Available online at www.sciencedirect.com
ScienceDirect
Automatica 44 (2008) 1642–1651
automatica
www.elsevier.com/locate/automatica

Applying a finite-horizon numerical optimization method to a periodic optimal control problem☆

Jeffrey Azzato[a], Jacek B. Krawczyk[b,*]

[a] Faculty of Science, Victoria University of Wellington, PO Box 600, Wellington, New Zealand
[b] Faculty of Commerce and Administration, Victoria University of Wellington, PO Box 600, Wellington, New Zealand

Received 15 May 2007; received in revised form 13 November 2007; accepted 27 December 2007
Available online 25 April 2008

Abstract

Computing a numerical solution to a periodic optimal control problem can be difficult, especially when the period is unknown. A method of approximating a solution to a stochastic optimal control problem using Markov chains was developed in [Krawczyk, J. B. (2001). A Markovian approximated solution to a portfolio management problem. *Information Technology for Economics and Management*, *1*, http://www.item.woiz.polsl.pl/issue/journal1.htm]. This paper describes the application of that method to a periodic optimal control problem formulated in [Gaitsgory, V. & Rossomakhine, S. (2006). Linear programming approach to deterministic long run average problems of optimal control. *SIAM Journal on Control and Optimization, 44*(6), 2006–2037]. As a result, approximately optimal feedback rules are computed that can control the system both on and off the optimal orbit.

Keywords: Periodic control; Numerical optimization; Dynamic systems; Markov chain approximation

1. Introduction

Gaitsgory and Rossomakhine (2006), develop a specialized linear programming technique for approximating solutions to long-run average optimal control problems. They illustrate this technique using a pair of un-discounted periodic optimal control problems, for which they obtain *feedback* rules for states belonging to the optimal orbit.

Periodic optimal control is especially applicable to situations requiring sensible long-term management of resources or externalities. For example, sustainable harvesting of an ecological resource is often best performed cyclically (He, Leung, & Stojanovic, 1995; Xiao, Cheng, & Qin, 2006). Similarly, the cost of controlling an established invasive species may be reduced by periodic use of countermeasures. Power generation and a variety of other industrial processes are either necessarily periodic or most efficient when operated periodically (Bausa & Tsatsaronis, 2001; Maurer, Büskens, & Feichtinger, 1998; van Noorden, Verduyn Lunel, & Bliek, 2003). The improvements in efficiency yielded by periodic operation appear increasingly valuable in the current world.

The uncertainty intrinsic to such problems requires optimal solutions which allow for stochastic perturbations. This paper proposes a practical feedback approach for dealing with these.

It was probably in the 1970s and early 1980s that applying the maximum principle became a standard approach to periodic optimal control (see Bittanti and Guardabassi (1985), Colonius (1985) and Noldus (1974)). A series of papers were published in which optimal *open-loop* solutions were computed and the methods yielding such solutions extensively discussed (see, for example, Han, Feichtinger, and Hartl (1994) and Maurer et al. (1998)). These solutions remain popular and can be found in papers from a variety of disciplines (see, for example, Bausa and Tsatsaronis (2001), He et al. (1995), Mombaur, Bock, Schlöder, and Longman (2005), Xiao et al. (2006), and Zhao (2004)).

☆ This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Editor Alain Haurie.

* Corresponding author. Tel.: +64 4 463 5352.
E-mail address: j.krawczyk@vuw.ac.nz (J.B. Krawczyk).

0005-1098/$ - see front matter
doi:10.1016/j.automatica.2007.12.022

As feedback solutions are more robust[1] than their open-loop counterparts (and thus of particular interest to control practitioners), research on the former has been growing. For example, Grüne and Semmler (2004)'s method is amenable to periodic controls. It produces feedback rules for states not necessarily confined to the optimal orbit. However, as the method is based on policy improvement, it requires that the optimization problem be *discounted*.[2]

Value iteration is a solution method applicable to finite-horizon (stochastic, discounted or otherwise) optimal control problems. In this paper, we use value iteration to compute a feedback solution to one of Gaitsgory and Rossomakhine (2006)'s periodic optimal control problems. This solution is used to obtain approximately optimal decision rules for controlling the system both on and off the optimal orbit. There is a very high degree of coincidence between the optimal performance indices and periods computed by our method and that of Gaitsgory and Rossomakhine (2006), and we conjecture that the former can be successfully applied to problems of this type.[3] Further pursuit of this thread of research (see also Gaitsgory and Rossomakhine (2006) and Grüne and Semmler (2004)) appears worthwhile.

The importance of the numerical methods proposed here and in Gaitsgory and Rossomakhine (2006) (also in Bausa and Tsatsaronis (2001), Han et al. (1994) and Grüne and Semmler (2004) and several other "discipline oriented" publications — see below) stems from the role that long-run average optimal control problems and their periodic solutions play in optimization. These problems arise from models finding application in such diverse areas as chemical reaction engineering (van Noorden et al., 2003), ecology management (He et al., 1995; Xiao et al., 2006), flight path planning (Speyer, 1996; Zhao, 2004), production planning (Maurer et al., 1998), robotics (Mombaur et al., 2005) and vibration damping (Kasturi & Dupont, 1998). The mathematical model discussed here and in Gaitsgory and Rossomakhine (2006) is representative of this class of problems.

This paper is organized as follows. In Section 2, we describe a useful class of long-run average control problems having periodic optimal solutions. A problem representative of this class is then defined and discussed in Section 3. In Section 5, we show how `SOCSol4L` – specialized software that converts an optimal control problem into a Markov decision chain (see Azzato and Krawczyk (2006) and Krawczyk (2001); for a brief introduction to `SOCSol4L` see the Appendix) – should be used to solve this periodic optimal control problem. The method used by `SOCSol4L` is outlined in Section 4, while Section 6 summarizes our findings.

[1] *E.g.*, to parameter uncertainty and shocks.

[2] The "discounted" method of Grüne and Semmler (2004) could also be used indirectly to solve un-discounted problems (see Camilli, Grüne, and Wirth (2001) and Grüne (1998)).

[3] *I.e.*, un-discounted and where a feedback solution is sought. In fact, our implementation of the method (see Section 4) allows for the possibility of a Gaussian diffusion term in the system dynamics.

2. Long-run average control problems

Consider a model having state variables $x_1, x_2 \in \mathbb{R}$ and horizon $[0, \Theta)$, where $\Theta > 0$ can be finite or infinite. Associate the first variable with a "generalized distance" and the second with the rate of change of the first ("speed"). Suppose that the second variable's rate of change is controlled through $v\colon [0, \Theta) \to [a_1, a_2]$ and depends on both the magnitude and rate of change of x_1. Then we have the system:

$$\left.\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -bx_1 - cx_2 - d + v \end{aligned}\right\} \qquad (1)$$

where $a_1 \geqslant 0$; $a_2, b, c > 0$ and the constant forcing factor d can be of either sign.

This system can be used to model a range of processes exhibiting periodic behaviour. For example, if x_1 is distance a vibration dampener (see Kasturi and Dupont (1998)) travels due to a shock, then the further from a reference point $x_1 = 0$ the slower the acceleration $\ddot{x}_1 = \dot{x}_2$. Similarly, biomass x_1 grows increasingly slowly when large due to scarcity of nutrients, and also, possibly, if growth x_2 has been too rapid. In macroeconomics, the amount of trade can undergo similar oscillations (Hicks, 1950). In each of these processes, timely control v can quickly help overcome undesirable values of x_1.

The system (1) can also describe some processes of a "qualitative" nature. Your business' production capacities x_1 decay with time, and may do so "faster" if you have just purchased new machinery. Only after x_1 has passed some critical value (but not beforehand) will you upgrade again. This reasoning applies to many managerial problems — including the prosaic task of when to repaint your house.

It is intuitively apparent that if control v is not only constrained, but also costly, and if the control's aim is to maintain a good average value of your capital (or a high amplitude of the dampener for good cushioning; or large or small but not medium biomass[4]), then you will not upgrade your capital every day "a little bit" but do it occasionally when it has decayed sufficiently. This results in the system's optimal evolution being cyclical and control v being of the *bang-bang* type.

The performance index introduced in Section 3.2 captures this behaviour, which reflects a desire to minimize resource utilization (u^2 term — similar to an LQ problem) balanced against a "greed" for the satisfaction yielded by large values of the generalized distance ($-y_1^2$ term — dissimilar to an LQ problem).

Analytically determining *when* control should be applied (and possibly its sign) may be difficult — hence the importance of the availability of numerical methods capable of delivering reliable results.

[4] Presumably, if biomass is large then harvest is cheap; if it is small you do not harvest at all. However, if the biomass is of medium size, harvesting is expensive.