

Annual Reviews in Control 28 (2004) 23-35

Annual Reviews in Control

www.elsevier.com/locate/arcontrol

Recent developments in stochastic MPC and sustainable development

B. Kouvaritakis*, M. Cannon, V. Tsachouridis

Department of Engineering Science, Oxford University, Parks Road, Oxford OX1 3PJ, UK Received 20 October 2003; received in revised form 20 January 2004; accepted 24 January 2004

Abstract

Despite the extensive literature that exists on predictive control and robustness to uncertainty, both multiplicative (e.g. parametric) and additive (e.g. exogenous), very little attention has been paid to the case of stochastic uncertainty. Yet this arises naturally in many control applications, for example when models are identified using least squares procedures. More generally, stochastic uncertainty is a salient feature in other key areas of human endeavour, such as sustainable development. Sustainability refers to the strategy of encouraging development at current time without compromising the potential for development in the future. Inevitably, modelling the effects of sustainable development policy over a horizon of say 30 years involves a very significant random element, which has to be taken into account when assessing the optimality of any proposed policy. Model Predictive Control (MPC) is ideally suited for generating constrained optimal solutions and as such would be an ideal tool for policy assessment. However, this calls first for suitable extensions to the stochastic case. The aim of this paper is to review some of the recent advances in this area, and to provide a pilot study that demonstrates the efficacy of stochastic predictive control as a tool for assessing policy in a sustainable development problem concerning allocation of public research and development budgets between alternative power generation technologies. This problem has been considered in earlier work, but only in the context of a single-shot, open-loop optimisation. Similarly, the consideration of stochastic predictive control methodologies has previously been restricted to general hypothetical control problems. The current paper brings together this body of work, proposes suitable extensions, and concludes with a closed-loop study of predictive control applied to a sustainable development policy assessment problem.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Stochastic MPC; Sustainable development; Closed-loop

1. Introduction

Model Predictive Control (MPC) solves, in a receding horizon manner, a series of open-loop optimisation problems, and thus provides tractable solutions to an infinite horizon constrained optimal control problem. As such it is ideally suited for use in a range of control applications, most of which are subject to constraints (both dynamic equality constraints and physical inequality constraints). However, MPC has enormous potential for application to a much wider class of problems in human endeavour where the aim is to maximise suitably defined benefit while keeping risk and cost within constraints. Many such examples can be found in applications of economics and finance, and also in a related area that is progressively gaining in importance: quantitative assessment of policy in sustainable development. Despite the grave concern about the effects that reckless development now will have on future generations (e.g. through the accumulation of atmospheric CO₂ emissions and the consequent damaging influence on climate), most attempts made so far at decision-making for sustainability are based on qualitative or "quasi-quantitative" assessments. What is needed is a purely quantitative assessment that allows a direct objective comparison of one policy against another. MPC has the potential to provide such an assessment but this development has been restricted due to its predominantly deterministic formulation. This is a major handicap because of the strongly stochastic nature of problems in sustainable development (and also in finance and many other problems in economics). Predicting how current development will, in say 30 years, affect the potential for development by our children or even generations beyond, inevitably introduces randomness which, given some statistical regularity, leads to a modelling exercise with a strong stochastic flavour.

It is somewhat curious to note that early development of MPC, such as self-tuning controllers for minimum variance (Astrom & Wittenmark, 1973) or generalized minimum variance (Clarke & Gawthrop, 1975), were specifically de-

^{*} Corresponding author. Tel.: +44-1865-273-100;

fax: +44-1865-273-906.

E-mail address: basil.kouvaritakis@eng.ox.ac.uk (B. Kouvaritakis).

 $^{1367\}text{-}5788/\$$ – see front matter © 2004 Elsevier Ltd. All rights reserved. doi:10.1016/j.arcontrol.2004.01.003

signed to handle random processes. However, optimisation with respect to a single statistic, be it the variance or the expected value (dealt with for example in a more recent paper by Batina, Stoorvogel, & Weiland, 2001), converts the problem to what amounts to a deterministic optimisation and therefore largely removes the stochastic element of the problem. To a significant degree this is avoided in the statistical framework of Whittle (1971), which later on was shown to lead to a convex second order cone program (e.g. Lobo, Vandenberghe, Boyd, & Lebret, 1998). This formulation considered linear inequality constraints with random coefficients and required that the constraints hold true with a probability greater than a given threshold. A closely related approach was adopted by Van Hessem and Bosgra (2002), who consider the problem of disturbance rejection in constrained MPC. The presence of constraints invalidates the assumptions of the separation principle, which would only allow consideration of the effects of past uncertainty during estimation and would ignore uncertainty altogether during prediction. The methodology of Van Hessem and Bosgra (2002) was based on confidence ellipsoids in the space of the random variables defining the stochastic disturbance. Statistical confidence has of course been used before to examine the effects of probabilistic uncertainty; for example, Cloud and Kouvaritakis (1986) considered the use of confidence ellipsoids in representing the effects of identification noise on model parameters. However, irrespective of whether the uncertainty is in the model parameters or future disturbances, it becomes necessary to use a projection transformation from the confidence ellipsoid to predicted outputs (in Van Hessem & Bosgra, 2002) or to the Nyquist plane (in Cloud & Kouvaritakis, 1986), and this transformation renders the results conservative to a significant degree.

To overcome these shortcomings, and also to consider the stochastic prediction and optimisation problem in its most general form, recent work (Cannon, Kouvaritakis, & Huang, 2004; Kouvaritakis, Cannon, & Tsachouridis, 2004) followed the methodology developed in the context of sustainable development in Kouvaritakis (2000). In the context of an integrated programme, this work developed a Tool for Integrated Policy Assessment (TIPA), which identified measurable indices (outputs) and common instruments (inputs). The approach therefore enabled quantitative comparisons of one policy against another; and it was possible to consider, in a non-conservative way, the probabilistic effects of uncertainty on both constraints and predicted cost. The context was the assessment of sustainable development policy in respect of allocation of public research and development budgets between alternative technologies for electrical power generation. Under the SAPIENT programme (Kouvaritakis, 2000), the starting point for this work was the derivation of a full econometric model, Prometheus, which modelled the effect of shocks (the discrete time equivalent of impulse changes to the common instruments superimposed on a baseline) applied to the budget allocation for each of 15 alternative technologies (listed in Appendix A) on eight measurable indicators (listed in Appendix B). Monte Carlo simulations, which used Prometheus to predict the effect of impacts on indicators at the end of a 30-year horizon, established the validity of a linear model in which the impact coefficients were random variables. Further investigation determined that the distribution of coefficients could be modelled as jointly normal (with a given mean and co-variance matrix); despite its convenience, this assumption was found not to have a significant effect on the optimal solution of the problem. It was then possible in Kouvaritakis (2002) to state the optimisation problem as a static constrained maximisation of the probability that the cumulative value (over the 30-year horizon) of an indicator, identified as the primary indicator, is greater than a given threshold. By analogy the constraints on the remaining indicators, referred to as secondary, were also probabilistic, and required that the probability of the cumulative value (over 30 years) of the secondary indicators being greater than given thresholds should be greater than or equal to pre-specified probabilities. Under hedging conditions (namely for probability values greater than 0.5, as opposed to smaller probabilities which correspond to gambling), it was conjectured that the above stochastic optimisation was convex and thus admitted a unique solution. This seminal work, although derived in the different context of sustainable development, lays the foundations for a meaningful formulation of stochastic MPC. It clearly needed to be extended in several ways: (i) the relationship considered between the shocks and the cumulative values of the primary and secondary indices (at the end of the 30-year horizon) was static; (ii) the optimisation was single-shot in the sense that it allowed only for a single budget allocation over the entire prediction horizon; (iii) although it was assumed that such policy would be implemented in a receding horizon fashion, this assumption was only implicit in that no concern was expressed for the effect of such implementation on closed-loop feasibility, stability and performance.

Appropriate extensions with respect to (i) and (ii) were proposed in Cannon et al. (2004), where attention was still focused on the cumulative values of indicators at the end of a 30-year horizon, but where Prometheus was used to identify a model which represents the dynamic effects of shocks on the indicators. It was then possible to account for several predicted future budget allocations within the open-loop optimisation framework. The aim of the work was not to carry out an exhaustive study of the full sustainable development with 15 inputs and 8 outputs, but rather to illustrate the benefits of dynamic models and the exploitation of the extra degrees of freedom introduced through the use of a multiple-shot optimisation. Accordingly Monte Carlo simulations performed on Prometheus were deployed to derive a 2×2 dynamic model linking the application of budget shocks on Wind Turbine and Combined Cycle Gas Turbine technologies to two indicators, one measuring cumulative CO2 emissions and the other Energy Costs in a particular world region. Simulations, repeated several thousands of times, were performed to identify the mean value and covariance matrix for the vec-

دريافت فورى 🛶 متن كامل مقاله

- امکان دانلود نسخه تمام متن مقالات انگلیسی
 امکان دانلود نسخه ترجمه شده مقالات
 پذیرش سفارش ترجمه تخصصی
 امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
 امکان دانلود رایگان ۲ صفحه اول هر مقاله
 امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
 دانلود فوری مقاله پس از پرداخت آنلاین
 پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات
- ISIArticles مرجع مقالات تخصصی ایران