, the eigenvalues of $S_{upd}^{-1}S$ satisfy

$$(3.17) \quad \gamma_1(\tilde{J}) \leq \lambda(S_{upd}^{-1}S) \leq \gamma_{n+m_1}(\tilde{J}),$$

and the following result holds.

COROLLARY 3.3. *Let \mathcal{A} , \mathcal{P}_{upd} and S_{upd} be the matrices in (2.1), (3.4) and (3.16), respectively, and λ an eigenvalue of $\mathcal{P}_{upd}^{-1}\mathcal{A}$. Let $\gamma_1(\tilde{J})$ and $\gamma_{n+m_1}(\tilde{J})$ be the smallest and the largest element of $\gamma(\tilde{J})$ and let X be the matrix in (2.7). If $2I_n - X$ is positive definite, then (2.12)–(2.14) hold with*

$$(3.18) \quad \bar{\lambda} \leq \gamma_{n+m_1}(\tilde{J}) \max\{2 - \lambda_{\min}(X), 1\},$$

$$(3.19) \quad \underline{\lambda} \geq \gamma_1(\tilde{J}) \min\{2 - \lambda_{\max}(X), 1\}.$$

Furthermore,

$$(3.20) \quad |\mathcal{I}(\lambda)| \leq \sqrt{\gamma_{n+m_1}(\tilde{J})} \|I_n - X\|.$$

Proof. Inequalities (3.18) and (3.19) follow directly from Theorem 2.1 and (3.17). Since $\Theta^{(2)}$ is positive semidefinite, for any vector $w \in \mathbb{R}^n$ we have

$$w^T S_{upd}^{-\frac{1}{2}} A G^{-1} A^T S_{upd}^{-\frac{1}{2}} w \leq w^T S_{upd}^{-\frac{1}{2}} (A G^{-1} A^T + \Theta^{(2)}) S_{upd} w.$$

Then, by matrix similarity,

$$(3.21) \quad \lambda_{\max}(S_{upd}^{-1} A G^{-1} A^T) \leq \lambda_{\max}(S_{upd}^{-1} S)$$

and (3.20) follows by using Theorem 2.1 and (3.17). \square

On the basis of the previous results, the generalization of the updating procedure to the case $\Theta^{(2)} \neq 0$ is straightforward. If $S_{seed} = LDL^T$, the factorization of \mathcal{P}_{upd} can be computed by invoking procedure *lr_update* with input data \tilde{H} , \tilde{G} and \tilde{A} in place of H , G and A .

4. Numerical results. We tested the effectiveness of our updating procedure by solving sequences of KKT systems arising in the solution of convex QP problems where $m \leq n$ and A is full rank.

To this end, we implemented *lr_update* within PRQP, a Fortran 90 solver for convex QP problems based on a primal-dual inexact Potential Reduction IP method [13, 16, 19]. For comparison purpose, we implemented the preconditioner \mathcal{P}_{ex} too, in the form specified in (2.5). We used the CHOLMOD library [21] to compute the sparse LDL^T factorization of S_{seed} and S , and to perform the low-rank updates and dndates required by S_{upd} . For the solution of the KKT systems, we developed an implementation of the left-preconditioned SQMR method without look-ahead [28], taking into account the block structure of the system matrices and of the exact and updated preconditioners. All the new code was written in Fortran 90, with interfaces to the functions of CHOLMOD, written in C. We note that only one matrix-vector product per iteration is performed in our SQMR implementation, except in the last few iterations, where an additional matrix-vector product per iteration is computed to use the residual instead of the preconditioned BCG-residual in the stopping criterion, as in the QMRPACK code [29]. This keeps the computational cost per iteration comparable with that of the Conjugate Gradient method. We also observe that, although no theoretical convergence estimates are available for SQMR, this method has shown good performance in all our experiments.

PRQP was run on several test problems, either taken from the CUTEst collection [35] or obtained by modifying CUTEst problems, as explained later in this section. The starting point was chosen as explained in [20] and the IP iterations were stopped when the relative duality gap and suitable measures of the primal and dual infeasibilities became lower than 10^{-7} and 10^{-8} , respectively (see [14] for the details). The zero vector was used as starting guess in SQMR. An adaptive criterion was applied to stop the iterations [14], 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]. A maximum number of 1000 SQMR iterations was considered too, but it was never reached in our runs.

Concerning the choice of q , some comments are in order. As the value of q increases, the updated preconditioner is expected to improve its effectiveness, reducing the number of linear iterations. On the other hand, its computational cost is also expected to increase, because of the growing cost of the low-rank modification. In order to reduce the time for the solution of the overall KKT sequence, the updating strategy must realize a tradeoff between effectiveness and cost, and our experience has shown that q must be much smaller than the dimension of the Schur complement. On the basis of these considerations, we also think that heuristic rules for choosing q dynamically, such as the one proposed in [45], may have limited impact of our procedure because only a small range of values of q is affordable. However, we postpone a systematic study of this issue to future work. In the experiments we set $q = 50, 100$ and $q_1 = q_2 = q/2$, which is much smaller than the dimension of the Schur complement in our test problems. We also considered “the limit case” $q = 0$, corresponding to $S_{upd} = S_{seed}$ in the updated preconditioner \mathcal{P}_{upd} . Preliminary experiments with larger values of q , i.e., $q = 150, 200$, did not lead to any performance improvement and therefore we decided to discard these values. Furthermore, since we had not obtained practical benefits by including in Γ indices corresponding to values of $\gamma_{l_i}(H)$ close to 1, we set $\mu_\gamma = 10$ and $\nu_\gamma = 0.1$. When the number of elements $\gamma_{l_i}(H) > \mu_\gamma$ or the number of elements $\gamma_{l_i}(H) < \nu_\gamma$ was less than $q/2$, we chose q_1 and q_2 to get the largest possible value of $q_1 + q_2$. The same comments hold for $\gamma_{l_i}(\tilde{H})$. For simplicity, in the following the notation $\gamma_{l_i}(H)$ is used also to indicate $\gamma_{l_i}(\tilde{H})$, as it will be clear from the context.

On the basis of numerical experiments, we decided to refresh the preconditioner, i.e., to build \mathcal{P}_{ex} instead of \mathcal{P}_{upd} , when the time for computing \mathcal{P}_{upd} and solving the linear system exceeded 90% of the time for building the last exact preconditioner and solving the corresponding system. When for a specific system of the sequence this situation occurred, the next system of the sequence was solved using the preconditioner \mathcal{P}_{ex} . We also set a maximum number, k_{max} , of consecutive preconditioner updates, after which the refresh was performed anyway. This strategy aims at avoiding possible situations in which the time saved by updating the Schur complement, instead of re-factorizing it, is offset by an excessive increase in the number of SQMR iterations, due to deterioration of the quality of the preconditioner. In the experiments discussed here $k_{max} = 5$ was used for all the test problems.

We observe that we did not apply any scaling to the matrix X in (2.7), although our supporting theory lies on the assumption that the eigenvalues of X are smaller than 2. Nevertheless, the results generally appear to be in agreement with the theory. It is also worth noting that since each KKT system is built from the approximation of the optimal solution computed at the previous IP iteration, different preconditioners produce changes in the sequence of KKT systems. On the other hand, testing our updating technique inside an IP solver allows to better evaluate its impact on the performance of the overall optimization method, providing a more general evaluation of our approach. In our experiments we did not observe strong differences in the behaviour of the IP method by using \mathcal{P}_{ex} or \mathcal{P}_{upd} with different values of q , and the number of IP iterations was generally unaffected by the choice of the preconditioner (a small variation of the IP iterations was observed in few cases, as shown in the tables reported in the next pages).

We performed the numerical experiments on an Intel Core 2 Duo E7300 processor with clock frequency of 2.66 GHz, 4 GB of RAM and 3 MB of cache memory, running Debian GNU/Linux 6.0.7 (kernel version 2.6.32-5-amd64). All the software

was compiled using the the GNU C and Fortran compilers (version 4.4.5).

Since the performance of the updating strategy was expected to depend on the cost of the factorization of the Schur complement, we selected test problems requiring different factorization costs. As a first test set, we considered some problems taken from the CUTEst collection [35] or obtained by modifying problems available in this collection, as explained next. Most of the large CUTEst convex QP problems with inequality constraints (corresponding to $\Theta^{(2)} \neq 0$) were not useful for our experiments, because of the extremely low cost of the factorization of their Schur complements. Therefore, we also modified two CUTEst QP problems with linear constraints $Ax = b$ and non-negligible factorization costs, by changing $Ax = b$ into $Ax \geq b$. The problems of this test set are listed in Table 4.1, along with their dimensions and the number of nonzero entries of their Schur complements. The modified problems are identified by appending “-M” to their original names. In the first three problems $\Theta^{(2)} = 0$, while in the remaining ones $\Theta^{(2)} \neq 0$. For all the problems the Schur complements are very sparse; furthermore, they are banded for STCQP2, MOSARQP1 and QPBAND (diagonal for the latter problem).

A comparison among the exact and updated preconditioners on this set of problems is presented in Table 4.1. For each preconditioner we report the total number of PRQP iterations (*IPits*), the total number of SQMR iterations (*its*) and the overall computation time, in seconds, needed to solve the KKT sequence. The results shows that using the updating strategy is not beneficial on these problems. Since the exact factorization of the Schur complement is not significantly more expensive than SQMR, the time saved by applying the updating strategy is not enough to offset the time required by the larger number of SQMR iterations resulting from the use of an approximate CP. Nevertheless, for STCQP2 and MOSARQP1, the updating strategy shows the same performance as the exact preconditioner. We also see that in many cases the number of SQMR iterations increases with q , which seems contrary to our expectations. This behaviour has been explained with a deeper analysis of the execution of PRQP on the selected test problems. The convergence histories of the IP method enlighten that the IP iterations at which the refresh takes place vary with q ; this does not allow a fair comparison among the updating rules using different values of the parameter q . This behaviour is ascribed to the fact that the computation of the exact preconditioner is not expensive and therefore the refresh strategy is very sensitive to the choice of q . In particular, in our experiments the choice $q = 0$ often leads to recomputing the exact preconditioners before the number of SQMR iterations increases too much, thus reducing the iteration count with respect to larger values of q . We also note that in some cases the number of SQMR iterations is practically constant as q varies, because either the number of elements $\gamma_{l_i}(H) \notin [\nu_\gamma, \mu_\gamma]$ is much smaller than q , or the values $\gamma_{l_i}(H)$ excluded by the updating strategy are not well separated from $\gamma_{q_2}(H)$ and $\gamma_{n-q_1+1}(H)$.

Despite these first unfavourable results, since the updating strategy does not excessively increase the number of SQMR iterations, we can still expect a significant time reduction on problems with Schur complements requiring large factorization times. In order to investigate this issue, we built a second set of test problems with less sparse Schur complements, by modifying the problems in Table 4.1 as follows. In problems CVXQP1, CVXQP1-M, CVXQP3 and CVXQP3-M, we added four nonzero entries per row in the matrix A , while in problem QPBAND we added two nonzero entries per row. In problem MOSARQP1, we introduced nonzeros in the positions (i, n) of the constraint matrix A , where i is such that $\text{mod}(i, 10) = 1$. These new

Problem	n m $nmz(S)$	\mathcal{P}_{ex}			$\mathcal{P}_{upd} (q=0)$			$\mathcal{P}_{upd} (q=50)$			$\mathcal{P}_{upd} (q=100)$		
		IPits	its	time	IPits	its	time	IPits	its	time	IPits	its	time
CVXQP1	20000	16	209	2.07e+0	16	298	2.35e+0	16	335	2.55e+0	16	323	2.49e+0
	10000												
	67976												
CVXQP3	20000	35	523	8.04e+0	35	800	9.32e+0	35	755	8.86e+0	35	757	8.90e+0
	15000												
	155942												
STCQP2	16385	12	226	1.46e+0	12	235	1.41e+0	12	235	1.43e+0	12	235	1.43e+0
	8190												
	114660												
CVXQP1-M	20000	26	1015	7.65e+0	29	1562	1.07e+1	26	1812	1.21e+1	26	1845	1.22e+1
	10000												
	67976												
CVXQP3-M	15000	30	1261	1.47e+1	30	1654	1.71e+1	30	2073	2.11e+1	30	2135	2.18e+1
	11250												
	155942												
MOSARQP1	22500	16	66	4.65e+0	16	193	4.53e+0	16	189	4.83e+0	16	215	5.21e+0
	20000												
	257166												
QPBAND	50000	12	757	7.13e+0	12	1596	1.37e+1	12	1600	1.43e+1	12	1599	1.37e+1
	25000												
	25000												

TABLE 4.1
Comparison between \mathcal{P}_{ex} and \mathcal{P}_{upd} on the first set of test problems.

problems are identified by appending “-D” to the names of the problems they come from, as listed in Table 4.2 (“D” stands for “denser”). Finally, starting from problems CVXQP3 and CVXQP3-M we generated two further problems, named CVXQP3-D2 and CVXQP3-M-D2, respectively. They were obtained by adding only one nonzero entry per row in the matrix A . The densities of the resulting Schur complements are between the Schur complements densities of the corresponding original and -D versions.

The results in Table 4.2 show that when the Schur complement is denser, the updating procedure provides a significant reduction in the overall computation time, because the increase in the SQMR iterations is largely offset by the time saving obtained by updating the factors of the Schur complement instead of recomputing them. For the problems under consideration the reduction ranges from 21%, for CVXQP3-M-D2 with $q = 0$, to 80%, for STCQP2-D with all the three values of q . Comparing the behaviour of the updating strategy on CVXQP3-M-D and CVXQP3-M-D2, we see that the percentage of time saved with the updating strategy drops from 53-54% to 21-24% when going from CVXQP3-M-D to CVXQP3-M-D2 (the latter has a sparser Schur complement). A similar behaviour can be observed by comparing the results obtained on CVXQP3-D and CVXQP3-D2. In this case the best time reduction for CVXQP3-D amounts to 68% ($q = 0$), while the best one for CVXQP3-D is 47% ($q = 50$). Furthermore, the time reduction also holds when the number of IP iterations corresponding to the updating strategy is greater than the number of IP iterations obtained with the exact preconditioner (see CVXQP3-D2). Surprisingly, also the reverse may happen, i.e., the updating procedure may slightly reduce the number of IP iterations (see CVXQP1-M-D and MOSARQP1-D).

We further note that the number of iterations obtained with \mathcal{P}_{upd} generally decreases as q increases; thus, for the second set of problems, updating the Schur complement by low-rank information appears to be beneficial in terms of iterations. There are also some cases where the number of SQMR iterations is practically constant as q varies (see STCQP2-D and QPBAND-D). In these cases, as for the corresponding problems in the first test set, we verified that either the number of elements $\gamma_{l_i}(H)$ with indices in Γ is very small or even zero, or those values of $\gamma_{l_i}(H)$ are not well separated from the remaining ones, thus making the low-rank modification ineffective. For similar reasons the reduction of the number of iterations from $q = 50$ to $q = 100$ is generally less significant than from $q = 0$ to $q = 50$. Finally, the best results in terms of execution time are mostly obtained with $q > 0$.

To provide more insight into the behaviour of the updated preconditioners, in Tables 4.3-4.6 we show some details concerning the solution of the sequences of KKT systems arising from four problems, i.e., CVXQP3, CVXQP3-D, MOSARQP1 and MOSARQP1-D. For each IP iteration we report the number, its , of SQMR iterations, as well as the time, T_{prec} , for building the preconditioner, the time, T_{solve} , for solving the linear system, and their sum, T_{sum} . The last row contains the total number of SQMR iterations and the total times, over all IP iterations, while the rows in bold correspond to the IP iterations at which the preconditioner is refreshed. These tables clearly support the previous observation that the updating strategy is efficient when the computation of \mathcal{P}_{ex} is expensive, as it is for CVXQP3-D and MOSARQP1-D. Conversely, when the time for building \mathcal{P}_{ex} is modest, recomputing \mathcal{P}_{ex} is a natural choice. It also appears that the refresh strategy plays a significant role in achieving efficiency, since it prevents the preconditioner from excessive deterioration. Finally, when the time for computing \mathcal{P}_{ex} is not dominant, the refresh tends to occur more

Problem	n m $nnz(S)$	\mathcal{P}_{ex}			$\mathcal{P}_{upd} (q = 0)$			$\mathcal{P}_{upd} (q=50)$			$\mathcal{P}_{upd} (q=100)$		
		IPits	its	time	IPits	its	time	IPits	its	time	IPits	its	time
CVXQP1-D	2000	15	239	2.95e+2	15	759	9.83e+1	15	616	9.91e+1	15	602	1.03e+2
	10000												
	240494												
CVXQP3-D	20000	15	192	1.03e+3	15	778	3.30e+2	15	526	4.40e+2	15	481	4.55e+2
	15000												
	542296												
CVXQP3-D2	20000	15	288	9.95e+1	18	1009	6.02e+1	17	819	5.26e+1	17	802	5.46e+1
	15000												
	224396												
STCQP2-D	16385	12	238	6.08e+2	12	262	1.22e+2	12	262	1.22e+2	12	262	1.22e+2
	8190												
	5003908												
CVXQP1-M-D	20000	28	1090	5.85e+2	28	4704	3.63e+2	28	3665	3.23e+2	27	3514	3.24e+2
	10000												
	240494												
CVXQP3-M-D	20000	25	910	1.93e+3	25	3605	9.08e+2	25	3416	8.89e+2	25	3317	9.07e+2
	15000												
	542296												
CVXQP3-M-D2	20000	25	822	1.66e+2	25	2782	1.32e+2	25	2645	1.33e+2	25	2148	1.25e+2
	15000												
	224396												
MOSARQP1-D	22500	24	93	4.94e+1	23	881	3.47e+1	22	599	3.00e+1	22	440	2.78e+1
	20000												
	573216												
QPBAND-D	50000	11	717	1.06e+3	11	2614	4.26e+2	11	2619	4.36e+2	11	2612	4.51e+2
	25000												
	149988												

TABLE 4.2
Comparison between \mathcal{P}_{ex} and \mathcal{P}_{upd} on the second set of test problems.

<i>IP it</i>	\mathcal{P}_{ex}				$\mathcal{P}_{upd} (q=50)$			
	<i>its</i>	T_{prec}	T_{solve}	T_{sum}	<i>its</i>	T_{prec}	T_{solve}	T_{sum}
1	23	8.82e-2	2.50e-1	3.38e-1	23	8.79e-2	2.49e-1	3.37e-1
2	8	7.86e-2	9.17e-2	1.70e-1	15	2.21e-2	1.70e-1	1.92e-1
3	6	7.06e-2	7.35e-2	1.44e-1	20	2.17e-2	2.17e-1	2.39e-1
4	5	7.86e-2	6.16e-2	1.40e-1	28	1.92e-2	2.93e-1	3.12e-1
5	5	7.86e-2	6.16e-2	1.40e-1	5	7.45e-2	5.85e-2	1.33e-1
6	5	7.46e-2	5.90e-2	1.34e-1	9	5.54e-3	9.98e-2	1.05e-1
7	7	7.86e-2	8.39e-2	1.62e-1	17	1.37e-2	1.84e-1	1.98e-1
8	7	7.46e-2	8.00e-2	1.55e-1	7	7.46e-2	7.96e-2	1.54e-1
9	9	7.46e-2	1.01e-1	1.75e-1	11	1.05e-2	1.20e-1	1.31e-1
10	9	7.86e-2	9.81e-2	1.77e-1	20	1.48e-2	2.01e-1	2.16e-1
11	11	7.86e-2	1.19e-1	1.98e-1	11	7.46e-2	1.19e-1	1.93e-1
12	12	7.46e-2	1.31e-1	2.05e-1	14	1.05e-2	1.42e-1	1.52e-1
13	12	7.86e-2	1.26e-1	2.04e-1	34	1.88e-2	3.34e-1	3.53e-1
14	12	7.86e-2	1.25e-1	2.04e-1	12	7.45e-2	1.24e-1	1.99e-1
15	12	7.06e-2	1.25e-1	1.96e-1	16	1.06e-2	1.59e-1	1.70e-1
16	12	7.86e-2	1.26e-1	2.04e-1	28	1.17e-2	2.72e-1	2.84e-1
17	14	7.86e-2	1.47e-1	2.25e-1	14	7.47e-2	1.45e-1	2.20e-1
18	14	7.46e-2	1.44e-1	2.19e-1	20	1.04e-2	1.99e-1	2.10e-1
19	14	7.86e-2	1.44e-1	2.23e-1	14	7.06e-2	1.43e-1	2.13e-1
20	14	7.46e-2	1.41e-1	2.16e-1	21	1.06e-2	2.01e-1	2.12e-1
21	16	7.86e-2	1.65e-1	2.43e-1	16	7.46e-2	1.63e-1	2.37e-1
22	14	7.86e-2	1.43e-1	2.22e-1	23	1.02e-2	2.31e-1	2.41e-1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
34	28	7.46e-2	2.68e-1	3.43e-1	28	7.46e-2	2.65e-1	3.40e-1
35	28	7.46e-2	2.71e-1	3.45e-1	42	1.14e-2	3.97e-1	4.08e-1
	523	2.69e+0	5.36e+0	8.04e+0	755	1.33e+0	7.53e+0	8.86e+0

TABLE 4.3
CVXQP3: details for \mathcal{P}_{ex} and \mathcal{P}_{upd} with $q = 50$.

frequently, since a small increase in the number of iterations obtained with \mathcal{P}_{upd} may easily raise the execution time over 90% of the time corresponding to the last application of the exact preconditioner.

5. Conclusion. We have proposed a preconditioner updating procedure for the solution of sequences of KKT systems arising in IP methods for convex QP problems. The preconditioners built by this procedure belong to the class of inexact CPs and are obtained by updating a given seed CP. The updates are performed through low-rank corrections of the Schur complement of the (1,1) block in the seed preconditioner and generate factorized preconditioners. The rule for identifying the low-rank corrections is based on new bounds on the eigenvalues of the preconditioned matrix. The numerical experiments show that our updated preconditioners, combined with a suitable preconditioner refreshing, can be rather successful. More precisely, the higher the cost of the Schur complement factorization, the more advantageous the updating procedure becomes. Finally, we believe that the updating strategy proposed here paves the way to the definition of preconditioner updating procedures for sequences of KKT systems where the Hessian and constraint matrices change from one iteration to the next.

<i>IP it</i>	\mathcal{P}_{ex}				$\mathcal{P}_{upd} (q=50)$			
	<i>its</i>	T_{fact}	T_{solve}	T_{sum}	<i>its</i>	T_{prec}	T_{solve}	T_{sum}
1	30	5.18e+0	1.19e+0	6.37e+0	30	5.24e+0	1.18e+0	6.42e+0
2	12	5.16e+0	4.87e-1	5.65e+0	14	5.52e-1	5.54e-1	1.11e+0
3	8	5.16e+0	3.40e-1	5.50e+0	15	5.49e-1	6.01e-1	1.15e+0
4	5	5.12e+0	2.24e-1	5.35e+0	13	5.11e-1	5.29e-1	1.04e+0
5	5	5.14e+0	2.27e-1	5.37e+0	36	5.52e-1	1.37e+0	1.92e+0
6	8	5.16e+0	3.37e-1	5.50e+0	48	6.00e-1	1.79e+0	2.39e+0
7	10	5.15e+0	4.15e-1	5.57e+0	10	5.24e+0	4.17e-1	5.66e+0
8	12	5.16e+0	4.93e-1	5.65e+0	15	1.56e-1	5.89e-1	7.45e-1
9	14	5.13e+0	5.61e-1	5.70e+0	22	2.76e-1	8.39e-1	1.11e+0
10	14	5.18e+0	5.58e-1	5.74e+0	41	4.82e-1	1.54e+0	2.02e+0
11	16	5.14e+0	6.33e-1	5.78e+0	78	4.90e-1	2.90e+0	3.39e+0
12	17	5.17e+0	6.68e-1	5.83e+0	139	4.66e-1	5.09e+0	5.56e+0
13	19	5.14e+0	7.40e-1	5.88e+0	19	5.25e+0	7.44e-1	5.99e+0
14	21	5.15e+0	8.11e-1	5.97e+0	31	1.95e-1	1.16e+0	1.36e+0
15	24	5.15e+0	9.24e-1	6.08e+0	62	4.68e-1	2.32e+0	2.79e+0
16	26	5.13e+0	1.39e+0	6.51e+0	86	4.72e-1	3.17e+0	3.64e+0
17	47	5.27e+0	1.76e+0	7.03e+0	160	4.61e-1	5.83e+0	6.29e+0
	288	8.77e+1	1.18e+1	9.95e+1	819	2.20e+1	3.06e+1	5.26e+1

TABLE 4.4
CVXQP3-D: details for \mathcal{P}_{ex} and \mathcal{P}_{upd} with $q = 50$.

<i>IP it</i>	\mathcal{P}_{ex}				$\mathcal{P}_{upd} (q=100)$			
	<i>its</i>	T_{prec}	T_{solve}	T_{sum}	<i>its</i>	T_{prec}	T_{solve}	T_{sum}
1	1	2.11e-1	2.95e-2	2.41e-1	1	2.09e-1	2.95e-2	2.39e-1
2	2	2.09e-1	4.53e-2	2.54e-1	5	3.48e-2	9.29e-2	1.28e-1
3	2	2.09e-1	4.84e-2	2.58e-1	9	6.82e-2	1.54e-1	2.22e-1
4	3	2.09e-1	6.55e-2	2.75e-1	3	2.05e-1	6.15e-2	2.66e-1
5	3	2.13e-1	6.29e-2	2.76e-1	18	9.65e-3	2.77e-1	2.86e-1
6	4	2.09e-1	8.21e-2	2.91e-1	4	2.01e-1	7.30e-2	2.74e-1
7	4	2.13e-1	7.84e-2	2.92e-1	16	4.92e-2	2.56e-1	3.06e-1
8	4	2.09e-1	7.76e-2	2.87e-1	4	2.05e-1	7.30e-2	2.78e-1
9	4	2.13e-1	7.51e-2	2.88e-1	19	7.03e-2	2.92e-1	3.63e-1
10	5	2.09e-1	9.20e-2	3.01e-1	5	2.01e-1	8.67e-2	2.87e-1
11	5	2.17e-1	8.93e-2	3.06e-1	15	2.31e-2	2.35e-1	2.58e-1
12	6	2.13e-1	1.09e-1	3.23e-1	62	6.90e-2	9.12e-1	9.81e-1
13	5	2.13e-1	8.95e-2	3.03e-1	5	2.09e-1	8.42e-2	2.93e-1
14	6	2.13e-1	1.07e-1	3.20e-1	21	4.69e-2	3.22e-1	3.69e-1
15	6	2.13e-1	1.07e-1	3.20e-1	6	2.01e-1	1.01e-1	3.02e-1
16	6	2.09e-1	1.04e-1	3.13e-1	22	2.73e-2	3.35e-1	3.63e-1
	66	3.38e+0	1.27e+0	4.65e+0	215	1.83e+0	3.39e+0	5.21e+0

TABLE 4.5
MOSARQP1: details for \mathcal{P}_{ex} and \mathcal{P}_{upd} with $q = 100$.

Acknowledgments. We are indebted to Tim Davis for his valuable help in interfacing PRQP with CHOLMOD. We also express our thanks to Miroslav Tuma for making available a Fortran 90 implementation of SQMR used in preliminary numerical experiments. Finally, we wish to thank the anonymous referees for their useful comments, which helped us to improve the quality of this work.

<i>IP it</i>	\mathcal{P}_{ex}				$\mathcal{P}_{upd} (q=100)$			
	<i>its</i>	T_{prec}	T_{solve}	T_{sum}	<i>its</i>	T_{prec}	T_{solve}	T_{sum}
1	1	1.95e+0	5.99e-2	2.01e+0	1	1.90e+0	5.93e-2	1.96e+0
2	1	1.92e+0	5.92e-2	1.98e+0	4	4.78e-2	1.34e-1	1.82e-1
3	1	1.91e+0	5.85e-2	1.97e+0	8	8.98e-2	2.46e-1	3.36e-1
4	2	1.93e+0	9.02e-2	2.02e+0	15	7.31e-1	4.36e-1	1.17e+0
5	2	1.91e+0	9.01e-2	2.00e+0	22	7.31e-1	6.31e-1	1.36e+0
6	3	1.91e+0	1.18e-1	2.03e+0	57	5.05e-1	1.56e+0	2.07e+0
7	3	1.91e+0	1.19e-1	2.03e+0	3	1.89e+0	1.13e-1	2.00e+0
8	4	1.90e+0	1.51e-1	2.05e+0	15	5.10e-2	4.39e-1	4.90e-1
9	4	1.90e+0	1.48e-1	2.05e+0	20	5.09e-2	5.74e-1	6.25e-1
10	4	1.91e+0	1.47e-1	2.06e+0	28	3.82e-1	7.92e-1	1.17e+0
11	4	1.90e+0	1.50e-1	2.05e+0	56	7.17e-1	1.55e+0	2.27e+0
12	4	1.91e+0	1.46e-1	2.05e+0	4	1.89e+0	1.42e-1	2.03e+0
13	4	1.91e+0	1.47e-1	2.06e+0	5	9.13e-2	1.64e-1	2.55e-1
14	4	1.93e+0	1.48e-1	2.08e+0	9	1.03e-1	2.69e-1	3.72e-1
15	3	1.91e+0	1.19e-1	2.03e+0	17	1.03e-1	4.86e-1	5.89e-1
16	4	1.92e+0	1.46e-1	2.06e+0	48	7.10e-1	1.32e+0	2.03e+0
17	4	1.94e+0	1.45e-1	2.08e+0	6	1.88e+0	1.94e-1	2.07e+0
18	6	1.92e+0	2.03e-1	2.13e+0	16	5.13e-2	4.57e-1	5.08e-1
19	6	1.91e+0	2.07e-1	2.11e+0	47	7.16e-1	1.30e+0	2.02e+0
20	5	1.91e+0	1.75e-1	2.08e+0	6	1.89e+0	1.97e-1	2.08e+0
21	6	1.91e+0	2.06e-1	2.11e+0	20	2.21e-1	5.69e-1	7.90e-1
22	6	1.91e+0	2.03e-1	2.12e+0	33	4.98e-1	9.19e-1	1.42e+0
23	6	1.93e+0	2.03e-1	2.13e+0				
24	6	1.91e+0	2.02e-1	2.12e+0				
	93	4.60e+1	3.44e+0	4.94e+1	440	1.52e+1	1.26e+1	2.78e+1

TABLE 4.6
MOSARQP1-D: details for \mathcal{P}_{ex} and \mathcal{P}_{upd} with $q = 100$.

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, D. Bertaccini, B. Morini, *Nonsymmetric preconditioner updates in Newton-Krylov methods for nonlinear systems*, SIAM Journal on Scientific Computing, 33 (2011), pp. 2595–2619.
- [4] 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.
- [5] 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.
- [6] S. Bellavia, B. Morini, M. Porcelli, *New updates of incomplete LU factorizations and applications to large nonlinear systems*, Optimization Methods and Software, 29 (2014), pp. 321–340.
- [7] M. Benzi, D. Bertaccini, *Approximate inverse preconditioning for shifted linear systems*, BIT, 43 (2003), pp. 231–244.
- [8] M. Benzi, V. Simoncini, *On the eigenvalues of a class of saddle point matrices*, Numerische Mathematik, 103 (2006), pp. 173–196.
- [9] L. Bergamaschi, *Eigenvalue distribution of constraint-preconditioned symmetric saddle point matrices*, Numerical Linear Algebra with Applications, 19 (2012), pp. 754–772.
- [10] L. Bergamaschi, R. Bru, A. Martinez, M. Putti, *Quasi-Newton preconditioners for the inexact Newton method*, Electronic Transactions on Numerical Analysis, 23 (2006) pp. 76–87.
- [11] L. Bergamaschi, J. Gondzio, G. Zilli, *Preconditioning Indefinite Systems in Interior Point Methods for Optimization*, Computational Optimization and Applications, 28 (2004), pp. 149–171.
- [12] 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.
- [13] 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.
- [14] S. Cafieri, M. D’Apuzzo, V. De Simone, 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.
- [15] S. Cafieri, M. D’Apuzzo, V. De Simone, D. di Serafino, *On the Use of an Approximate Constraint Preconditioner in a Potential Reduction Algorithm for Quadratic Programming*, in “Applied and Industrial Mathematics in Italy II”, V. Cutello, G. Fotia and L. Puccio eds., Series on Advances in Mathematics for Applied Sciences, 75, World Scientific, 2007, pp. 220–230.
- [16] 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.
- [17] C. Calgareo, J.P. Chehab, Y. Saad, *Incremental incomplete ILU factorizations with applications*, Numerical Linear Algebra with Applications, 17 (2010), pp. 811–837.
- [18] B. Carpentieri, I.S. Duff, L. Giraud, *A class of spectral two-level preconditioners*, SIAM Journal on Scientific Computing, 25 (2003), pp. 749–765.
- [19] 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.
- [20] 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.
- [21] T.A. Davis, W.W. Hager, *Dynamic supernodes in sparse Cholesky update/downdate and triangular solves*, ACM Transactions on Mathematical Software, 35 (2009), article 27.
- [22] H.S. Dollar, *Constraint-style preconditioners for regularized saddle-point problems*, SIAM Journal on Matrix Analysis and Applications, 29 (2007), pp. 672–684.
- [23] H.S. Dollar, N.I.H. Gould, W.H.A. Schilders, A.J. Wathen, *Implicit-factorization preconditioning and iterative solvers for regularized saddle-point systems*, SIAM Journal on Matrix

- Analysis and Applications, 28 (2006), pp. 170–189.
- [24] H.S. Dollar, A.J. Wathen, *Approximate factorization constraint preconditioners for saddle-point matrices* SIAM Journal on Scientific Computing, 27, (2006), pp. 1555–1572.
 - [25] J. Duintjer Tebbens, M. Tüma, *Efficient Preconditioning of Sequences of Nonsymmetric Linear Systems*, SIAM Journal on Scientific Computing, 29 (2007), pp. 1918–1941.
 - [26] 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.
 - [27] R.W. Freund, N.M. Nachtigal, *QMR: a quasi-minimal residual method for non-Hermitian linear systems*, Numerische Mathematik, 60 (1991), pp. 315–339.
 - [28] R.W. Freund, N.W. Nachtigal, *Software for simplified Lanczos and QMR algorithms*, Applied Numerical Mathematics, 19 (1995), pp. 319–341.
 - [29] R.W. Freund, N.W. Nachtigal, *QMRPACK: a package of QMR algorithms*, ACM Transactions on Mathematical Software, 22 (1996), pp. 46–77.
 - [30] G. Fasano, M. Roma, *Preconditioning Newton–Krylov Methods in Non-Convex Large Scale Optimization*, Computational Optimization and Applications, 56 (2013), pp. 253–290.
 - [31] 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.
 - [32] L. Giraud, S. Gratton, E. Martin, *Incremental spectral preconditioners for sequences of linear systems*, Applied Numerical Mathematics, 57 (2007), pp. 1164–1180.
 - [33] G.H. Golub, C.F. van Loan, *Matrix Computations*, third edition, The Johns Hopkins University Press, 1996.
 - [34] J. Gondzio, *Interior point methods 25 years later*, European Journal of Operational Research, 218 (2012), pp. 587–601.
 - [35] 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.
 - [36] S. Gratton, A. Sartenaer, J. Tshimanga, *On a class of limited memory preconditioners for large scale linear systems with multiple right-hand sides*, SIAM Journal on Optimization, 21 (2011), pp. 912–935.
 - [37] 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.
 - [38] D. Loghin, D. Ruiz, A. Touhami, *Adaptive preconditioners for nonlinear systems of equations*, Journal of Computational and Applied Mathematics, 189 (2006), pp. 362–374.
 - [39] 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.
 - [40] G. Meurant, *On the incomplete Cholesky decomposition of a class of perturbed matrices*, SIAM J. Sci. Comput., 23 (2001), pp. 419–429.
 - [41] J.L. Morales, J. Nocedal, *Automatic preconditioning by limited memory quasi-Newton updating*, SIAM Journal on Optimization, 10 (2000), pp. 1079–1096.
 - [42] 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.
 - [43] Y. Saad, M.H. Schultz, *GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.
 - [44] 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.
 - [45] 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.
 - [46] S.J. Wright, *Primal-Dual Interior-Point Methods*, SIAM, Philadelphia, 1997.