MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Preconditioning of conjugate gradient iterations in interior point MPC method

Malyshev, A.; Quirynen, R.; Knyazev, A.

TR2018-120 August 25, 2018

Abstract

There are several efficient direct solvers for structured systems of linear equations defining search directions in primal-dual interior point methods applied to constrained model predictive control problems. We propose reusing matrix decompositions of direct solvers as preconditioners in Krylov-subspace methods applied to subsequent iterations of the interior point method, which results in at least halving its asymptotic computational complexity. We also analyze sensitivity of direct solvers to the regularization parameters.

IFAC Conference on Nonlinear Model Predictive Control

Preconditioning of conjugate gradient iterations in interior point MPC method

Alexander Malyshev * **Rien Quirynen** ** **Andrew Knyazev** **

* *University of Bergen, Department of Mathematics, Postbox 8700, 5020 Bergen, Norway. (e-mail: alexander.malyshev@uib.no)*
** *Mitsubishi Electric Research Laboratories, 201 Broadway, 8th floor, Cambridge, MA 02139, USA. (e-mails: quirynen@merl.com, knyazev@merl.com)*

Abstract: There are several efficient direct solvers for structured systems of linear equations defining search directions in primal-dual interior point methods applied to constrained model predictive control problems. We propose reusing matrix decompositions of direct solvers as preconditioners in Krylov-subspace methods applied to subsequent iterations of the interior point method, which results in at least halving its asymptotic computational complexity. We also analyze sensitivity of direct solvers to the regularization parameters.

Keywords: model predictive control, interior point, preconditioning, conjugate gradient.

1. INTRODUCTION

In this paper, preconditioning techniques are proposed that may be useful for numerical solution of optimal control problems in model predictive control (MPC); see (Rawlings et al., 2017). We consider MPC over a receding horizon in the form of a quadratic program (QP), where the cost function is given by a convex quadratic performance index, the equality constraints are the linear model dynamics, the inequality constraints on the state and control inputs are also linear; similar formulations are used in (Wang and Boyd, 2010; Shahzad et al., 2012; Frison et al., 2014; Quirynen et al., 2018). Such structured QPs appear as subproblems within a sequential QP for nonlinear MPC; see, e.g., (Diehl et al., 2009).

The inequality constraints in the QP problem are usually treated by the interior point (IP) method or by the active-set method; cf. (Bartlett et al., 2000). The IP method is characterized by fast convergence, which is almost independent of the problem size; see, e.g., (Mehrotra, 1992; Gondzio, 2012, 2013). The active-set solvers for MPC are highly competitive owing to rather small changes of the active set during continuation steps, reducing arithmetic costs by exploiting low-rank updates of a Cholesky factorization when changing the active set as shown in (Gill et al., 1974; Kirches et al., 2011).

Matrices in systems of linear equations defining search directions in the Newton method are sparse, consisting of blocks with banded matrices. Several efficient direct methods have been proposed for solving such systems; see, e.g., (Rao et al., 1998; Wang and Boyd, 2010; Domahidi et al., 2012; Frison and Jørgensen, 2013). Let n_x, n_u, n_g, and N denote the dimensions of the state variable and the control input, the number of the inequality constraints, and the horizon length, respectively. Arithmetic complexity of block structured direct solvers used in the IP method typically amounts to $O(N(n_x^3 + n_u^3 + n_g^3))$ flops.

Numerical solution of linear systems of the Newton method, e.g., using direct solvers, is the most time-consuming online part of MPC. Computational complexity can be reduced if saddle-point linear systems of KKT conditions are replaced by linear systems having symmetric positive definite (SPD) block structured matrices. Direct solvers for SPD linear systems are commonly based on the Cholesky factorization, often implemented to take advantage of the BLAS3 matrix operations, which is highly efficient on contemporary computer architectures with advanced memory hierarchies. The active-set methods may achieve theoretical complexity $O(N(n_x^2+n_u^2+n_g^2))$, if based on cost-efficient low-rank updates, e.g., of the Cholesky factorization, although the low-rank updates can only use BLAS2; see (Anderson et al., 1999).

Well-known fast convergence of the IP method, its resilience to inexact implementations, and high performance of BLAS3 implementations of the direct solvers can make the IP method competitive with the active-set methods. We have considered the IP method in (Malyshev et al., 2018) using each of the three direct solvers: Riccati recursion, Schur complement, and augmented Lagrangian, with regularization by adding the identity matrix multiplied by a small parameter to some blocks. The factorized Riccati recursion and the augmented Lagrangian regularization have been found the best in terms of arithmetic operations.

The sparse linear systems of the Newton methods can be solved by iterative methods; see, e.g., (Golub and Van Loan, 2013, Ch. 11). If an efficient preconditioner is available, the theoretical complexity of an iterative method is reduced to $O(N(n_x^2 + n_u^2 + n_g^2))$, and the preconditioned iterative method outperforms direct solvers for relatively large values of n_x, n_u, n_g.

The rest of this paper is organized as follows. Section 2 introduces the notation and the MPC QP problem. Section 3 discusses reductions to SPD systems. We revisit reg-

ularization of the Schur complement in Section 4 and the augmented Lagrangian in Section 5. We show theoretically in Section 6 and numerically in Section 7 that the former is less sensitive to the choice of the regularization parameter, compared to the latter. Section 7 also illustrates behavior of the proposed preconditioning in the conjugate gradient (CG) method, where Cholesky factorizations computed at odd IP iterations are reused for preconditioning at even IP iterations to halve the computational cost.

Our main contributions include (a) reuse of the Cholesky factorization for preconditioning subsequent IP iterations and (b) analysis of sensitivity of the regularized direct solvers from Sections 4 and 5 to the choice of regularization parameter values ϵ and τ.

2. QUADRATIC PROGRAMMING FOR MODEL PREDICTIVE CONTROL

We consider the following quadratic program (QP) with respect to a decision variable $d \in \mathbb{R}^{n_h}$,

$$\min_d \frac{1}{2} d^T H d + h^T d \quad \text{subject to } Fd = f,\ Gd \le g, \qquad (1)$$

where $n_h = (n_x + n_u)N + n_x$, $n_f = n_x(N+1)$, and the matrices $H \in \mathbb{R}^{n_h \times n_h}$, $F \in \mathbb{R}^{n_f \times n_h}$, $G \in \mathbb{R}^{n_g \times n_h}$ and the vectors h, f, and g have the form

$$H = \begin{bmatrix} Q_0 & S_0^T & \cdots & 0 & 0 & 0 \\ S_0 & R_0 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & Q_{N-1} & S_{N-1}^T & 0 \\ 0 & 0 & \cdots & S_{N-1} & R_{N-1} & 0 \\ 0 & 0 & \cdots & 0 & 0 & Q_N \end{bmatrix},$$

$$F = \begin{bmatrix} -I & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ A_0 & B_0 & -I & 0 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & A_{N-1} & B_{N-1} & -I \end{bmatrix},$$

$$G = \begin{bmatrix} G_0^x & G_0^u & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & G_{N-1}^x & G_{N-1}^u & 0 \\ 0 & 0 & \cdots & 0 & 0 & G_N^x \end{bmatrix},$$

$$h = \begin{bmatrix} h_0^x \\ h_0^u \\ \vdots \\ h_{N-1}^u \\ h_N^x \end{bmatrix}, \quad f = \begin{bmatrix} f_0 \\ f_1 \\ \vdots \\ f_N \end{bmatrix}, \quad g = \begin{bmatrix} g_0 \\ g_1 \\ \vdots \\ g_N \end{bmatrix}.$$

The blocks I in F denote identity matrices of order n_x. We assume that the matrix H is symmetric positive semidefinite. We note that the matrix H as well as the blocks G_i^x and G_i^u of the matrix G are often diagonal.

Optimization problem (1) is solved by the primal-dual IP method; see, e.g., (Wright, 1997; Shahzad et al., 2012). The IP method uses the following Karush-Kuhn-Tucker (KKT) optimality conditions

$$Hd + F^T y + G^T z + h = 0, \qquad (2)$$
$$Fd - f = 0, \qquad (3)$$
$$Gd - g + s = 0, \qquad (4)$$
$$Z\mathbf{1} - \mu S^{-1}\mathbf{1} = 0, \qquad (5)$$

where $\mathbf{1} = [1 \dots 1]^T \in \mathbb{R}^{n_g}$, and the matrices $Z = \operatorname{diag}(z)$ and $S = \operatorname{diag}(s)$ are diagonal. The vectors $y \in \mathbb{R}^{n_f}$ and $z \in \mathbb{R}^{n_g}$ denote the Lagrange multipliers. The slack vector $s \in \mathbb{R}^{n_g}$ with positive components and the small parameter $\mu > 0$ are required by the IP method.

The KKT conditions consist of linear equations (2)-(4) and nonlinear equation (5). The latter is usually substituted by $ZS\mathbf{1} - \mu\mathbf{1} = 0$ to improve scaling the residuals in the Newton method. The search direction at iteration k of the Newton method for solving (2)-(5) is determined by the system of linear equations

$$\begin{bmatrix} H & F^T & G^T & 0 \\ F & 0 & 0 & 0 \\ G & 0 & 0 & I \\ 0 & 0 & S^k & Z^k \end{bmatrix} \begin{bmatrix} \Delta d^k \\ \Delta y^k \\ \Delta z^k \\ \Delta s^k \end{bmatrix} = - \begin{bmatrix} r_H^k \\ r_F^k \\ r_G^k \\ r_S^k \end{bmatrix}, \qquad (6)$$

with the residuals

$$r_H^k = Hd^k + F^T y^k + G^T z^k + h,$$
$$r_F^k = Fd^k - f,$$
$$r_G^k = Gd^k - g + s^k,$$
$$r_S^k = Z^k S^k \mathbf{1} - \sigma\mu^k \mathbf{1}, \text{ where } \sigma \in (0,1).$$

The scalar σ in the last formula is called a *centering parameter*. The value of $\mu^k = (z^k)^T s^k / n_g$ is related to a *duality gap*; e.g., (Wright, 1997; Gondzio, 2012).

Linear system (6) must be solved online repeatedly for $k = 1, 2, \dots, k_{\max}$, where $k_{\max}$ is determined by a stopping criterion, e.g., based on the 2-norm of the residual vector in the right-hand side of (6). The matrix of linear system (6) becomes increasingly ill-conditioned when some inequality constraints become active. Nevertheless, the Newton IP method performs well, if system (6) is solved by direct methods, as discussed and justified in (Wright, 1998). We note that the matrix of system (6) depends on the diagonal matrices S^k and Z^k in the last block row. Therefore, a direct solver has to factorize the matrix of (6) at each iteration k, which is the most time-consuming operation of the IP method. Reducing costs of solving (6) is crucial.

3. PRELIMINARY REDUCTIONS BY BLOCK GAUSSIAN ELIMINATIONS

The slack variable Δs^k is usually eliminated from (6), which leads to the system of linear equations

$$\begin{bmatrix} H & F^T & G^T \\ F & 0 & 0 \\ G & 0 & -W^k \end{bmatrix} \begin{bmatrix} \Delta d^k \\ \Delta y^k \\ \Delta z^k \end{bmatrix} = - \begin{bmatrix} r_H^k \\ r_F^k \\ r_W^k \end{bmatrix} \qquad (7)$$

with the diagonal matrix $W^k = (Z^k)^{-1} S^k$. The residual r_G^k is substituted with the residual $r_W^k = r_G^k - (Z^k)^{-1} r_S^k$. The eliminated variable Δs^k is recovered by the formula

$$\Delta s^k = -(Z^k)^{-1} \left(r_S^k + S^k \Delta z^k \right).$$

In (Shahzad et al., 2010a,b, 2012), the block $\begin{bmatrix} H & F^T \\ F & 0 \end{bmatrix}$ is assumed to be nonsingular and (7) is reduced to

$$\left\{ W^k + [G\ 0] \begin{bmatrix} H & F^T \\ F & 0 \end{bmatrix}^{-1} \begin{bmatrix} G^T \\ 0 \end{bmatrix} \right\} \Delta z^k =$$

$$r_M^k = r_W^k - [G\ 0] \begin{bmatrix} H & F^T \\ F & 0 \end{bmatrix}^{-1} \begin{bmatrix} r_H^k \\ r_F^k \end{bmatrix}.$$

The main advantage of such a reduction is that the matrix $M = [G\ 0] \begin{bmatrix} H & F^T \\ F & 0 \end{bmatrix}^{-1} \begin{bmatrix} G^T \\ 0 \end{bmatrix}$ is computed only once for all IP iterations. The SPD system $(W^k + M)\Delta z^k = r_M^k$ is then solved by preconditioned CG. However, cheap preconditioning via the diagonal part of $W^k + M$ may be not sufficient to provide fast convergence, while preconditioning using the block diagonal or tridiagonal parts of $W^k + M$ require expensive computations of the Cholesky factorization at every iteration.

In contrast, we eliminate the variable Δz^k from linear system (7). The resulting 2×2 block linear system

$$\begin{bmatrix} \tilde{H}^k & F^T \\ F & 0 \end{bmatrix} \begin{bmatrix} \Delta d^k \\ \Delta y^k \end{bmatrix} = - \begin{bmatrix} r_E^k \\ r_F^k \end{bmatrix} \tag{8}$$

has the residual $r_E^k = r_H^k + G^T (W^k)^{-1} r_W^k$ and the block

$$\tilde{H}^k = H + G^T (W^k)^{-1} G, \tag{9}$$

which is symmetric positive semidefinite and block diagonal with the same sparsity pattern as the matrix H.

The eliminated variable Δz^k is recovered by the formula

$$\Delta z^k = (W^k)^{-1} \left(G \Delta d^k + r_W^k \right). \tag{10}$$

Systems in the form (8) are called *saddle point problems*.

Now that we have reduced (6) to (8), further reductions of the linear system (8) with respect to the variables Δd^k or Δy^k are possible. In the next two sections, we describe two popular alternative reductions, Schur complement- and augmented Lagrangian-based, both supplemented with regularization by adding the identity matrix multiplied by a small parameter to some blocks.

4. INEXACT NEWTON METHOD BY THE SCHUR COMPLEMENT APPROACH

The elimination of Δd^k is most popular in the primal-dual IP method; see (Wright, 1997). If the matrix $\tilde{H}^k$ in (9) is singular, it is usually replaced by a regularized SPD block $\hat{H}^k = \tilde{H}^k + \epsilon I$, where ϵ is a small positive constant. A justification of such a replacement is due to the fact that the IP algorithm can use an inexact Newton method, i.e., the system (6) may be solved only approximately; see, e.g., (Nocedal and Wright, 2006).

In the following regularized analog of system (8),

$$\begin{bmatrix} \hat{H}^k & F^T \\ F & 0 \end{bmatrix} \begin{bmatrix} \Delta d^k \\ \Delta y^k \end{bmatrix} = - \begin{bmatrix} r_E^k \\ r_F^k \end{bmatrix}, \tag{11}$$

the first equation of (11) is equivalent to

$$\Delta d^k + (\hat{H}^k)^{-1} F^T \Delta y^k = -(\hat{H}^k)^{-1} r_E^k$$

and hence

$$\left[F (\hat{H}^k)^{-1} F^T \right] \Delta y^k = r_F^k - F (\hat{H}^k)^{-1} r_E^k. \tag{12}$$

Equation (12), often called a *normal equation*, has the SPD matrix $F(\hat{H}^k)^{-1}F^T$, which is the *Schur complement* for linear system (11). SPD linear system (12) can be solved by preconditioned CG.

For nonlinear MPC, preconditioning is proposed, e.g., in (Knyazev et al., 2015; Knyazev and Malyshev, 2015, 2016) for the Newton method. Preconditioning in the IP context is widely discussed; see, e.g., the survey in (Gondzio, 2012). Interesting relations between preconditioners for normal (12) and indefinite (8) equations are discovered in (Oliveira and Sorensen, 2005). We propose using the Cholesky factorization of the matrix $F(\hat{H}^{k-1})^{-1}F^T$, i.e. based on the values from the previous IP iteration $k-1$, as a preconditioner for CG in iteration k.

The arithmetic complexity of the Cholesky factorization of the matrix $F(\hat{H}^{k-1})^{-1}F^T$ is

$$N \left(\frac{19}{3} n_x^3 + 4 n_x^2 n_u + 2 n_x n_u^2 + \frac{1}{3} n_u^3 \right),$$

where only the leading terms are taken into account.

5. INEXACT NEWTON METHOD BY AUGMENTED LAGRANGIAN REGULARIZATION

In this section, we investigate elimination of the variable Δy^k from system (8). Since the IP iterations allow the inexact Newton method, we consider the following regularized analog of system (8),

$$\begin{bmatrix} \tilde{H}^k & F^T \\ F & -\tau I \end{bmatrix} \begin{bmatrix} \Delta d^k \\ \Delta y^k \end{bmatrix} = - \begin{bmatrix} r_E^k \\ r_F^k \end{bmatrix}, \tag{13}$$

where the positive regularization parameter τ is sufficiently small. Such a regularization is often used in iterative methods for numerical solution of saddle point problems, especially for preconditioning, e.g., (Rees and Greif, 2007; Benzi and Wathen, 2008), and is also connected to the augmented Lagrangian method in numerical optimization; see, e.g., (Nocedal and Wright, 2006).

Linear system (13) can be equivalently transformed to the block triangular system

$$\begin{bmatrix} \tilde{H}^k + \frac{1}{\tau} F^T F & 0 \\ F & -\tau I \end{bmatrix} \begin{bmatrix} \Delta d^k \\ \Delta y^k \end{bmatrix} = - \begin{bmatrix} r_E^k + \frac{1}{\tau} F^T r_F^k \\ r_F^k \end{bmatrix}. \tag{14}$$

The solution of (14) is given by the formulas

$$\Delta d^k = -(\tilde{H}^k + \frac{1}{\tau} F^T F)^{-1} (r_E^k + \frac{1}{\tau} F^T r_F^k), \tag{15}$$

$$\Delta y^k = \frac{1}{\tau} (r_F^k + F \Delta d^k). \tag{16}$$

The main computational burden in formulas (15)-(16) is solving the system with the SPD matrix $\tilde{H}^k + \frac{1}{\tau} F^T F$ in (15), which depends on k. Arithmetic complexity of the Cholesky factorization for the matrix $\tilde{H}^k + \frac{1}{\tau} F^T F$ equals

$$N \left(\frac{7}{3} n_x^3 + 4 n_x^2 n_u + 2 n_x n_u^2 + \frac{1}{3} n_u^3 \right).$$

By analogy with the previous section, we propose evaluating (15) by preconditioned CG, where the preconditioner is based on the Cholesky factorization of the system matrix $\tilde{H}^{k-1} + \frac{1}{\tau} F^T F$ from the previous IP iteration.

6. ANALYSIS OF REGULARIZATION IN INEXACT NEWTON METHODS

Regularization of the saddle point matrices $\begin{bmatrix} \tilde{H} & F^T \\ F & 0 \end{bmatrix}$ is a well-known topic in optimization theory. For example, (Vanderbei, 1995) investigates LDL^T factorizations of the following matrices, with both diagonal blocks regularized,

$$\begin{bmatrix} \tilde{H} + \epsilon I & F^T \\ F & -\tau I \end{bmatrix}. \tag{17}$$

(Saunders and Tomlin, 1996) consider the IP iterations, where the inexact Newton method uses regularized linear systems with the matrices of form (17). They give a formulation of the initial QP problem, which is equivalent to the regularized IP iteration, and present numerous tests showing high efficiency of the regularization approach.

In this section, we contribute into theoretical analysis of the regularizations given by equations (11) and (13).

Since the matrix $\tilde{H}^k$ is symmetric positive semidefinite, it has its symmetric positive semidefinite square root $\tilde{H}^{1/2}$. The convex saddle point problem can be analyzed by the aid of the generalized singular value decomposition (GSVD) of the matrix pair $\tilde{H}^{1/2}$ and F; see, e.g., (Golub and Van Loan, 2013). In our case, there exist orthogonal matrices U_1 and U_2 and a nonsingular matrix V such that

$$\tilde{H}^{1/2} = U_1 \begin{bmatrix} C & \\ & I \end{bmatrix} V^T, \quad F = U_2 \begin{bmatrix} S & 0 \end{bmatrix} V^T,$$

where the diagonal matrices

$$C = \begin{pmatrix} c_1 & & & \\ & c_2 & & \\ & & \ddots & \\ & & & c_{n_f} \end{pmatrix}, \quad S = \begin{pmatrix} s_1 & & & \\ & s_2 & & \\ & & \ddots & \\ & & & s_{n_f} \end{pmatrix}$$

have nonnegative diagonal elements c_i and s_i such that $c_i^2 + s_i^2 = 1$ for $i = 1, \ldots, n_f$. We consider below the regularization $\tilde{H}^k + \epsilon V V^T$ instead of our actual regularization $\hat{H}^k = \tilde{H}^k + \epsilon I$ for simplicity of presentation.

Theorem 1. The residual matrices

$$\mathcal{R}_S = \begin{pmatrix} \tilde{H}^k & F^T \\ F & 0 \end{pmatrix} \begin{pmatrix} \tilde{H}^k + \epsilon V V^T & F^T \\ F & 0 \end{pmatrix}^{-1} - I$$

and

$$\mathcal{R}_L = \begin{pmatrix} \tilde{H}^k & F^T \\ F & 0 \end{pmatrix} \begin{pmatrix} \tilde{H}^k & F^T \\ F & -\tau I \end{pmatrix}^{-1} - I$$

satisfy the bounds

$$\|\mathcal{R}_S\| \le \epsilon \cdot \mathrm{cond}(V) \max\left(1, \|S^{-1}\|\right), \tag{18}$$

$$\|\mathcal{R}_L\| \le \tau \cdot \mathrm{cond}(V) \cdot \max\left(\|S(S^2 + \tau C^2)^{-1}\|, \|C^2(S^2 + \tau C^2)^{-1}\|\right). \tag{19}$$

Proof. Bounds (18) and (19) follow from the identities

$$\mathcal{R}_S = \begin{pmatrix} V & \\ & U_2 \end{pmatrix} \begin{pmatrix} 0 & & -\epsilon S^{-1} \\ & -\frac{\epsilon}{1+\epsilon} I & \\ 0 & & 0 \end{pmatrix} \begin{pmatrix} V & \\ & U_2 \end{pmatrix}^{-1}$$

and

$$\mathcal{R}_L = \begin{pmatrix} V & \\ & U_2 \end{pmatrix} \cdot \begin{pmatrix} 0 & & 0 \\ & 0 & \\ \tau S(S^2 + \tau C^2)^{-1} & & -\tau C^2 (S^2 + \tau C^2)^{-1} \end{pmatrix} \begin{pmatrix} V & \\ & U_2 \end{pmatrix}^{-1}.$$

□

Theorem 1 implies the following asymptotic results in the situation where the smallest singular value $\sigma_{\min}(S)$ vanishes. The residuals for the Schur complement and augmented Lagrangian regularizations are small if $\epsilon \ll \sigma_{\min}(S)$ and $\tau \ll \sigma_{\min}(S)^2$. We thus conclude that the Schur based regularization has much wider interval of ϵ, where the regularization works, than the working interval of τ in the augmented Lagrangian regularization.

7. CASE STUDY: CHAIN OF OSCILLATING MASSES

Let us present preliminary results from numerical experiments for the classical linear control example of a chain of $n/2$ unit masses, connected by springs of unit stiffness without damping, where the ends are attached to fixed walls. We analyze computational performance and numerical robustness solving one QP at a particular sampling instant—a representative case for performing full closed-loop MPC simulations with cold initializations.

7.1 Problem Formulation and Algorithm Implementation

The continuous-time state-space system is discretized using a sample time $\Delta\tau = 0.5$ while keeping the m inputs constant between the sampling instants. The discrete-time dynamics are $x_{i+1} = A_i x_i + B_i u_i$, where

$$A_i = \exp\left(\Delta\tau A_c\right), \quad B_i = A_c^{-1}(A_i - I_n) \begin{bmatrix} 0_{n/2} \\ I_m \\ 0_{n/2-m} \end{bmatrix},$$

and

$$A_c = \begin{bmatrix} 0_{n/2} & I_{n/2} \\ T_{n/2} & 0_{n/2} \end{bmatrix}, \quad T = \begin{bmatrix} -2 & 1 & & \\ 1 & -2 & 1 & \\ & \ddots & \ddots & \ddots \\ & & 1 & -2 \end{bmatrix}.$$

The control inputs and states are subject to the inequality constraints

$$-0.5 \le u(i) \le 0.5, \qquad i = 0, \ldots, N-1,$$
$$-3.5 \le x(i) \le 3.5, \qquad i = 1, \ldots, N,$$

where $x(\cdot)$ denotes the position of the masses.

The cost matrices in the MPC formulation are $R_i = 10^{-6} I$, $S_i = 0$, and $Q_i = C^T C = \begin{bmatrix} I_{n/2} & 0 \end{bmatrix}^T \begin{bmatrix} I_{n/2} & 0 \end{bmatrix}$. The initial values of the state correspond to

$$-f_0 = 3.5 \begin{bmatrix} 1 & 1 & 0 & \cdots & 0 \end{bmatrix}^T.$$

Our numerical experiments in MATLAB are performed on a test problem of the following size: we consider 6 masses such that the state dimension equals $n = 12$, the number of control inputs is $m = 3$ and the control horizon length is $N = 30$. The number of inequality constraints at each control interval is $n + 2m = 18$.

Our MATLAB code implements the variant of the IP method described in more detail by (Shahzad et al., 2012) and uses the same algorithmic constants and stopping criterion such as $\sigma = 0.1$, $\gamma = 10^{-3}$, $\beta = 2$, $\epsilon = 10^{-7}$. We report the results of solving QP from a cold start at the initial time of the MPC controller, where the initial values for x, u, y, z, and s are equal to $\mathbf{1}$. The IP method for QP generally does not benefit much from warm starting, especially for S^k and Z^k, (Bartlett et al., 2000).

As a reference for further comparisons, we initially run the IP method in which the linear system (8) is solved exactly by the backslash operator in MATLAB. Our implementation returns a solution to the quadratic program after 20 IP iterations and with the 2-norm of the IP residual equal to $\|[r_H^T, r_F^T, r_G^T, r_S^T]^T\|_2 = 9.84 \cdot 10^{-8}$.

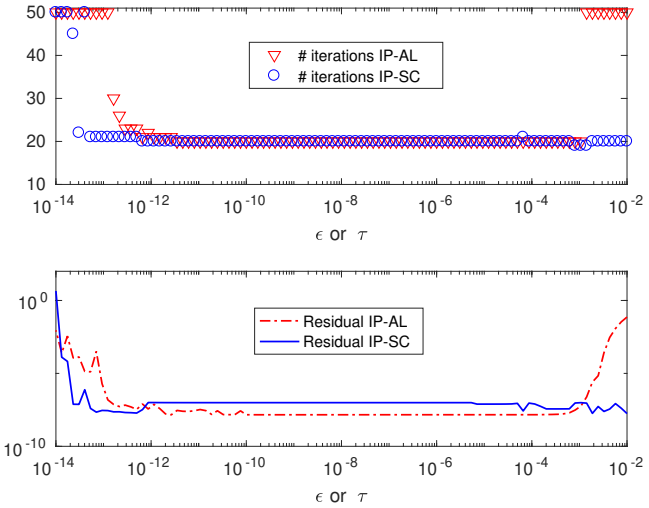

Fig. 1. The number of IP iterations and the 2-norm of the residual for the solution with respect to ϵ or τ.

Our proposed implementation of the IP method uses direct solvers based on the Schur complement (IP-SC) or the augmented Lagrangian (IP-AL) regularization approach. The IP-SC and IP-AL algorithms use regularization of the problem with a small value of the parameter ϵ and τ, respectively. Smaller values of ϵ and τ give better approximations to the KKT matrix, but too small values may lead to prohibitively large condition numbers, for which the IP method breaks down in our implementation using double precision arithmetic. Choosing the regularization parameter carefully becomes increasingly important when performing computations with a lower precision, e.g., using single precision arithmetic.

7.2 Numerical Stability and Parameter Sensitivity

Let us investigate the dependence of IP-SC on ϵ and IP-AL on τ. For this purpose, we choose 100 values of ϵ and τ uniformly distributed between 10^{-2} and 10^{-14} in the logarithmic scale. Figure 1 displays the number of IP iterations and the 2-norm of the IP residual with respect to the value of the parameter $\epsilon = \tau$.

Our numerical results for ϵ or τ between approximately 10^{-13} and 10^{-3} almost coincide with those for the backslash operator, i.e., the number of IP iterations equals 20 and the 2-norm of the IP residual is very close to 10^{-7}. Both algorithms have similar behavior in Figure 1 for the range 10^{-11}–10^{-3} of values of ϵ and τ. The IP-SC method is clearly preferential for practically important small values, below 10^{-11}, consistent with our analysis in Section 6. IP-SC and IP-AL are dramatically different for the larger values of ϵ and τ within the interval $[10^{-3.5}, 10^{-2}]$. The algorithm IP-SC continues executing well, while IP-AL has very large residuals and breaks down.

7.3 Preconditioned CG within the IP Method

We report our preliminary results performing experiments with preconditioned CG inside the IP method. The preconditioner is applied at even IP iterations and uses the Cholesky factorization from the previous odd IP iteration. Odd IP iterations are resolved exactly, which technically

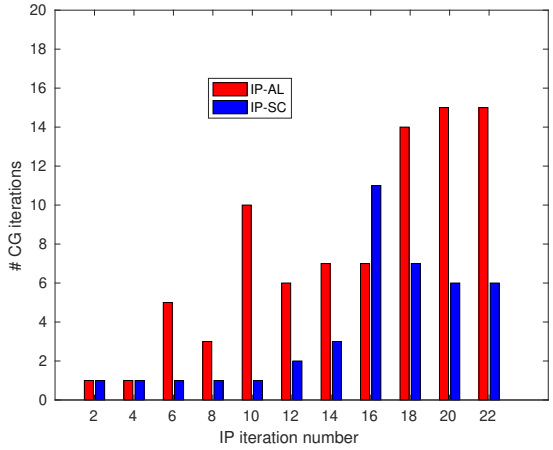

Fig. 2. The number of preconditioned CG iterations in the IP method for one MPC time step.

corresponds to one CG iteration. The dimension of the state is $n_x = 240$, corresponding to 120 masses, and the number of the control inputs is $n_u = 3$. We increase the dimension of the state 20 times, compared to the previous experiments, in order to make the test case representative for using CG, since running CG on small-size linear systems may result in atypical convergence behavior.

Figure 2 shows that the number of CG iterations in IP-AL with $\tau = 10^{-8}$ sometimes is growing with the number of IP iterations. This behavior can be observed also less strongly for the results of the IP-SC method in Figure 2. We anticipate that tuning of IP-AL and IP-SC parameters may further improve the performance of the preconditioned CG method in both cases.

Practical embedded MPC implementations require solving each QP within tight timing constraints. Therefore, it is common to limit the maximum number of IP iterations to a predetermined fixed value. Similarly, one may limit the number of CG iterations, e.g., up to 10 iterations in Figure 2, at the potential cost of an increase in the overall number of IP iterations.

Finally, one can change the rule of when to construct a new preconditioner, i.e. to compute the new structured Cholesky factorization. Instead of constructing the new preconditioner at each odd IP iteration, one could, e.g., apply the same preconditioner until the number of CG iterations passes a threshold, further increasing performance.

8. CONCLUSION

We consider preconditioned CG for computation of the search direction in the inexact Newton based IP method to solve each QP in MPC. We propose preconditioning by reusing matrix factors of direct solvers for the Schur complement of the normal equations or for the augmented Lagrangian regularization approach. We can apply preconditioned CG, e.g., at every even iteration of the IP method with the Cholesky factors computed in the previous odd iteration. Such a variant could allow computing optimal control solutions twice as fast compared to the traditional use of the direct solvers at every IP/Newton iteration.

We have carried out numerical tests of the two proposed regularizations for saddle point linear systems. On the one hand, smaller computational complexity of the preconditioners makes the augmented Lagrangian regularization based approach somewhat preferential. On the other hand, our perturbation analysis and numerical experiments demonstrate that preconditioning based on the more computationally expensive Schur complement is generally more robust to round-off errors, which is advantageous for low-precision embedded MPC implementations.

REFERENCES

Anderson, E., Bai, Z., Bischof, C., Blackford, L.S., Demmel, J., Dongarra, J., Croz, J.D., Greenbaum, A., Hammarling, S., McKenney, A., and Sorensen, D. (1999). *LAPACK Users' Guide*. SIAM, Philadelphia.

Bartlett, R., Wächter, A., and Biegler, L. (2000). Active set vs. interior point strategies for model predictive control. In *Proc. American Control Conf.*, 4229–4233. Chicago, IL.

Benzi, M. and Wathen, A.J. (2008). Some preconditioning techniques for saddle point problems. In W.H.A. Schilders et al. (ed.), *Model Order Reduction: Theory, Research Aspects and Applications*, Mathematics in Industry, vol 13., 195–211. Springer, Berlin, Heidelberg, Germany.

Diehl, M., Ferreau, H.J., and Haverbeke, N. (2009). Efficient numerical methods for nonlinear MPC and moving horizon estimation. In L. Magni et al. (ed.), *Nonlinear Model Predictive Control*, LNCIS 384, 391–417. Springer, Heidelberg, Germany.

Domahidi, A., Zgraggen, A.U., Zeilinger, M.N., Morari, M., and Jones, C.N. (2012). Efficient interior point methods for multistage problems arising in receding horizon control. In *Proc. 51st IEEE Conf. Decision and Control*, 668–674. Maui, Hawaii, USA.

Fletcher, R. (1987). *Practical Methods of Optimization*. Wiley, 2nd edition.

Frison, G. and Jørgensen, J.B. (2013). Efficient implementation of the Riccati recursion for solving linear-quadratic control problems. In *Proc. IEEE Multi-conference on Systems and Control*, 1117–1122. India.

Frison, G., Sorensen, H.B., Dammann, B., and Jørgensen, J.B. (2014). High-performance small-scale solvers for linear model predictive control. In *Proc. European Control Conf. (ECC)*, 128–133. France.

Gill, P., Golub, G., Murray, W., and Saunders, M.A. (1974). Methods for modifying matrix factorizations. *Math. Comp.*, 28(126), 505–535.

Golub, G. and Van Loan, C.V. (2013). *Matrix computations*. The Johns Hopkins Univ. Press, Baltimore, 4th edition.

Gondzio, J. (2012). Interior point methods 25 years later. *European J. Operational Research*, 218, 587–601.

Gondzio, J. (2013). Convergence analysis of an inexact feasible interior point method for convex quadratic programming. *SIAM J. Optim.*, 23(3), 1510–1527.

Kirches, C., Bock, H.G., Schlöder, J.P., and Sager, S. (2011). A factorization with update procedures for a KKT matrix arising in direct optimal control. *Math. Prog. Comp.*, 3(4), 319–348.

Knyazev, A., Fujii, Y., and Malyshev, A. (2015). Preconditioned continuation model predictive control. In *Proc. SIAM Conf. Control Appl.*, 101–108, Paris, France.

Knyazev, A. and Malyshev, A. (2015). Preconditioning for continuation model predictive control. *IFAC-PapersOnLine*, 48(23), 191–196.

Knyazev, A. and Malyshev, A. (2016). Sparse preconditioning for model predictive control. In *Proc. American Control Conf. (ACC)*, 4494–4499. Boston, MA, USA.

Knyazev, A.V. (2007). Observations on degenerate saddle point problems. *Computer Methods in Applied Mechanics and Engineering*, 196(37), 3742–3749.

Malyshev, A., Quirynen, R., Knyazev, A., , and Cairano, S.D. (2018). A regularized newton solver for linear model predictive control. In *Proc. European Control Conf. (ECC)*. Limassol, Cyprus.

Mehrotra, S. (1992). On the implementation of a primal-dual interior point method. *SIAM J. Optim.*, 2(4), 575–601.

Nocedal, J. and Wright, S.J. (2006). *Numerical Optimization*. Springer, second edition.

Oliveira, A.R.L. and Sorensen, D.C. (2005). A new class of preconditioners for large-scale linear systems from interior point methods for linear programming. *Linear Algebra Appl.*, 394, 1–24.

Quirynen, R., Knyazev, A., and Cairano, S.D. (2018). Block structured preconditioning within an active-set method for real-time optimal control. In *Proc. European Control Conf. (ECC)*. Limassol, Cyprus.

Rao, C.V., Wright, S.J., and Rawlings, J.B. (1998). Application of interior-point methods to model predictive control. *J. Optimiz. Theory Appl.*, 99, 723–757.

Rawlings, J.B., Mayne, D.Q., and Diehl, M. (2017). *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, Madison, WI, second edition.

Rees, T. and Greif, C. (2007). A preconditioner for linear systems arising from interior point optimization methods. *SIAM J. Sci. Comp.*, 29(5), 1992–2007.

Saunders, M. and Tomlin, J.A. (1996). Solving regularized linear programs using barrier methods and KKT systems. Tech. rep. sol 96-4, Systems Optimization Laboratory, Department of Operations Research, Stanford University, Stanford, CA 94305, USA.

Shahzad, A., Kerrigan, E.C., and Constantinides, G.A. (2010a). A fast well-conditionend interior point method for predictive control. In *Proc. IEEE Conf. Decision Control*, 508–513. Atlanta, GA, USA, December 15–17.

Shahzad, A., Kerrigan, E.C., and Constantinides, G.A. (2010b). Preconditioners for inexact interior point methods for predictive control. In *Proc. American Control Conf.*, 5714–5719. Baltimore, MD, USA.

Shahzad, A., Kerrigan, E.C., and Constantinides, G.A. (2012). A stable and efficient method for solving a convex quadratic program with application to optimal control. *SIAM J. Optim.*, 22(4), 1369–1393.

Vanderbei, R.J. (1995). Symmetric quasidefinite matrices. *SIAM J. Optimization*, 5, 100–113.

Wang, Y. and Boyd, S. (2010). Fast model predictive control using online optimization. *IEEE Trans. Control Syst. Technology*, 18(2), 267–278.

Wright, M.H. (1998). Ill-conditioning and computational error in interior point methods for nonlinear programming. *SIAM J. Optim.*, 9(1), 84–111.

Wright, S.J. (1997). *Primal-Dual Interior Point Methods*. SIAM, Philadelphia.