G Model
CACE-4824; No. of Pages 8
ARTICLE IN PRESS
Computers and Chemical Engineering xxx (2013) xxx–xxx
Contents lists available at ScienceDirect
Computers and Chemical Engineering
journal homepage: www.elsevier.com/locate/compchemeng

From robust model predictive control to stochastic optimal control and approximate dynamic programming: A perspective gained from a personal journey

Jay H. Lee*

Korea Advanced Institute of Science and Technology, Daejeon, Republic of Korea

ARTICLE INFO

Article history:
Received 28 May 2013
Received in revised form
15 September 2013
Accepted 13 October 2013
Available online xxx

Keywords:
Robust model predictive control
Min–max control
Approximate dynamic programming
Multistage decision-making under uncertainty

ABSTRACT

Developments in robust model predictive control are reviewed from a perspective gained through a personal involvement in the research area during the past two decades. Various min–max MPC formulations are discussed in the setting of optimizing the "worst-case" performance in closed loop. One of the insights gained is that the conventional open-loop formulation of MPC is fundamentally flawed to address optimal control of systems with uncertain parameters, though it can be tailored to give conservative solutions with robust stability guarantees for special classes of problems. Dynamic programming (DP) may be the only general framework for obtaining closed-loop optimal control solutions for such systems. Due to the "curse of dimensionality (COD)," however, exact solution of DP is seldom possible. Approximate dynamic programming (ADP), which attempts to overcome the COD, is discussed with potential extensions and future challenges.

1. Introduction

This paper reviews efforts in combining robust control with model predictive control. It is not meant to be a comprehensive survey but rather a review of personal sort that describes my own involvement in the research area and perspectives gained from it. It is thought to be appropriate for the current special issue, which is being put together in behalf of Prof. Manfred Morari, because the journey began at Caltech where I did my Ph.D. research under his supervision. It is not overboard to say that my research interest and style were shaped and molded largely during those times. Caltech during the mid and late 1980s was a vibrant place for those engaged in robust control research. The students who were fortunate to be there during that period enjoyed free access to some of the most recognized authorities on the topic. New theories and tools like the H_∞ control and the structured singular value (SSV) μ were taught and discussed as they were being developed. My home department was Chemical Engineering, so the group may not have participated in pioneering the *avant garde* theories but enjoyed the privilege of having a firsthand chance to learn and apply them to chemical systems.

Though the topic of robust control was dominating the Caltech's research activities, Prof. Morari's group was certainly aware of an important new development within the process control community called model predictive control (MPC). MPC was making a big splash among the industrial process control leaders like those at Shell Development, owing largely to its generality and ability to handle constraints, and was quickly becoming a hot topic among the academic researchers in the community. Given the group's deep engagement and expertise in robust control, it was natural for us to seek ways to incorporate the various concepts of robust control into MPC, in order to impart better robustness behavior to the MPC controllers. It was believed that the lack of explicit consideration by MPC was a major barrier to better and wider use of MPC by industry. This was not straightforward as it seemed because MPC, at least in its basic form embraced by the industrialists, is inherently a time-domain technique, whereas most of the robust control theories, *e.g.*, SSV and loop shaping, had been developed in the frequency domain. Some of the analysis efforts therefore were limited to unconstrained linear MPC, which could be translated into a linear time-invariant (LTI) control law and therefore yielded to frequency-domain analysis tools. In addition, an important milestone work that came out of these efforts was the min–max MPC formulation (Campo & Morari, 1987), which later would serve as a cornerstone for a large volume of research activities that followed in the next two decades.

By the time I left Caltech in 1991 to start my own academic career, I had developed a real interest in the problem of designing a model predictive controller for systems with uncertain parameters, e.g., those bounded within a polytope. Though my thesis research at

* Tel.: +82 42 350 3926; fax: +82 42 350 3910.
E-mail address: jayhlee@kaist.ac.kr

0098-1354/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.compchemeng.2013.10.014

Caltech started in the robust control area (applying the SSV theory to the sensor selection problem), my research had moved significantly towards MPC by the time of graduation (formulating MPC in state space and coupling it with a state estimator). This interest has led to a journey that lasted more than 20 years, almost the entire period of the author's academic career up to now. This paper describes the journey and attempts to provide some perspective on the problem, including its importance, difficulty, current status, and future challenges.

It is not easy to admit after working on it for two decades that the problem remains largely unsolved but it is the case for this problem. Along the way, a lot of insights have been gained and some partial solutions have come along. In fact, it was soon realized that this problem connects to the more general problem of stochastic optimal control and Markov Decision Process (MDP). It was also noticed that MPC may be inherently flawed to address the general class of the problem due to its open-loop optimal control formulation. Dynamic programming (DP) may be the only general method for it but it has its own problem known as the "curse of dimensionality (COD)." Therefore my works in this area for the second decade have attempted to connect the problems of robust MPC and stochastic optimal control with a new class of theories and techniques collectively known as approximate dynamic programming (ADP) (Bertsekas, 2012; Powell, 2011). Research efforts in ADP have mostly been driven by the computer science community and therefore translations and refinements as well as testing of these techniques became a central part of the my research efforts.

The remainder of the paper is organized as follows. In Section 2, a brief historical perspective into the related topics including robust MPC and ADP will be provided. In Section 3, a representative problem will be defined so that the various methods can be discussed in a more technical manner. Model form, uncertainty dynamics, and objective function will be mathematically defined and applicable methodologies will be categorized. In Section 4, several different robust MPC formulations will be given and their strengths and limitations will be discussed. In Section 5, the more general approach of DP and ADP will be brought in. We will focus on key concepts rather than specific methodologies in order to provide a sense for the motivation and the current state of the ADP development. Section 6 concludes the paper with some final perspectives. It should be apparent from reading the introduction that this paper carries somewhat of a personal tone, both in its contents and style. Though an excuse can be made for such a choice in a special issue paying a personal tribute, an apology is asked for nevertheless.

2. Historical perspective

One of the earliest works on robust MPC is a conference paper by Campo and Morari (1987), who adopted a finite impulse response model with upper and lower bounded response coefficients and formulated a min–max optimization over a finite time window to be solved at each time. This is very much in the same spirit as the traditional MPC used by industrialists at the time but the main difference was that the optimization minimized the "worst-case" control error with respect to a predefined impulse response parameter set. They showed that, when 1-norm/∞-norm is used as the error measure, the min–max optimization reduced to a linear program (LP). The formulation seemed sensible but the closed-loop results were found to be surprisingly unrobust at times. In addition, the closed-loop responses obtained seemed too slow due to the conservativeness of the "worst-case" objective. For this reason, the approach was not pursued further. In retrospect, the consequence of solving an open-loop optimal control problem at each time in the feedback control's context was not fully appreciated at the time. This is not surprising as even the conventional MPC based on a single fixed model had not been fully understood from the control-theoretic viewpoint at the time. For example, a formulation with guaranteed stability was lacking for the constrained case, mainly due to the nonlinear nature of the optimization-based controller. Some researchers chose to focus on the case of unconstrained linear MPC, which could be translated into a linear time-invariant (LTI) control system and therefore yielded to the existing stability and robustness analysis methods like the SSV and loopshaping (Bitmead, Gevers, & Wertz, 1990; Lee & Yu, 1994). Zarifiou on the other hand suggested to use the framework of nonlinear operator control theory and contraction mapping to analyze the stability and robustness of the industrial MPC algorithms (Zafiriou, 1990).

The min–max formulation got another look in the early 1990s. Lee and Yu (1997) took the min–max formulation by Campo and Morari and generalized to the case of a state space system of which matrices are parameterized by a vector belonging to a compact set, e.g., an ellipsoid or axis-aligned polyhedron. The parameters could also be time-invariant or time-varying (varying within the set in an arbitrary manner or according to some dynamic system description). They provided an important insight that in order for the min–max formulation to be closed-loop optimal in the worst-case error sense, one needs to solve a series of min–max optimizations, each one embedded into next one. It is unlikely that such an optimization can be solved on-line, and this consideration naturally led to a dynamic programming (DP) formulation. They showed that the resulting closed-loop behavior using the closed-loop optimal control formulation was entirely different from that using the original open-loop control formulation, in contrast to the conventional MPC case. For special cases of time-varying parameters, the open-loop min–max formulation can be made to be robustly stable by adding appropriate end constraints (Genceli & Nikolau, 1993; Lee & Yu, 1997) or using ∞-norm formulation with a sufficiently large horizon choice (Zheng & Morari, 1993). However, these formulations are significantly suboptimal in the closed-loop sense and oftentimes show very conservative behavior.

One notable approach that has been explored to circumvent the shortcomings of the open-loop optimal control formulation of MPC is to assume a specific parameterized form of feedback control law. Kothare, Balakrishnan, and Morari (1996) presented a min–max robust MPC technique for linear time-varying systems with a large class of plant uncertainty description, e.g., a convex hull of multiple linear models. Their approach represented a significant departure from the previous MPC approaches in that a linear state feedback law was assumed to be in force throughout the prediction horizon and the gain matrix elements, rather than the input values, were optimized at each time step. By optimizing a feedback policy rather than control inputs, the idea of closed-loop control was naturally incorporated into the optimization. They showed that the problem can be formulated as a convex optimization involving linear matrix inequalities and robust stability can be assured by stretching the prediction horizon to ∞ or by adding a terminal constraint. A limitation of the approach, of course, is that a special form of control law should be assumed *a priori*. For example, it may be possible to improve the performance significantly by relaxing the requirement of linear state feedback. In addition, even though the optimization is a convex semi-definite program, it is not straightforward to devise an algorithm that can solve large such problems reliably on-line.

In terms of the basic problem structure, the min–max MPC problem is closely related to the general problem of stochastic optimal control. For example, the topic of optimal control of linear systems with stochastic parameters had been studied for many decades prior to the min–max MPC research. Various ideas such as open-loop optimal feedback control (OLOFC), which is basically the MPC approach, appeared in the literature in the 1960s (Dreyfus, 1964, 1965). So did the idea of dual optimal control, which states that an optimal control law for a general stochastic system (e.g., linear

Please cite this article in press as: Lee, J.H. From robust model predictive control to stochastic optimal control and approximate dynamic programming: A perspective gained from a personal journey. *Computers and Chemical Engineering* (2013), http://dx.doi.org/10.1016/j.compchemeng.2013.10.014