Automatica 80 (2017) 340–348

Contents lists available at ScienceDirect

Automatica

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

Dual adaptive model predictive control ☆

Tor Aksel N. Heirung [a], B. Erik Ydstie [b], Bjarne Foss [c]

[a] *Department of Chemical and Biomolecular Engineering, University of California, Berkeley, CA 94720, United States*

[b] *Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA 15123, United States*

[c] *Department of Engineering Cybernetics, Norwegian University of Science and Technology, N-7491 Trondheim, Norway*

ARTICLE INFO

Article history:
Received 17 December 2015
Received in revised form
10 August 2016
Accepted 11 January 2017
Available online 30 March 2017

Keywords:
Dual control
Model predictive control
Adaptive control
Optimal control
Stochastic control
Probabilistic constraints
Parameter estimation
System identification
Excitation
Active learning

ABSTRACT

We present an adaptive dual model predictive controller (DMPC) that uses current and future parameter-estimation errors to minimize expected output error by optimally combining probing for uncertainty reduction with control of the nominal model. Our novel approach relies on orthonormal basis-function models to derive expressions for the predicted distributions for the output and unknown parameters, conditional on the future input sequence. Propagating the exact future statistics enables reformulating the original stochastic problem into a deterministic equivalent that illustrates the dual nature of the optimal control but is nonlinear and nonconvex. We further reformulate the nonlinear deterministic problem to pose an equivalent quadratically-constrained quadratic-programming (QCQP) problem that state-of-the-art algorithms can solve efficiently, providing the exact solution to the probabilistically constrained finite-horizon dual control problem. The adaptive DMPC solves this QCQP at each sampling time on a receding horizon; the adaptation is a result of updating the parameter estimates used by the DMPC to decide the control input. The paper demonstrates the application of DMPC to a single-input single-output (SISO) system with unknown parameters. In the simulation example, the parameter estimates converge quickly and the probing vanishes with increasing accuracy and precision of the estimates, improving the future control performance.

1. Introduction

This paper addresses the problem of optimal control and learning in the context of stochastic systems and models with stochastic parametric uncertainty and probabilistic constraints. Dual control, as introduced by Feldbaum (1961), is the optimal control under decision-relevant, reducible uncertainty. Dual control problems include the mechanisms for both control and learning in the formulation, and the solution optimally incorporates both aspects in the input to the process.

Using data to progressively reduce uncertainty is often framed as a learning process, in the control community primarily studied in the field of adaptive control. Most adaptive control algorithms are passively adaptive in the sense that learning takes place only as a side-effect of the control action. These controllers learn from normal operating data, which can contain very little information. Informally, a control that with nonzero probability affects not only the system state but also the uncertainty (specifically, error covariances or higher-order central moments) has a dual effect on the system; systems in which the control cannot affect this uncertainty are called neutral (see Bar-Shalom & Tse, 1974 for a rigorous definition). Note that dual effect and neutrality are properties of the system rather than the control algorithm. For systems in which the control has a dual effect, operating data can be made more informative by actively probing the process (Bar-Shalom, 1981), also known as excitation (Mareels et al., 1987), experimentation (Gevers & Ljung, 1986), exploration (Sutton & Barto, 1998), or active learning (Tse & Bar-Shalom, 1973). An actively adaptive controller is designed to improve the learning by accounting for the dual effect and increasing the amount of information generated. While active learning may fail to improve performance if the level of excitation is insufficient or excessive, the dual control is the optimal control with respect to expected system performance through endogenizing the dual effect in the problem formulation.

☆ We gratefully acknowledge the funding provided by the Center of Integrated Operations in the Petroleum Industry at NTNU, Trondheim, Norway (www.iocenter.no). The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Tongwen Chen under the direction of Editor Ian R. Petersen.

E-mail addresses: heirung@berkeley.edu (T.A.N. Heirung), ydstie@cmu.edu (B.E. Ydstie), bjarne.foss@ntnu.no (B. Foss).

http://dx.doi.org/10.1016/j.automatica.2017.01.030

Adaptive model predictive control (MPC) has received relatively little attention in the literature (Mayne, 2014). As with most adaptive control approaches, adaptive MPC may suffer from signals that are insufficiently exciting for the controller or model parameters to converge to appropriate values, which may lead to problems such as bursting (Anderson, 1985), pole-zero cancellations or inadmissible models (Mareels & Polderman, 1996), and turn-off (Wieslander & Wittenmark, 1971). One way of approaching this issue is to design a controller that actively explores the system by ensuring a certain level of excitation, either constantly or when needed. Larsson, Rojas, Bombois, and Hjalmarsson (2015), Marafioti, Bitmead, and Hovd (2014), and Shouche, Genceli, Vuthandam, and Nikolaou (1998) develop algorithms that ensure a prescribed amount of information or excitation is generated, with the potential disadvantage of suboptimal performance resulting from excessive excitation of the process.

Several proposed controllers generate excitation without a specific requirement. Rather, they include a function of information or uncertainty in the MPC cost function and optimize this function together with standard control objectives. Heirung, Foss, and Ydstie (2015a) propose and compare two such formulations that converge to a standard adaptive certainty-equivalence (CE; Åström & Wittenmark, 1995) MPC formulation as the uncertainty is reduced, and show that the excitation can improve closed-loop performance. Tanaskovic, Fagiano, Smith, and Morari (2014) suggest the addition of an exploring property as a possible extension of their adaptive MPC for finite-impulse-response (FIR) systems. Their approach involves modifying the nominally optimal input sequence by solving a second-stage optimization problem, the objective of which is decreasing the set of possible models at the next time step. Common to these approaches is that the excitation is a consequence of a heuristic modification of the controller, based on the assumption that the resulting excitation will improve overall performance. While this type of algorithm may work well in practice and improve performance over passive-learning approaches (see Heirung et al., 2015a), the excitation is not an implicit consequence of optimizing for performance, which is the case for dual control in the sense of Feldbaum. The algorithm type does, however, illustrate an important distinction: superimposing excitation on a nominally optimal control signal does not generally result in *optimal* performance, and the inputs are consequently not dual.

Feldbaum (1961) identified (stochastic) dynamic programming as an appropriate solution method for dual control problems in his pioneering papers on optimal integration of active learning with multistage decision-making under uncertainty. Åström and Helmersson (1986) solved a scalar dual control problem with one unknown parameter, but the "curse of dimensionality" prevents dynamic programming from being a viable solution approach for most dual control problems. This has motivated the use of modern approximate methods (Bayard & Schumitzky, 2010; Lee & Lee, 2009) that directly approximate the dynamic programming equations rather than the problem formulation.

In this article we derive an adaptive dual MPC (DMPC) for systems modeled with orthonormal basis functions with probabilistic parametric uncertainty and process noise. We formulate a stochastic optimal-control problem for minimizing expected performance cost, which involves the use of future information to evaluate the conditional expected future tracking error. This stochastic problem is transformed into an equivalent deterministic form that enables exact evaluation of both the objective function and the probabilistic constraints. The reformulation relies on the future decisions for propagation of the exact conditional distributions over the prediction horizon, which enables determination of the learning outcome of the decision sequence. Consequently, the learning is correctly rewarded in the control algorithm, avoiding heuristic additions to the cost function (as opposed to earlier work by Heirung et al., 2015a, e.g.). We transform the reformulated problem into a quadratically-constrained quadratic-programming (QCQP) problem that can be solved efficiently using state-of-the-art solvers. The proposed DMPC ensures that the system is sufficiently excited for accurate and precise parameter estimation but does not generate a persistently exciting input. Some results in this article are generalizations of ideas by Heirung, Ydstie, and Foss (2015b), a portion of which is given there without proof. Primarily, this paper considers a more general system type and includes probabilistic output constraints.

The act of exciting, or probing, a system for learning is often seen as conflicting with the control objective (see, e.g., Tse & Bar-Shalom, 1973), and a trade-off between control and probing is frequently discussed (Åström & Kumar, 2014, e.g.). However, based on the derivations in this article we argue that this is not a correct interpretation and show that excitation is an intrinsic part of the optimal control. That excitation is an inextricable part of the input in dual control means it cannot be derived or rewarded heuristically. Furthermore, the excitation and the nominal output error-minimization are not conflicting goals that can be traded off against each other; rather, they are inseparable components that together constitute the optimal control. Uncertainty reduction cannot be sacrificed for increased control performance.

This article is organized as follows: Section 2 provides the formulation of the stochastic control problem (P) and briefly reviews some necessary statistical background. The main contributions of the paper are in Section 3, where we state and prove the results necessary to reformulate the stochastic optimal-control problem as the equivalent deterministic problem (P′) and subsequently transform this formulation into the QCQP problem (P″). Section 4 contains the dual control algorithm, followed by a simulation example in Section 5. Section 6 concludes the paper with some thoughts for future work.

Notation: $\mathrm{E}[x \mid y]$ denotes the expected value of x, given y. $\Pr[A \mid y]$ is the probability of an event A, given y.

2. Problem formulation and background

This paper considers the output tracking problem for a class of systems of the form

$$\varphi(t+1) = A\varphi(t) + Bu(t) \tag{1a}$$

$$y(t) = \theta^\top \varphi(t) + v(t) \tag{1b}$$

where $\varphi(t) \in \mathbb{R}^{n_p}$ is a deterministic regression vector whose elements are functions of past control inputs (deterministic decision variables) u, and $A \in \mathbb{R}^{n_p \times n_p}$ and $B \in \mathbb{R}^{n_p}$ are known matrices determined by basis functions. The variable $y(t) \in \mathbb{R}^{n_p}$ is the process output and $v(t) \in \mathbb{R}^{n_p}$ is an additive, stationary process disturbance assumed to be a sequence of independent and identically distributed Gaussian random variables with zero mean and variance r. The vector $\theta \in \mathbb{R}^{n_p}$ contains the unknown parameters, $\theta = \left[\theta_1,\ \theta_2,\ \ldots,\ \theta_{n_p}\right]^\top$, where the set $\{\theta_j\}_{j=1}^{n_p}$ is drawn from a multivariate Gaussian distribution at initial time $t = t_0$ with mean $\hat{\theta}(t_0)$ and covariance $P(t_0)$. The model (1b) is often referred to as a linear regression.

The system (1) is a linear, time-invariant, single-input, single-output (SISO) system, and we assume that the pair (A, B) is controllable and stable; this formulation includes systems modeled by orthogonal basis functions (OBFS). The most well-known member of this model class is the FIR model; other common formulations include the Laguerre (Wahlberg, 1991) and Kautz (Wahlberg, 1994) models; see also Finn, Wahlberg, and Ydstie (1993) for a combination of the FIR and Laguerre structures. Heuberger, Van den Hof, and Wahlberg (2005) provide a comprehensive overview of modeling and identification with