Automatica 50 (2014) 155–161

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

An active set solver for input-constrained robust receding horizon control☆

Johannes Buerger, Mark Cannon [1], Basil Kouvaritakis

Department of Engineering Science, University of Oxford, OX1 3PJ, UK

ARTICLE INFO

Article history:
Received 7 August 2012
Received in revised form
28 May 2013
Accepted 9 September 2013
Available online 30 October 2013

Keywords:
Robust control
Control of constrained systems
Dynamic programming
Min–max optimal control

ABSTRACT

An efficient optimization procedure is proposed for computing a receding horizon control law for linear systems with linearly constrained control inputs and additive disturbances. The procedure uses an active set approach to solve the dynamic programming problem associated with the min–max optimization of an $\mathcal{H}_\infty$ performance index. The active constraint set is determined at each sampling instant using first-order necessary conditions for optimality. The computational complexity of each iteration of the algorithm depends linearly on the prediction horizon length. We discuss convergence, closed loop stability and bounds on the disturbance l^2-gain in closed loop operation.

1. Introduction

The aim of robust control is to provide guarantees of stability and of performance with respect to a suitable measure, despite uncertainty in the model of the controlled system. Model Predictive Control (MPC) uses a receding horizon strategy to derive robust control laws by repeatedly solving a constrained optimization problem online, and consequently the approach is effective for systems with constraints and bounded disturbances.

Robust receding horizon control based on a worst-case optimization was proposed in Witsenhausen (1968). The approach employed a min–max optimization, which was subsequently adopted in Campo and Morari (1987) to derive an MPC law for linear systems with uncertain impulse response coefficients. In this strategy, and in the related work (Allwright & Papavasiliou, 1992), an open loop predicted future input sequence was used to minimize the worst-case predicted performance. It was argued in Lee and Yu (1997) that by optimizing instead over closed loop predicted input sequences, control laws with improved performance and larger regions of attraction could be obtained. However, unless a degree of optimality is sacrificed through the use of suboptimal controller parameterizations (such as, for example, those proposed in Goulart, Kerrigan, and Maciejowski (2006), Kothare, Balakrishnan, and Morari (1996) and Löfberg (2003)), strategies that involve a receding horizon optimization over predicted feedback policies generally require impractically large computational loads. For example Kerrigan and Maciejowski (2003) and Scokaert and Mayne (1998) apply a scenario-based approach to constrained linear systems with bounded additive uncertainty, which leads to an optimization problem in a number of variables which grows exponentially with the prediction horizon length.

Parametric solution methods aim to avoid the explosion in computational complexity of robust dynamic programming with horizon length by characterizing the solution of the receding horizon optimization problem offline, typically as a feedback law that is a piecewise affine function of the model state. In (Bemporad, Borrelli, & Morari, 2003; Diehl & Björnberg, 2004) this method was applied to linear systems with polytopic parametric uncertainty. However, whereas MPC typically solves an optimization problem for a single initial condition at each instant, this approach requires the solution at all points in state space, and moreover necessitates determining online which of a large number of polytopic regions contains the current state.

This paper extends the methodology developed in Cannon, Liao, and Kouvaritakis (2008), Ferreau, Bock, and Diehl (2008) and Best (1996) to the case of linear systems with bounded additive uncertainty and input constraints in order to derive a robust dynamic programming solver. An online active set method is described which avoids the need to compute the solution over the entire

☆ The material in this paper was partially presented at the 50th IEEE Conference on Decision and Control (CDC), December 12–15, 2011, Orlando, Florida, USA. This paper was recommended for publication in revised form by Associate Editor Martin Guay under the direction of Editor Frank Allgöwer.

E-mail addresses: johannes.buerger@bmw.de (J. Buerger), mark.cannon@eng.ox.ac.uk (M. Cannon), basil.kouvaritakis@eng.ox.ac.uk (B. Kouvaritakis).

[1] Tel.: +44 1865 273189; fax: +44 1865 273906.

0005-1098/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.automatica.2013.09.032

state space, and which forms the basis of an efficient line-search-based point location technique. The control law is optimal for a convex–concave min–max $\mathcal{H}_\infty$ performance index, which ensures closed loop stability and a specified l^2-disturbance gain bound. The algorithm's computational complexity per iteration grows only linearly with horizon length. We consider the case, which was presented in Buerger, Cannon, and Kouvaritakis (2011), of systems subject to constraints on control inputs alone; this paper provides further theoretical and numerical results and gives comparisons with max–min and open loop strategies, as well as numerical comparison with the suboptimal min–max strategy of Goulart, Kerrigan, and Alamo (2009).

2. Problem statement and notation

We consider linear discrete time systems with model

$$x_{t+1} = Ax_t + Bu_t + Dw_t, \quad t = 0, 1, \ldots \tag{1}$$

with state $x_t \in \mathbb{R}^{n_x}$, control input $u_t \in \mathbb{R}^{n_u}$ and disturbance input $w_t \in \mathbb{R}^{n_w}$ at time t. Here u_t and w_t are subject to constraints: $u_t \in \mathcal{U}$, $w_t \in \mathcal{W}$, and $\mathcal{U}$ and $\mathcal{W}$ are assumed to be convex polytopic sets defined by

$$\mathcal{U} \triangleq \left\{u \in \mathbb{R}^{n_u} : Fu \le \mathbf{1}\right\}$$

$$\mathcal{W} \triangleq \left\{w \in \mathbb{R}^{n_w} : Gw \le \mathbf{1}\right\}$$

for $F \in \mathbb{R}^{n_F \times n_u}$, $G \in \mathbb{R}^{n_G \times n_w}$, where $\mathbf{1} = [1 \ \cdots \ 1]^T$ denotes a vector of conformal dimensions.

We define the feedback law $u_N^*(x)$ as the solution to the following closed loop robust optimal control problem (Bemporad et al., 2003; Mayne, Raković, Vinter, & Kerrigan, 2006) over a finite horizon of N time-steps:

$$\left(u_m^*(x), w_m^*(x, u)\right) \triangleq \arg \min_{u \in \mathcal{U}} \max_{w \in \mathcal{W}} J_m(x, u, w) \tag{2a}$$

with J_m defined for $m = 1, 2, \ldots, N$ by

$$J_m(x, u, w) \triangleq \frac{1}{2}(\|x\|_Q^2 + \|u\|_R^2 - \gamma^2 \|w\|^2) + J_{m-1}^*(x^+) \tag{2b}$$

$$J_m^*(x) \triangleq J_m\big(x, u_m^*(x), w_m^*(x, u_m^*(x))\big) \tag{2c}$$

where $x^+ = Ax + Bu + Dw$, with the terminal cost:

$$J_0^*(x) \triangleq \frac{1}{2}\|x\|_P^2. \tag{2d}$$

Here R is a positive-definite matrix (denoted as $R \succ 0$), Q is a positive-semidefinite matrix ($Q \succeq 0$), $\|x\|_Q^2$ denotes $x^T Qx$, and the scalar γ is chosen (as discussed in Section 3.1) to be sufficiently large that (2a) defines a strictly convex–concave min–max problem. We make the assumption that P is chosen so that $\|x_0\|_P^2 = \sum_{t=0}^{\infty}(\|x_t\|_Q^2 + \|u_t\|_R^2 - \gamma^2\|w_t\|^2)$ with $u_t = u_\infty^f(x_t)$ and $w_t = w_\infty^f(x_t, u_t)$, where $u_\infty^f(\cdot)$, $w_\infty^f(\cdot)$ are the optimal solutions of (2a)–(2c) in the limit as $N \to \infty$ and in the absence of constraints $u \in \mathcal{U}$, $w \in \mathcal{W}$. In order to guarantee the existence of this solution we assume that (A, B) is controllable and $(Q^{1/2}, A)$ observable. Note that $u_\infty^f(\cdot)$, $w_\infty^f(\cdot)$ can be computed by solving a semidefinite programming problem (see e.g. Boyd, El Ghaoui, Feron, & Balakrishnan, 1994, for details).

The problem defined in (2a)–(2d) is formulated under the assumption that the disturbance w_t is unknown when the control input u_t is chosen at time t. Since the solutions $u_m^*(\cdot)$ and $w_m^*(\cdot)$ depend on x and on (x, u) respectively, (2a)–(2d) defines a closed loop optimal control problem (see e.g. Lee & Yu, 1997). The sequential nature of this min–max problem and the fact that the optimization is performed over arbitrary feedback laws $\{u_m^*(x), w_m^*(x, u),\ m = 1, \ldots, N\}$ imply that, unlike open-loop formulations of robust MPC (e.g. Campo & Morari, 1987), (2a)–(2d) cannot be solved exactly by a single quadratic program.

For given x_0, we denote the optimal state, input and disturbance sequences as $\mathbf{x}(x_0) = \{x_0, \ldots, x_N\}$, $\mathbf{u}(x_0) = \{u_0, \ldots, u_{N-1}\}$, $\mathbf{w}(x_0) = \{w_0, \ldots, w_{N-1}\}$, where, for $t = 0, \ldots, N-1$ we define $u_t = u_{N-t}^*(x_t)$, $w_t = w_{N-t}^*(x_t, u_t)$ and $x_{t+1} = Ax_t + Bu_t + Dw_t$.

As discussed in Section 4, a receding horizon control law is obtained by setting $u_t = u_N^*(x_t)$ at each time t.

3. Active set solution via Riccati recursion

This section describes a method of solving (2a)–(2d) in order to determine $u_N^*(x)$ for a given plant state, $x = x^p$. Therefore we aim at determining a local solution to the closed-loop formulation of problem (2). For a given active set, we use a Riccati recursion to solve the Karush–Kuhn–Tucker (KKT) conditions (Nocedal & Wright, 2006) providing first-order necessary optimality conditions for problem (2). The optimal control and disturbance inputs for the corresponding equality constrained problem are obtained as a sequence of affine state feedback functions. We give necessary and sufficient conditions for optimality of these policies with respect to problem (2). For the given active constraint set, our approach then determines state, control, disturbance and multiplier sequences as functions of the initial state x_0 using the system model (1). As in Cannon et al. (2008), we use a line-search through x_0-space to update the active set, and the process is repeated until $x_0 = x^p$. This line-search is based on homotopy of solutions to problem (2) and the solution can either be initialized using the unconstrained optimal control law with $x_0 = 0$ or warm started using the optimal solution for the plant state at the preceding time instant. Finally we discuss how the computation required by this approach depends on the problem size.

3.1. First order optimality conditions

Let λ_t and $\hat{\lambda}_t$ denote the Lagrange multipliers associated with the constraints $x_{t+1} = \hat{x}_{t+1} + Dw_t$ and $\hat{x}_{t+1} = Ax_t + Bu_t$, and let μ_t and η_t denote the Lagrange multipliers for the constraints $u_t \in \mathcal{U}$ and $w_t \in \mathcal{W}$ respectively.

Define Lagrangian functions for the stage-wise maximization and minimization subproblems recursively as follows:

$$\begin{aligned}\hat{H}_t(\hat{x}_{t+1}, w_t, \eta_t, \lambda_t, x_{t+1}, \ldots, x_N) &\triangleq -\frac{1}{2}\gamma^2\|w_t\|^2 + \eta_t^T(\mathbf{1} - Gw_t) \\ &\quad - \lambda_t^T(x_{t+1} - (\hat{x}_{t+1} + Dw_t)) \\ &\quad + H_{t+1}(x_{t+1}, u_{t+1}, \mu_{t+1}, \hat{\lambda}_{t+1}, \hat{x}_{t+2}, \ldots, x_N)\end{aligned}$$

for the maximization subproblem and

$$\begin{aligned}H_t(x_t, u_t, \mu_t, \hat{\lambda}_t, \hat{x}_{t+1}, \ldots, x_N) &\triangleq \frac{1}{2}\|x_t\|_Q^2 + \frac{1}{2}\|u_t\|_R^2 - \mu_t^T(\mathbf{1} - Fu_t) \\ &\quad - \hat{\lambda}_t^T(\hat{x}_{t+1} - (Ax_t + Bu_t)) + \hat{H}_t(\hat{x}_{t+1}, w_t, \eta_t, \lambda_t, x_{t+1}, \ldots, x_N)\end{aligned}$$

for the minimization subproblem, with terminal condition $H_N(x_N) = J_0^*(x_N) = \frac{1}{2}\|x_N\|_P^2$.

Lemma 1. *The solution of problem (2) at time-step t satisfies*

$$J_{N-t}^*(x_t) = \min\left\{\frac{1}{2}\|x_t\|_Q^2 + \frac{1}{2}\|u_t\|_R^2 + \hat{H}_t(\hat{x}_{t+1}, w_t, \eta_t, \lambda_t, x_{t+1}, \ldots, x_N)\right\} \tag{3}$$

where the minimization is over variables u_t, $\hat{x}_{t+1}$, and w_j, η_j, λ_j, x_{j+1} for $j = t, \ldots, N-1$ and u_j, μ_j, $\hat{\lambda}_j$, $\hat{x}_{j+1}$ for $j = t+1, \ldots, N-1$ and subject to the constraints: $\hat{x}_{t+1} = Ax_t + Bu_t$, $Fu_t \le \mathbf{1}$, and $\nabla_{w_j}\hat{H}_j = 0$, $\nabla_{x_{j+1}}\hat{H}_j = 0$ for $j = t, \ldots, N-1$, and $\nabla_{u_j}H_j = 0$, $\nabla_{\hat{x}_{j+1}}H_j = 0$ for $j = t+1, \ldots, N-1$. Furthermore the objective of the maximization in (2) at time-step t, defined by $\hat{J}_{N-t}^*(\hat{x}_{t+1})$