

Execution time certification for gradient-based optimization in model predictive control

Pontus Giselsson

Department of Automatic Control LTH

Lund University

Box 118, SE-221 00 Lund, Sweden

pontusg@control.lth.se

Abstract— We consider model predictive control (MPC) problems with linear dynamics, polytopic constraints, and quadratic objective. The resulting optimization problem is solved by applying an accelerated gradient method to the dual problem. The focus of this paper is to provide bounds on the number of iterations needed in the algorithm to guarantee a prespecified accuracy of the dual function value and the primal variables as well as guaranteeing a prespecified maximal constraint violation. The provided numerical example shows that the iteration bounds are tight enough to be useful in an inverted pendulum application.

I. INTRODUCTION

Model predictive control (MPC) is an optimization based control methodology that can handle state and control constraints (see [9], [10] for thorough descriptions of MPC). In the optimization problem a cost function is minimized based on predicted future state and control trajectories and subject to constraints. Optimal control and state trajectories are obtained and the first element in the input trajectory is applied to the system. This procedure is repeated every sampling instant which sets requirements on the execution time of the optimization problem. The topic of this paper is to provide certificates for the execution time of the optimization algorithm such that for every feasible initial condition the optimization algorithm provides a solution within the sampling time. We consider linear time-invariant systems with polytopic constraints and quadratic cost and a dual accelerated gradient method [6] is used to solve the resulting optimization problem.

For accelerated gradient methods there are convergence rate results [13], [2], [17], [6] that depend explicitly on the norm of the difference between the optimal solution and the initial iterate. If this norm can be bounded, a bound on the number of iterations to achieve a prespecified accuracy of the function value can be computed. This was done in [15] where input constrained MPC was considered. The condensed problem, i.e., the problem with all state variables eliminated, was solved using a fast gradient method. An iteration bound was obtained since the norm of the difference between the optimal solution and the initial iterate is bounded by the size of the input constraint set. Accelerated gradient methods can also be applied to the dual problem [16], [6]. To compute a bound on the number of iterations to achieve a prespecified accuracy, a bound on the norm of the difference

between the optimal dual variables and initial dual iterate is needed. This is more involved in the dual space than in a constrained primal space since dual variables are not chosen from a compact set. This is addressed in [16] where the equality constraints are dualized and a bound on the norm of the optimal dual variables is obtained using a recent result in [5]. The obtained bounds turn out to be quite conservative. Another method to provide computation time certificates in MPC is to bound the search time in the look-up table in explicit MPC [3], [1]. Practically this method is limited to small or medium-sized problems. For interior point methods, iteration bounds are available [11], these are, however, reported to be quite conservative [16], [11].

In this paper we consider the dual to the condensed problem, i.e., the dual to the problem where the state variables are eliminated. The resulting optimization problem has only inequality constraints and we apply the accelerated gradient method to the dual problem. To compute an iteration bound, we need a bound on the norm of the optimal dual variables. Using a result in [12] a bound to this norm can be computed if a Slater vector to the optimization problem is known. Computation of the norm bound requires that the distance from equality in the inequality constraints for the Slater vector is known, as well as the primal cost for the Slater vector. We will see that such a Slater vector can be constructed for almost all feasible initial conditions in the MPC case. The provided numerical example shows that the presented bounds are tight enough to give useful bounds in an inverted pendulum application.

II. PROBLEM SETUP AND PRELIMINARIES

We consider the problem of controlling a linear dynamical system to the origin subject to polytopic constraints. To achieve this we use MPC in which the following finite horizon optimization problem is solved at the current state $\bar{x} \in \mathbb{R}^n$:

$$\begin{aligned}
 V_N(\bar{x}) := & \min_{\mathbf{x}, \mathbf{u}} \quad \frac{1}{2} \sum_{t=0}^{N-1} (x_t^T Q x_t + u_t^T R u_t) + \frac{1}{2} x_N^T Q_N x_N \\
 \text{s.t.} \quad & (x_t, u_t) \in \mathcal{X} \times \mathcal{U}, \quad t = 0, \dots, N-1 \\
 & x_{t+1} = A x_t + B u_t, \quad t = 0, \dots, N-1 \\
 & x_N \in \mathcal{X}_f, \quad x_0 = \bar{x}
 \end{aligned} \tag{1}$$

where $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^m$ and $\mathbf{x} = [x_1^T, \dots, x_N^T]^T$ and $\mathbf{u} = [u_0^T, \dots, u_{N-1}^T]^T$. We use the standard assumptions that $Q \succeq 0$, $Q_N \succeq 0$ and $R \succ 0$. The constraint sets are assumed to be polytopes

$$\mathcal{X} = \{x \in \mathbb{R}^n \mid C_x x \leq d_x\}, \quad \mathcal{X}_f = \{x \in \mathbb{R}^n \mid C_f x \leq d_f\}$$

$$\mathcal{U} = \{u \in \mathbb{R}^m \mid C_u u \leq d_u\}.$$

Throughout this paper we assume that the sets \mathcal{X} , \mathcal{X}_f , and \mathcal{U} are non-empty and compact and that $0 \in \text{int } \mathcal{X}$, $0 \in \text{int } \mathcal{X}_f$, and $0 \in \text{int } \mathcal{U}$ which implies that $d_x, d_f, d_u > 0$. By introducing the following matrices

$$\mathbf{A} = \begin{pmatrix} A \\ \vdots \\ A^N \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} B & 0 & \cdots & 0 \\ AB & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ A^{N-1}B & \cdots & AB & B \end{pmatrix}$$

the predicted future state variables can be described in the current state \bar{x} and control variables \mathbf{u} as

$$\mathbf{x} = \mathbf{A}\bar{x} + \mathbf{B}\mathbf{u}.$$

We further define

$$\mathbf{Q} := \text{blkdiag}(Q, \dots, Q, Q_N), \quad \mathbf{R} := \text{blkdiag}(R, \dots, R),$$

$$\mathbf{C}_x := \text{blkdiag}(C_x, \dots, C_x, C_f), \quad \mathbf{d}_x := [d_x^T, \dots, d_x^T, d_f^T]^T,$$

$$\mathbf{C}_u := \text{blkdiag}(C_u, \dots, C_u), \quad \mathbf{d}_u := [d_u^T, \dots, d_u^T]^T.$$

The optimization problem (1) can, using these matrices, equivalently be written as

$$V_N(\bar{x}) = \min_{\mathbf{u}} J_N(\bar{x}, \mathbf{u}) := \frac{1}{2} \mathbf{u}^T \mathbf{H} \mathbf{u} + \bar{x}^T \mathbf{G} \mathbf{u} + \frac{1}{2} \bar{x}^T \mathbf{F} \bar{x}$$

$$\text{s.t. } g(\bar{x}, \mathbf{u}) \leq 0 \quad (2)$$

where $\mathbf{H} = \mathbf{B}^T \mathbf{Q} \mathbf{B} + \mathbf{R}$, $\mathbf{G} = \mathbf{A}^T \mathbf{Q} \mathbf{B}$, $\mathbf{F} = \mathbf{A}^T \mathbf{Q} \mathbf{A} + Q$, $g(\bar{x}, \mathbf{u}) = \mathbf{C} \mathbf{u} - d(\bar{x})$ and

$$\mathbf{C} = \begin{pmatrix} \mathbf{C}_u \\ \mathbf{C}_x \mathbf{B} \end{pmatrix}, \quad d(\bar{x}) = \begin{pmatrix} \mathbf{d}_u \\ \mathbf{d}_x - \mathbf{C}_x \mathbf{A} \bar{x} \end{pmatrix}.$$

To solve (2) we introduce dual variables $\boldsymbol{\mu} \in \mathbb{R}_{\geq 0}^p$ for the inequality constraints. The first $p_u \leq p$ dual variables in the dual variable vector $\boldsymbol{\mu}$ correspond to the input constraints and the last $p - p_u$ dual variables correspond to the state constraints. If Slater's condition holds, we get the following dual problem (cf. [4])

$$V_N(\bar{x}) = \max_{\boldsymbol{\mu} \geq 0} \min_{\mathbf{u}} \frac{1}{2} \mathbf{u}^T \mathbf{H} \mathbf{u} + \bar{x}^T \mathbf{G} \mathbf{u} + \boldsymbol{\mu}^T (\mathbf{C} \mathbf{u} - d(\bar{x})).$$

As shown in [6] the dual problem becomes

$$\max_{\boldsymbol{\mu} \geq 0} -\frac{1}{2} (C^T \boldsymbol{\mu} + G^T \bar{x})^T H^{-1} (C^T \boldsymbol{\mu} + G^T \bar{x}) - \boldsymbol{\mu}^T d(\bar{x}). \quad (3)$$

We define the dual function

$$D_N(\bar{x}, \boldsymbol{\mu}) = -\frac{1}{2} (C^T \boldsymbol{\mu} + G^T \bar{x})^T H^{-1} (C^T \boldsymbol{\mu} + G^T \bar{x}) - \boldsymbol{\mu}^T d(\bar{x})$$

which satisfies the following properties (cf. [6]).

Proposition 1: The dual function has Lipschitz continuous gradient with Lipschitz constant $L = \|CH^{-1}C^T\|$ and the

gradient is given by $\nabla D_N(\bar{x}, \boldsymbol{\mu}) = -CH^{-1}(C^T \boldsymbol{\mu} + G^T \bar{x}) - d(\bar{x})$.

This implies that the dual function can be maximized using an accelerated gradient method [13], [2], [17], [6]. The algorithm presented in [6] with a cold-starting strategy, i.e., $\boldsymbol{\mu}^0 = 0$ is presented below.

Algorithm 1: Accelerated gradient algorithm

Initialize $\boldsymbol{\mu}^0 = \boldsymbol{\mu}^{-1} = 0$ and $\mathbf{u}^{-1} = -H^{-1}G^T \bar{x}$.

For $k \geq 0$

$$\mathbf{u}^k = -H^{-1}(C^T \boldsymbol{\mu}^k + G^T \bar{x})$$

$$\tilde{\mathbf{u}}^k = \mathbf{u}^k + \frac{k-1}{k+2}(\mathbf{u}^k - \mathbf{u}^{k-1})$$

$$\boldsymbol{\mu}^{k+1} = \max \left\{ 0, \boldsymbol{\mu}^k + \frac{k-1}{k+2}(\boldsymbol{\mu}^k - \boldsymbol{\mu}^{k-1}) + \frac{1}{L} (C \tilde{\mathbf{u}}^k - d(\bar{x})) \right\}$$

Before we state the convergence rate properties of the algorithm, we introduce the set of optimal dual variables

$$M^*(\bar{x}) = \{\boldsymbol{\mu} \in \mathbb{R}_{\geq 0}^p \mid D_N(\bar{x}, \boldsymbol{\mu}) \geq V_N(\bar{x})\}.$$

We also introduce \mathbb{X}_N which is the steerable set defined as

$$\mathbb{X}_N = \{\bar{x} \in \mathbb{R}^n \mid \text{there exist } \mathbf{u} \text{ s.t. } C \mathbf{u} \leq d(\bar{x})\}.$$

Remark 1: From [14], we know that the steerable set \mathbb{X}_N is convex and that $0 \in \mathbb{X}_N$.

We also denote by $\mathbf{u}^*(\bar{x})$ the optimal solution to (2) with initial condition \bar{x} and $\underline{\sigma}(H)$ the smallest eigenvalue to H .

Proposition 2: Suppose that $\bar{x} \in \mathbb{X}_N$. For any $\boldsymbol{\mu}^* \in M^*(\bar{x})$ Algorithm 1 has the following convergence rate properties:

1) The dual function converges as

$$D_N(\bar{x}, \boldsymbol{\mu}^*) - D_N(\bar{x}, \boldsymbol{\mu}^k) \leq \frac{2L\|\boldsymbol{\mu}^*\|^2}{(k+1)^2}, \forall k \geq 1. \quad (4)$$

2) The primal variable rate of convergence is

$$\|\mathbf{u}^k - \mathbf{u}^*(\bar{x})\|^2 \leq \frac{4L\|\boldsymbol{\mu}^*\|^2}{\underline{\sigma}(H)(k+1)^2}, \forall k \geq 1. \quad (5)$$

3) The constraint violation is bounded by

$$\|g(\bar{x}, \mathbf{u}^k) - g(\bar{x}, \mathbf{u}^*(\bar{x}))\|^2 \leq \frac{4L^2\|\boldsymbol{\mu}^*\|^2}{(k+1)^2}, \forall k \geq 1.$$

Proof. Argument 1 is proven in [2], [17], [6] and argument 2 is proven in [6]. To prove the third argument we have

$$\begin{aligned} \|g(\bar{x}, \mathbf{u}^k) - g(\bar{x}, \mathbf{u}^*(\bar{x}))\|^2 &= \\ &= \|C \mathbf{u}^k - d(\bar{x}) - (C \mathbf{u}^*(\bar{x}) - d(\bar{x}))\|^2 \\ &= \|\nabla D_N(\bar{x}, \boldsymbol{\mu}^k) - \nabla D_N(\bar{x}, \boldsymbol{\mu}^*)\|^2 \\ &\leq 2L(-\langle -\nabla D_N(\bar{x}, \boldsymbol{\mu}^*), \boldsymbol{\mu}^k - \boldsymbol{\mu}^* \rangle \\ &\quad + D_N(\bar{x}, \boldsymbol{\mu}^*) - D_N(\bar{x}, \boldsymbol{\mu}^k)) \\ &\leq 2L(D_N(\bar{x}, \boldsymbol{\mu}^*) - D_N(\bar{x}, \boldsymbol{\mu}^k)). \end{aligned}$$

The first inequality comes from [13, Theorem 2.1.5] since $-D_N$ is convex. The second inequality is due to first order

optimality condition [13, Theorem 2.2.5] for the convex function $-D_N$. It is left to apply Argument 1 to prove the result. \square

The objective of the paper is to, a priori, compute bounds on the number of iterations needed to achieve a prespecified dual function, primal variable, and constraint satisfaction tolerance when initializing the algorithm with $\mu^0 = 0$. These bounds should ideally hold for any initial state $\bar{x} \in \mathbb{X}_N$. In this paper we will show how to compute bounds that hold for any $\bar{x} \in \beta\mathbb{X}_N$ where $\beta \in (0, 1)$ and $\beta\mathbb{X}_N$ is defined as

$$\beta\mathbb{X}_N := \{\bar{x} \in \mathbb{R}^n \mid \frac{1}{\beta}\bar{x} \in \mathbb{X}_N\}.$$

From the definition and Remark 1 we conclude that $\beta\mathbb{X}_N \subseteq \mathbb{X}_N$ and that $0 \in \beta\mathbb{X}_N$. Before we proceed with the presentation we introduce the following definition.

Definition 1: We define $\kappa \geq 1$ as the smallest scalar such that for every $\bar{x} \in \mathbb{X}_N$ the following holds

$$V_N(\bar{x}) \leq \kappa \min_{\mathbf{u}} J_N(\bar{x}, \mathbf{u}).$$

Remark 2: The optimal solution to $\min_{\mathbf{u}} J_N(\bar{x}, \mathbf{u})$ is $\mathbf{u}_{uc}^*(\bar{x}) = -H^{-1}G^T\bar{x}$. The corresponding cost becomes

$$\begin{aligned} \min_{\mathbf{u}} J_N(\bar{x}, \mathbf{u}) &= \frac{1}{2}\bar{x}^T GH^{-1}G^T\bar{x} - \bar{x}^T GH^{-1}G^T\bar{x} + \frac{1}{2}\bar{x}^T F\bar{x} \\ &= \frac{1}{2}\bar{x}^T (F - GH^{-1}G^T)\bar{x}. \end{aligned}$$

By defining $P := F - GH^{-1}G^T$ where $P \succ 0$ we get

$$V_N(\bar{x}) \leq \kappa \min_{\mathbf{u}} J_N(\bar{x}, \mathbf{u}) = \frac{\kappa}{2}\bar{x}^T P\bar{x}.$$

Also, note that we have

$$V_N(\bar{x}) \geq \min_{\mathbf{u}} J_N(\bar{x}, \mathbf{u}) = \frac{1}{2}\bar{x}^T P\bar{x}.$$

A. Notation

We denote by \mathbb{R} the real numbers and by $\mathbb{R}_{\geq 0}$ non-negative real numbers. The norm $\|\cdot\|$ refers to the Euclidean norm or the induced Euclidean norm unless otherwise is specified and $\langle x, y \rangle = x^T y$. Further $\bar{\sigma}(H)$ denotes the largest singular value of H and $\underline{\sigma}(H)$ denotes the smallest singular value of H . Further $[\cdot]_i$ denotes the i :th element in the vector.

III. LAGRANGE MULTIPLIER NORM BOUNDS

All quantities in the bounds in Proposition 2 are known except for $\|\mu^*\|$ where $\mu^* \in M^*(\bar{x})$. This section is devoted to bounding the norm of the optimal dual variables in (3) for any $\bar{x} \in \beta\mathbb{X}_N$ where $\beta \in (0, 1)$. The following result is used to achieve this.

Lemma 1: Assume that $\bar{\mathbf{u}}(\bar{x})$ is a Slater vector, i.e., that $\bar{\mathbf{u}}(\bar{x})$ satisfies $C\bar{\mathbf{u}}(\bar{x}) < d(\bar{x})$. Then

$$\max_{\mu \in M^*(\bar{x})} \|\mu\| \leq \frac{1}{\gamma(\bar{x}, \bar{\mathbf{u}}(\bar{x}))} (J_N(\bar{x}, \bar{\mathbf{u}}(\bar{x})) - V_N(\bar{x}))$$

where $\gamma(\bar{x}, \bar{\mathbf{u}}(\bar{x})) := \min_{1 \leq j \leq p} [-g(\bar{x}, \bar{\mathbf{u}}(\bar{x}))]_j$.

Proof. A proof is provided in [12]. \square

Thus, if we can find a Slater vector for any initial condition $\bar{x} \in \beta\mathbb{X}_N$ we can bound the norm of the optimal Lagrange multipliers, μ^* . In the following lemma we show how to construct a Slater vector to (2) for any initial state $\bar{x} \in \beta\mathbb{X}_N$. Before we present the lemma the following notation is introduced; $\mathbf{d} := [\mathbf{d}_u^T, \mathbf{d}_x^T]^T$ and $d_{\min} := \min_j [\mathbf{d}]_j$ which implies that $d_{\min} > 0$.

Lemma 2: For every $\bar{x} \in \beta\mathbb{X}_N$ with $\beta \in (0, 1)$, $\bar{\mathbf{u}}(\bar{x}) = \beta\mathbf{u}^*(\bar{x}/\beta)$ is a Slater vector to the optimization problem (2). The Slater vector satisfies $\gamma(\bar{x}, \bar{\mathbf{u}}(\bar{x})) \geq (1 - \beta)d_{\min}$.

Proof. Since $\bar{x} \in \beta\mathbb{X}_N$ we have by definition that $\bar{x}/\beta \in \mathbb{X}_N$. The optimal control trajectory at \bar{x}/β is $\mathbf{u}^*(\bar{x}/\beta)$. Since $\bar{x}/\beta \in \mathbb{X}_N$ the optimal control trajectory is feasible, i.e., the following holds

$$g(\frac{\bar{x}}{\beta}, \mathbf{u}^*(\frac{\bar{x}}{\beta})) = \begin{pmatrix} \mathbf{C}_u \mathbf{u}^*(\bar{x}/\beta) - \mathbf{d}_u \\ \mathbf{C}_x (\mathbf{A}\bar{x}/\beta + \mathbf{B}\mathbf{u}^*(\bar{x}/\beta)) - \mathbf{d}_x \end{pmatrix} \leq 0.$$

For any $\bar{x} \in \beta\mathbb{X}_N$ we have for the chosen Slater vector $\bar{\mathbf{u}}(\bar{x}) = \beta\mathbf{u}^*(\bar{x}/\beta)$ that

$$\begin{aligned} g(\bar{x}, \beta\mathbf{u}^*(\frac{\bar{x}}{\beta})) &= \begin{pmatrix} \mathbf{C}_u \beta\mathbf{u}^*(\bar{x}/\beta) - \mathbf{d}_u \\ \mathbf{C}_x (\mathbf{A}\bar{x} + \mathbf{B}\beta\mathbf{u}^*(\bar{x}/\beta)) - \mathbf{d}_x \end{pmatrix} \\ &= \begin{pmatrix} \beta(\mathbf{C}_u \mathbf{u}^*(\bar{x}/\beta) - \mathbf{d}_u) + (\beta - 1)\mathbf{d}_u \\ \beta(\mathbf{C}_x (\mathbf{A}\bar{x}/\beta + \mathbf{B}\mathbf{u}^*(\bar{x}/\beta)) - \mathbf{d}_x) + (\beta - 1)\mathbf{d}_x \end{pmatrix} \\ &= \beta \begin{pmatrix} \mathbf{C}_u \mathbf{u}^*(\bar{x}/\beta) - \mathbf{d}_u \\ \mathbf{C}_x (\mathbf{A}\bar{x}/\beta + \mathbf{B}\mathbf{u}^*(\bar{x}/\beta)) - \mathbf{d}_x \end{pmatrix} - \begin{pmatrix} (1 - \beta)\mathbf{d}_u \\ (1 - \beta)\mathbf{d}_x \end{pmatrix} \\ &\leq - \begin{pmatrix} (1 - \beta)\mathbf{d}_u \\ (1 - \beta)\mathbf{d}_x \end{pmatrix}. \end{aligned}$$

This gives

$$\begin{aligned} \gamma(\bar{x}, \bar{\mathbf{u}}(\bar{x})) &= \min_{1 \leq j \leq p} [-g(\bar{x}, \beta\mathbf{u}^*(\bar{x}/\beta))]_j \\ &\geq (1 - \beta) \min([\mathbf{d}]_j) = (1 - \beta)d_{\min}. \end{aligned}$$

This completes the proof. \square

By limiting the set of initial states, a Slater vector can be constructed with a certain distance to equality in the inequality constraints. Using this result the following theorem provides a bound on the norm of the optimal dual variables.

Theorem 1: For every $\bar{x} \in \beta\mathbb{X}_N$ we have that

$$\max_{\mu \in M^*(\bar{x})} \|\mu\| \leq \frac{\kappa - 1}{2(1 - \beta)d_{\min}} \bar{x}^T P\bar{x}. \quad (6)$$

Proof. We will show that Lemma 1 gives (6) using the Slater vector $\bar{\mathbf{u}}(\bar{x}) = \beta\mathbf{u}^*(\bar{x}/\beta)$. We have

$$\begin{aligned} J_N(\bar{x}, \beta\mathbf{u}^*(\bar{x}/\beta)) &= \\ &= \frac{1}{2}(\beta\mathbf{u}^*(\bar{x}/\beta))^T H \beta\mathbf{u}^*(\bar{x}/\beta) + \bar{x}^T G \beta\mathbf{u}^*(\bar{x}/\beta) + \frac{1}{2}\bar{x}^T F\bar{x} \\ &= \frac{\beta^2}{2} \left((\mathbf{u}^*(\frac{\bar{x}}{\beta}))^T H \mathbf{u}^*(\frac{\bar{x}}{\beta}) + 2 \left[\frac{\bar{x}}{\beta} \right]^T G \mathbf{u}^*(\frac{\bar{x}}{\beta}) + \left[\frac{\bar{x}}{\beta} \right]^T F \left[\frac{\bar{x}}{\beta} \right] \right) \\ &= \beta^2 V_N(\bar{x}/\beta) \leq \frac{\kappa}{2} \beta^2 \left[\frac{\bar{x}}{\beta} \right]^T P \left[\frac{\bar{x}}{\beta} \right] = \frac{\kappa}{2} \bar{x}^T P\bar{x} \end{aligned}$$

where the inequality comes from Remark 2. From Lemma 1 and Lemma 2 we have

$$\begin{aligned} \max_{\mu \in M^*(\bar{x})} \|\mu\| &\leq \frac{J_N(\bar{x}, \beta \mathbf{u}^*(\bar{x}/\beta)) - V_N(\bar{x})}{\gamma(\bar{x}, \beta \mathbf{u}^*(\bar{x}/\beta))} \\ &\leq \frac{1}{(1-\beta)d_{\min}} \left(\frac{\kappa}{2} \bar{x}^T P \bar{x} - V_N(\bar{x}) \right) \\ &\leq \frac{\kappa-1}{2(1-\beta)d_{\min}} \bar{x}^T P \bar{x} \end{aligned}$$

where the last inequality is due to Remark 2. This completes the proof. \square

Remark 3: If Definition 1 is changed such that κ_{β_1} is the smallest scalar such that for all $\bar{x} \in \beta_1 \mathbb{X}_N$ and for some $\beta_1 \in (0, 1)$ we have an upper bound $V_N(\bar{x}) \leq \frac{\kappa_{\beta_1}}{2} \bar{x}^T P \bar{x}$. Then for every $\bar{x} \in \beta_2 \mathbb{X}_N$ where $\beta_2 \in (0, \beta_1)$ it is straightforward to verify that

$$\max_{\mu \in M^*(\bar{x})} \|\mu\| \leq \frac{\kappa_{\beta_1} - 1}{2(1 - \beta_2/\beta_1)d_{\min}} \bar{x}^T P \bar{x}.$$

If $\frac{\kappa_{\beta_1} - 1}{1 - \beta_2/\beta_1} < \frac{\kappa - 1}{1 - \beta_2}$ we get an improved bound on the norm of the dual variables compared to Theorem 1.

The provided bound on the norm of optimal dual variables can, together with Proposition 2, be used to bound the number of iterations to get a prespecified accuracy in the function value, primal variables and constraint violation. This is the topic of the following section.

IV. ALGORITHM ITERATION BOUNDS

In this section we provide bounds on the number of iterations within which a dual ϵ_d -solution, ϵ_c constraint violation and ϵ_p norm-distance to the primal optimal solution are guaranteed. The bounds are developed for the cold starting case, i.e., when the initial iterate is $\mu^0 = 0$.

A. Iteration bound to guarantee dual ϵ -solution

The first bound is on the number of iterations within which a dual ϵ -solution is guaranteed. To avoid that scaling the Q and R -matrices give different bounds we use a relative tolerance.

Theorem 2: Suppose that Algorithm 1 is initialized with $\mu^0 = 0$. Then for every $\bar{x} \in \beta \mathbb{X}_N$ with $\beta \in (0, 1)$ we have

$$V_N(\bar{x}) - D_N(\bar{x}, \mu^k) \leq \epsilon_d V_N(\bar{x}) \quad (7)$$

if

$$k \geq k_d(\bar{x}) := \sqrt{\frac{L \bar{x}^T P \bar{x}}{\epsilon_d}} \frac{\kappa - 1}{(1 - \beta)d_{\min}} - 1.$$

Proof. Inequality (7) is equivalent to

$$D_N(\bar{x}, \mu^*) - D_N(\bar{x}, \mu^k) \leq \epsilon_d D_N(\bar{x}, \mu^*)$$

for any $\mu^* \in M^*(\bar{x})$. From Proposition 2 and Theorem 1 we have that

$$\begin{aligned} D_N(\bar{x}, \mu^*) - D_N(\bar{x}, \mu^k) &\leq \frac{2L \|\mu^*\|^2}{(k+1)^2} \leq \\ &\leq \frac{2L(\kappa-1)^2}{4(1-\beta)^2 d_{\min}^2 (k+1)^2} (\bar{x}^T P \bar{x})^2. \end{aligned}$$

Since $\frac{1}{2} \bar{x}^T P \bar{x} \leq V_N(\bar{x})$, we have that

$$\frac{2L(\kappa-1)^2}{4(1-\beta)^2 d_{\min}^2 (k+1)^2} (\bar{x}^T P \bar{x})^2 \leq \epsilon_d \frac{1}{2} \bar{x}^T P \bar{x} \quad (8)$$

implies (7). Rearranging the terms gives the result. \square

Remark 4: By scaling the penalty-matrices $Q_a = aQ$, $R_a = aR$ we get $H_a = aH$ which implies $L_a = \|CH_a^{-1}C^T\| = \frac{1}{a}\|CH^{-1}C^T\| = \frac{1}{a}L$ and $P_a = aP$. Thus, using a relative tolerance the same bound is obtained for every scaling factor $a > 0$.

Remark 5: To get a bound that holds for all $\bar{x} \in \beta \mathbb{X}_N$, $k_d(\bar{x})$ should be maximized subject to $\bar{x} \in \beta \mathbb{X}_N$. An over-estimator is to maximize $k_d(\bar{x})$ subject to $\bar{x} \in \beta \mathcal{X}$ which is readily available. The resulting maximization problem depends affinely on $\sqrt{\bar{x}^T P \bar{x}}$, hence the maximizing argument can be found by maximizing $\bar{x}^T P \bar{x}$ on $\beta \mathcal{X}$. This is a quadratic maximization problem that can be rewritten as a mixed integer linear program (MILP) as shown in [8]. MILP software produce upper and lower bounds to the objective in each iteration and since an upper bound to the objective is enough to compute an iteration bound, the optimization can be stopped when sufficient accuracy is achieved.

B. Iteration bound for constraint violation

In this section we bound the number of iterations within which a prespecified constraint violation is guaranteed. We use the following relative tolerance $g(\bar{x}, \mathbf{u}^k) \leq \epsilon_d$.

Theorem 3: Suppose that Algorithm 1 is initialized with $\mu^0 = 0$. Then, $g(\bar{x}, \mathbf{u}^k) \leq \epsilon_c d$ holds for every $\bar{x} \in \beta \mathbb{X}_N$ if

$$k \geq k_c(\bar{x}) := \frac{L(\kappa-1)\bar{x}^T P \bar{x}}{(1-\beta)d_{\min}^2 \epsilon_c} - 1.$$

Proof. First note that if $\|g(\bar{x}, \mathbf{u}^k) - g(\bar{x}, \mathbf{u}^*)\| \leq \epsilon_c d_{\min}$ then $g(\bar{x}, \mathbf{u}^k) \leq \epsilon_c d$ since $g(\bar{x}, \mathbf{u}^*) \leq 0$. From Proposition 2 and Theorem 1 we get

$$\|g(\bar{x}, \mathbf{u}^k) - g(\bar{x}, \mathbf{u}^*(\bar{x}))\| \leq \frac{2L \|\mu^*\|}{k+1} \leq \frac{L(\kappa-1)\bar{x}^T P \bar{x}}{(1-\beta)d_{\min}(k+1)}.$$

Setting this $\leq \epsilon_c d_{\min}$ and rearranging the terms gives the iteration bound. \square

Remark 6: This result can be used in a constraint tightening approach to guarantee a feasible solution w.r.t. to the original constraint sets within $k_c(\bar{x})$ iterations.

C. Primal variable iteration bound

Using the same techniques it is also possible to bound the number of iterations needed to guarantee a primal solution that is within a prespecified distance to the optimal solution.

Theorem 4: Suppose that Algorithm 1 is initialized with $\mu^0 = 0$. Then, for every $\bar{x} \in \beta \mathbb{X}_N$ we have

$$\|\mathbf{u}^k - \mathbf{u}^*(\bar{x})\| \leq \epsilon_p$$

if

$$k \geq k_p(\bar{x}) := \sqrt{\frac{L}{\sigma(H)}} \frac{(\kappa-1)\bar{x}^T P \bar{x}}{\epsilon_p(1-\beta)d_{\min}} - 1.$$

Proof. From Proposition 2 and Theorem 1 we have

$$\begin{aligned}\|\mathbf{u}^k - \mathbf{u}^*(\bar{x})\| &\leq \sqrt{\frac{L}{\underline{\sigma}(H)}} \frac{2\|\boldsymbol{\mu}^*\|}{(k+1)} \\ &\leq \sqrt{\frac{L}{\underline{\sigma}(H)}} \frac{(\kappa-1)x^T P x}{(1-\beta)d_{\min}(k+1)}.\end{aligned}$$

Setting this $\leq \epsilon_p$ and rearranging gives the result. \square

V. PRECONDITIONING

There are two different ways of preconditioning the problem data to possibly achieve smaller iteration bounds. One is to do a variable change in the primal variables and another is to scale the matrices defining the inequality constraints. We start by considering scaling the matrices defining the inequality constraints.

A. Scaling inequality constraints

All iteration bounds k_d, k_c, k_p depend on \sqrt{L}/d_{\min} or L/d_{\min}^2 . By introducing $D = \text{diag}(\mathbf{d})$ and recalling the definition of L we get $L/d_{\min}^2 = \|CH^{-1}C^T\|/\lambda_{\min}^2(D)$. By scaling the inequality constraints, this ratio can be minimized to get less conservative bounds without affecting the solutions of the optimization problem. We introduce the scaling matrix $S = \text{blkdiag}(S_u, S_x)$ where $S_u = \text{diag}(s_1, \dots, s_{p_u})$, $S_x = \text{diag}(s_{p_u+1}, \dots, s_p)$ where $p_u \leq p$ and all elements $s_i > 0, i = 1, \dots, p$. We get the following scaling

$$SC\mathbf{u} \leq Sd(\bar{x}).$$

From the definition of C and $d(\bar{x})$ we see that this is equivalent to

$$S_u C_u \leq S_u \mathbf{d}_u \quad S_x C_x (\mathbf{A}\bar{x} + \mathbf{B}\mathbf{u}) \leq S_x \mathbf{d}_x.$$

The scaling of constraints will give as small bounds as possible if the scaling is chosen according to the following minimization

$$\min_S \frac{\|SCH^{-1}C^T S\|}{\lambda_{\min}^2(SD)}.$$

We introduce $\bar{S} = SD$ which is a diagonal matrix with strictly positive elements since it is a product of two diagonal matrices with strictly positive elements. This gives the equivalent minimization problem

$$\min_{\bar{S}} \frac{\|\bar{S}D^{-1}CH^{-1}C^T D^{-1}\bar{S}\|}{\lambda_{\min}^2(\bar{S})}. \quad (9)$$

It was shown in [16, Lemma 1] that for invertible \bar{S} , an optimal solution is $\bar{S} = I$. Since diagonal matrices with positive elements are a subset of all invertible matrices, we get that $\bar{S} = I$ minimizes (9). The optimal scaling becomes $S = \bar{S}D^{-1} = D^{-1}$.

B. Preconditioning of primal variables

When performing a linear change of variables in primal variables, i.e., set $q = T^{-1}\mathbf{u}$ where T is an invertible matrix, H , G and C must be changed accordingly to not affect the primal optimal solution. We get $H_q = T^T H T$, $G_q = G T$ and $C_q = C T$. The Lipschitz constant does not change since

$$\begin{aligned}L_q &= \|C_q H_q^{-1} C_q^T\| = \|C T T^{-1} H^{-1} T^{-T} T^T C^T\| \\ &= \|C H^{-1} C^T\| = L.\end{aligned}$$

Straightforward verification of the algorithm when initialized with $\boldsymbol{\mu}^0 = 0$ gives that the $\boldsymbol{\mu}^k$ -sequence is identical whether using the new variables q or the original variables \mathbf{u} . It is also straightforward to verify that the relation between the iterates in the new variable q^k and the iterates in the original variable \mathbf{u}^k is $q^k = T^{-1}\mathbf{u}^k$. Thus, we do not get better (or worse) convergence properties by preconditioning the primal variables.

VI. NUMERICAL EXAMPLE

We evaluate the conservatism of the iteration bounds by applying them to a double integrator system and a double integrator with a pendulum attached. We consider the pendulum in [7] with pendulum length $l = 0.4\text{m}$. The cart has inner control loops that make it behave as a double integrator. We choose sample time $h = 0.02\text{s}$ as in [7]. We get the following discrete time dynamics for the pendulum system when the pendulum is in its inverted position (cf. [7])

$$x^+ = \begin{pmatrix} 1 & 0.02 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1.0049 & 0.0200 \\ 0 & 0 & 0.4913 & 1.0049 \end{pmatrix} x + \begin{pmatrix} 0.0002 \\ 0.02 \\ -0.0005 \\ -0.0501 \end{pmatrix} u.$$

The state variables are $x = [p \ \dot{p} \ \theta \ \dot{\theta}]^T$ where p is cart position, \dot{p} is cart velocity, θ is pendulum angle and $\dot{\theta}$ is pendulum angular velocity. The double integrator system is the system consisting of only the first two states, $[p \ \dot{p}]$. We have the following constraints

$$\begin{aligned}-0.5 &\leq p \leq 0.5 & -1 &\leq \dot{p} \leq 1 & -5 &\leq u \leq 5 \\ -0.2 &\leq \theta \leq 0.2 & -0.5 &\leq \dot{\theta} \leq 0.5.\end{aligned}$$

The objective is to minimize

$$\sum_{t=0}^{N-1} (x_t^T Q x_t + u_t^T R u_t) + x_N^T P x_N$$

where $Q = \text{diag}(1, 0.3, 0.3, 0.1)$, $R = 0.1$ and P is the infinite horizon cost for the unconstrained LQ-problem with weighting matrices Q and R . Further we choose the terminal set $\mathcal{X}_f = \mathcal{X}$.

In Figures 1 and 2 we compare the iteration bounds with the worst case actual number of iterations and the maximum number of iterations, k_{req} , to guarantee an execution time less than $h = 0.02\text{s}$ on a machine with 1 Gflops/s computing power. If implemented wisely, the number of flops per iteration in Algorithm 1 is $2(pN)^2 + 7pN$ and we get

$$k_{\text{req}} = \frac{10^9 h}{2(pN)^2 + 7pN}.$$




Fig. 1. The figure shows the iteration bounds (k_d, k_c, k_p), the actual number of iterations ($k_d^{\text{act}}, k_c^{\text{act}}, k_p^{\text{act}}$), and the maximal number of iterations needed to certify execution time within $h=0.02s$ for the double integrator system (k_{req}).




Fig. 2. The figure shows the iteration bounds (k_d, k_c, k_p), the actual number of iterations ($k_d^{\text{act}}, k_c^{\text{act}}, k_p^{\text{act}}$), and the maximal number of iterations needed to certify execution time within $h=0.02s$ for the pendulum system (k_{req}).

In all examples, we use control horizon $N = 10$, accuracy requirements $\epsilon_d = 0.01$ and $\epsilon_c = \epsilon_p = 0.05$. In Figure 1, the results for the double integrator are presented. On the x-axis β in $\beta\mathbb{X}_N$ is plotted and on the y-axis the iterations bounds and the actual number of iterations are plotted. We are able to certify that the optimization algorithm will terminate with a close to optimal solution for all $\bar{x} \in 0.925\mathbb{X}_N$ within the sampling time, $h = 0.02s$. We also see that for $\bar{x} \in 0.825\mathbb{X}_N$ we can guarantee that the optimal solution is found in one iteration, i.e., that no constraints are active.

In Figure 2, the results for the inverted pendulum system are presented. Also here we have β in $\beta\mathbb{X}_N$ on the x-axis and the iteration bounds and the actual number of iterations on the y-axis. We are able to certify that for $\bar{x} \in 0.6\mathbb{X}_N$ that the required accuracy is achieved within the sampling time, $h = 0.02s$. We can also certify that a dual ϵ_d -solution is found within the sampling time for any $\bar{x} \in 0.9\mathbb{X}_N$. We see that for large parts of the steerable set, \mathbb{X}_N , the iteration bounds give meaningful results that can be used to certify the MPC-controller with respect to execution time.

VII. CONCLUSIONS AND FUTURE WORK

We solve the optimization problems arising in MPC with linear dynamics, polytopic constraints, and a quadratic cost using a dual accelerated gradient method [6]. By constructing

Slater vectors to the optimization problems, we are able to bound the norm of the optimal dual variables. This is used to compute iteration bounds on the number of iterations within which a certain accuracy of the dual function value, constraint violation, and primal variables is guaranteed. The provided numerical example shows that the bounds are tight enough to be useful in a pendulum application. A future work direction is to search for tighter iteration bounds when using warm-starting strategies.

REFERENCES

- [1] Alessandro Alessio and Alberto Bemporad. A survey on explicit model predictive control. In *Nonlinear Model Predictive Control*, pages 345–369. 2009.
- [2] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2(1):183–202, October 2009.
- [3] A. Bemporad, M. Morari, V. Dua, and E.N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, January 2002.
- [4] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- [5] O. Devolder, F. Glineur, and Y. Nesterov. Double smoothing technique for convex optimization problems in Hilbert spaces with applications to optimal control. Technical report, Center for Operations Research and Econometrics (CORE), Universit catholique de Louvain (UCL), Belgium, January 2011.
- [6] P. Giselsson, M. D. Doan, T. Keviczky, B. De Schutter, and A. Rantzer. Accelerated gradient methods and dual decomposition in distributed model predictive control. To appear. Available: <http://www.control.lth.se/Staff/PontusGiselsson.html>, 2012.
- [7] Pontus Giselsson. Model predictive control in a pendulum system. In *Proceedings of the 31st IASTED conference on Modelling, Identification and Control*, Innsbruck, Austria, February 2011.
- [8] C.N. Jones and M. Morari. Approximate explicit MPC using bilevel optimization. In *European Control Conference*, Budapest, Hungary, August 2009.
- [9] J. M. Maciejowski. *Predictive control with constraints*. Prentice Hall, Essex, England, 2002.
- [10] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789 – 814, June 2000.
- [11] L. K. McGovern. *Computational Analysis of Real-time Convex Optimization for Control Systems*. PhD thesis, Massachusetts Institute of Technology, Dept. of Aeronautics and Astronautics, 2000.
- [12] Angelia Nedić and Asuman E. Ozdaglar. Approximate primal solutions and rate analysis for dual subgradient methods. *SIAM Journal on Optimization*, 19(4):1757–1780, December 2009.
- [13] Yuri Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course (Applied Optimization)*. Springer Netherlands, 1 edition, 2003.
- [14] J.B. Rawlings, D. Bonne, J.B. Jorgensen, A.N. Venkat, and S.B. Jorgensen. Unreachable setpoints in model predictive control. *IEEE Transactions on Automatic Control*, 53(9):2209 –2215, October 2008.
- [15] S. Richter, C.N. Jones, and M. Morari. Real-time input-constrained MPC using fast gradient methods. In *Proceedings of the 48th Conference on Decision and Control*, pages 7387 – 7393, Shanghai, China, December 2009.
- [16] S. Richter, M. Morari, and C.N. Jones. Towards computational complexity certification for constrained MPC based on Lagrange relaxation and the fast gradient method. In *Conference on Decision and Control (CDC)*, pages 5223–5229, Orlando, USA, December 2011.
- [17] Paul Tseng. On accelerated proximal gradient methods for convex-concave optimization. Submitted to *SIAM J. Optim.* Available: <http://www.math.washington.edu/~tseng/papers/apgm.pdf>, May 2008.