ELSEVIER

Journal of Process Control 10 (2000) 283–289

JOURNAL OF PROCESS CONTROL

www.elsevier.com/locate/jprocont

Linear programming and model predictive control

Christopher V. Rao, James B. Rawlings*

Department of Chemical Engineering, University of Wisconsin-Madison, Madison WI, 53706-1691, USA

Abstract

The practicality of model predictive control (MPC) is partially limited by the ability to solve optimization problems in real time. This requirement limits the viability of MPC as a control strategy for large scale processes. One strategy for improving the computational performance is to formulate MPC using a linear program. While the linear programming formulation seems appealing from a numerical standpoint, the controller does not necessarily yield good closed-loop performance. In this work, we explore MPC with an l_1 performance criterion. We demonstrate how the non-smoothness of the objective function may yield either dead-beat or idle control performance.

Keywords: Model predictive control; Linear programming; Optimization

1. Introduction

Model predictive control (MPC) is an optimization based strategy that uses a plant model to predict the effect of potential control action on the evolving state of the plant. At each time step, an open-loop optimal control problem is solved and the input profile is injected into the plant until a new measurement becomes available. The updated plant information is used to formulate and solve a new open-loop optimal control problem.

The MPC methodology is appealing to the practitioner, because input and state constraints are explicitly accounted for in the controller. A practical disadvantage is the computational cost, which tends to limit MPC applications to linear processes with relatively slow dynamics. For such problems, the optimal control problem to be solved at each stage of MPC is a convex program. The necessity to solve the optimization problem in real time is especially troublesome for large-scale processes. While efficient software exists for the solution of convex programs, significant improvements are obtained by exploiting the structure of the MPC subproblem.

Traditionally model predictive control has been formulated using a quadratic criterion. Part of the popularity of the quadratic criterion from a theoretical standpoint is due to its mathematical convenience. From a numerical standpoint, the quadratic criterion is popular, because the resulting optimization can be cast as a quadratic program. For the unconstrained case, the linear quadratic optimal control problem is solved efficiently using dynamic programming. This solution technique has the desirable property that the computational cost scales linearly in the horizon length N as opposed to cubically for the general least squares solution. While the addition of constraints negates the possibility of a general analytic solution to the optimal control problem, the quadratic program may be structured in an analogous manner to the unconstrained problem, yielding linear growth in the horizon length N. Approaches to structuring the optimal control problem with a linear quadratic objective utilizing sparse matrix methods are available in the literature [2,18,21].

Recently Dave and co-workers [7] have advocated the use of an l_1/l_∞ norm as a performance criterion for MPC. One motivation is that the resulting optimal control problem is cast as a linear program. The solution of a linear program is less computationally demanding than the corresponding solution of a quadratic program of the same size and complexity, so it may be preferable to formulate MPC as a linear program. The concept of using linear programming is not new and has been considered by many authors in optimal control (e.g. [16,22]) and in MPC (e.g. [1,3–6,10,12,14,17]). A review of some MPC research with non-quadratic objectives can be found in the paper by Garcia and co-workers [9]. The main theoretical objection to linear programming formulations is that analytic

* Corresponding author. Tel.: +1-608-262-1092; fax: +1-608-262-5434.

0959-1524/00/$ - see front matter
PII: S0959-1524(99)00034-7

solutions are generally unavailable due to the non-smoothness of the objective function. The non-smoothness is one of the prime reasons why the analysis of the stability for linear programming formulations has been lacking. Notable exceptions include the works of Keerthi and Gilbert [13], who use an endpoint constraint, Genceli and Nikolaou [10], who consider finite impulse response models, and Shamma and Xiong [20], who provide a numerical test whether a given horizon is sufficiently long to guarantee stability for unconstrained MPC.

In this paper we examine linear programming formulations of MPC. We begin our discussion by presenting in Section 2 a stabilizing formulation of MPC with a general l_p criterion. In Section 3 we analyze the qualitative properties of MPC with an l_1 criterion. Unlike MPC with a quadratic criterion, the choice of the tuning parameters for the l_1 formulation may result in appreciably different closed-loop performance. In particular, we demonstrate how the non-smoothness of the objective may yield either dead-beat or idle control performance.

2. Stabilizing MPC with l_p criterion

Consider the regulation following linear discrete-time representation of the plant

$$x_{k+1} = Ax_k + Bu_k, \quad k \geqslant 0, \tag{1a}$$

$$y_k = Cx_k \tag{1b}$$

where $x_k \in \mathbb{R}^n$, $u_k \in \mathbb{R}^m$, and $y_k \in \mathbb{R}^q$. We formulate the regulator as the feedback law $\eta(\hat{x}_j)$ that generates the sequence $\{u_k\}_{k=0}^{\infty}$, where $\eta(\hat{x}_j) \triangleq u_0$, that minimizes the infinite horizon objective function

$$\Phi(\hat{x}_k) = \sum_{k=0}^{\infty} \| \overline{R}u_k \| \hat{p} + \| \overline{Q}y_k \|_p, \tag{2}$$

subject to Eq. 1(a and b), the initial condition $x_0 = \hat{x}_j$, and the constraints

$$u_{\min} \leqslant Du_k \leqslant u_{\max}, \tag{3a}$$

$$-\Delta_u \leqslant \Delta u_k \leqslant \Delta_u, \tag{3b}$$

$$y_{\min} \leqslant y_k \leqslant y_{\max}, \tag{3c}$$

where

$$\| x \|_p \triangleq \left(\sum_{i=1}^{n} |x^{(i)}|^p \right)^{1/p}$$

and $x^{(i)}$ denotes the ith entry of the vector x. Common examples of l_p norms are the sum norm (l_1 norm)

$$\| x \|_1 \triangleq |x^{(1)}| + \cdots + |x(n)|$$

and the max norm (l_∞ norm)

$$\| x \|_\infty \triangleq \max\left\{ |x^{(1)}|, \cdots, |x^{(n)}| \right\}.$$

The vector $\hat{x}_j$ denotes the current state estimate of the plant at time index j. By suitably adjusting the origin, the regulator can account for target tracking and disturbance rejection [15]. We make the following assumptions: (a) (A, B) is stabilizable and (C, A) is detectable; (b) $\overline{Q}$ and $\overline{R}$ are diagonal matrices with positive elements; and (c) the origin $(u_k, x_k) = 0$ is contained within the interior of the feasible region Eq. 3(a–c). If a feasible solution exists, then the origin is an asymptotically stable fixed point for the feedback controller [13].

With the notable exceptions discussed in Section 3, analytic solutions to Eq. (2) are generally unavailable, because the l_p norm has a kink at the origin (see Fig. 1). To circumvent the computational barrier imposed by the infinite horizon calculation, we employ a stable finite horizon approximation. Our method is analogous to the technique employed by Rawlings and Muske [19] for a quadratic criterion. The basic strategy is to consider only a finite number of decision variables, so that the infinite horizon problem reduces to a finite dimensional mathematical program.

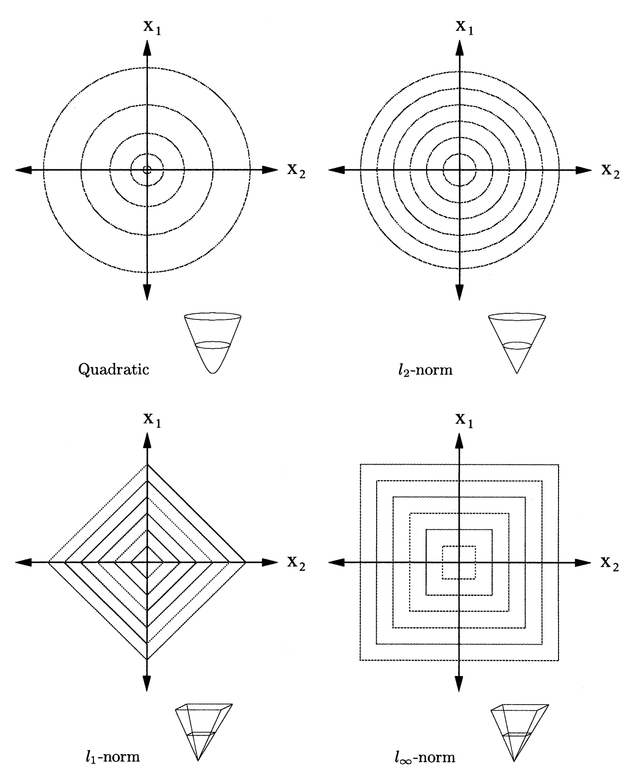

Fig. 1. Geometric interpretation of cost functions.