Available online at www.sciencedirect.com

ScienceDirect

IFAC-PapersOnLine 49-7 (2016) 139–144

Dual Control and Information Gain in Controlling Uncertain Processes ⋆

H. C. La * A. Potschka * J. P. Schlöder * H. G. Bock *

* Interdisciplinary Center for Scientific Computing (IWR), INF 368, Im Neuenheimer Feld, 69120 Heidelberg, Germany. E-mails: Huu.Chuong.La@iwr.uni-heidelberg.de, potschka@iwr.uni-heidelberg.de, schloeder@iwr.uni-heidelberg.de, bock@iwr.uni-heidelberg.de

Abstract: In controlling uncertain processes, it is decisive to utilize information provided by measurements in order to estimate parameters and states. Nonlinear Model Predictive Control (NMPC) is a popular method to implement feedback control and deal with uncertainties. Conventional NMPC or nominal control, however, sometimes does not provide enough information for system estimation, leading to unsatisfactory performance. Dual control attempts to strike a balance between the two goals of enhancing system estimation and optimizing the nominal objective function. In this paper, we analyze the performance of these strategies through the interplay between the performance control task and the information gain task in connection with Optimal Experimental Design. Examples illustrate the conflict and agreement between the two tasks and explain why in some cases nominal control performs well. It is also observed that measurement noise provides excitation helping to improve the quality of estimates.

Keywords: Nonlinear Model Predictive Control, Dual Control, Parameter and State Estimation

1. INTRODUCTION

In solving control problems under uncertainties, control and estimation are often carried out in parallel. The information obtained by measurements is utilized to estimate parameters and states. With the use of Optimal Experimental Design (OED), the quality of the estimates can be enhanced before the realization of measurements is known. Hence it is reasonable to incorporate some sort of future information into the current control actions. Here emerges the concept of dual control as proposed by Feldbaum (1960-1961). The theoretical solution based on dynamic programming (Åström (1970), Bertsekas (1995)) is not practical in general due to the burden of computing conditional distributions and the *curse of dimensionality*. Recent studies explore the use of OED, e.g., Lucia and Paulen (2014) for robust Nonlinear Model Predictive Control (NMPC), Heirung et al. (2015a,b) for input-output systems, and La et al. (2016) using a sensitivity approach, demonstrating the advantages of using dual NMPC. A real-time implementation of such a strategy, especially for nonlinear systems, is numerically challenging, since often covariance matrices and possibly the derivatives of the covariance matrix with respect to controls are required. With increasing computational power and advanced algorithms, e.g., Körkel (2002), Bock et al. (2007), Kühl et al. (2011) this is can be efficiently tackled.

⋆ This paper was not presented at any IFAC meeting. Corresponding author H. C. La. This work was supported by the German Research Foundation (DFG) within the Heidelberg Graduate School for Mathematical and Computational Methods for the Sciences. Support by the EU through S. Engell's and H.G. Bock's ERC Advanced Investigator Grant MOBOCON (291 458) is gratefully acknowledged.

There are two tasks that the controller should take into account: the performance control task, which aims to feasibly drive the process in an optimal way specified by the objective function, and the information gain task, which aims to improve the accuracy of parameter and state estimates. The interplay between the two tasks needs to be investigated in order to make a suitable balance.

Conventional NMPC with estimation procedures has a feedback property in nature, hence helps the process react to uncertainties. Good performance of NMPC has been widely reported, and theoretical results for stability and robustness have been established, e.g., Findeisen et al. (2003); Diehl et al. (2005). However in some cases, the performance of conventional NMPC may be degraded, even infeasibility can occur. This is due to a lack of information for accurately estimating states and parameters. It is then necessary to use dual NMPC techniques. In this paper, we would like to present a rather broad view of controlling systems under uncertainties. We illustrate the interplay between performance control and information gain, giving examples in which they are conflicting or harmonious regarding particular parameters. It is also shown that measurement noise can play a role in exciting the control to estimate critical parameters, and that without noise, nominal NMPC would give a bad performance.

The paper is organized as follows. Section 2 presents several approaches to controlling uncertain systems in the framework of NMPC. Section 3 deals with a rocket car problem in which we show that the performance control task and the information gain task conflict when estimating some parameters but completely agree when estimating others. A realistic example of a tractor follows in Sec-

Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2016.07.230

tion 4 which demonstrates the superiority of dual NMPC over nominal NMPC when the control time is short. It also reveals the excitation contained in measurement noise.

2. CERTAINTY EQUIVALENCE, OPEN-LOOP FEEDBACK AND DUAL CONTROL

The model under consideration reads as

$$\begin{cases} x_{k+1} = f_k(x_k, u_k), & k = 0, 1, 2, ... \\ y_k = \eta_k(x_k) + v_k, & k = 1, 2, ... \end{cases} \tag{1}$$

where $x_k \in \mathbb{R}^n$ are states and $u_k \in \mathbb{U} \subset \mathbb{R}^{n_u}$ are controls. Often states x_k come from sampling a continuous-time system at discrete times $t_0, t_1, ...$ The $f_k : \mathbb{R}^n \times \mathbb{R}^{n_u} \to \mathbb{R}^n$ are state transition functions; the $\eta_k : \mathbb{R}^n \to \mathbb{R}^{n_y}$ are measurement functions; the $y_k \in \mathbb{R}^{n_y}$ are noisy measurements; and the v_k are random measurement noise. The initial state x_0 is a random vector with a known probability distribution. Note that some components of states x_k can stay constant over k, which represent constant parameters. We introduce notations that express information up to and including time $t \geq 1$

$$\mathcal{Y}^t = (x_0, u_0, y_1, u_1, y_2, ..., u_{t-1}, y_t), \quad \mathcal{Y}^0 = x_0.$$

Furthermore, we consider $\mathcal{Y}^t$ as random vectors and denote by y^t their realizations. As usual, the expectation of a random vector Z is denoted by $\mathbb{E}Z$. By an admissible control we understand a sequence $u = (u_0, u_1, ..., u_{N-1})$ such that u_t is a function of $\mathcal{Y}^t$, i.e., $u_t = \mu_t(\mathcal{Y}^t) \in \mathbb{U}$ for each t, $0 \leq t \leq N-1$ and the evaluation of (1) is feasible (cf. Åström (1970), Bertsekas (1995)). We aim to find in the set of all admissible control sequences a control sequence u^* that solves

$$\min_u \mathbb{E} J_N(x_0, u) \tag{2}$$

with $J_N(x_0, u) = F_N(x_N) + \sum_{k=0}^{N-1} L_k(x_k, u_k)$ subject to (1). Here $L_k : \mathbb{R}^n \times \mathbb{R}^{n_u} \to \mathbb{R}$ and $F_N : \mathbb{R}^n \to \mathbb{R}$ are cost functions. Define the value function

$$V(y^t, t) = \min_{u_t, ..., u_{N-1}} \mathbb{E}\Big[F_N(x_N) + \sum_{k=t}^{N-1} L_k(x_k, u_k)\Big|\mathcal{Y}^t = y^t\Big]$$

for $t = 1, 2, ..., N-1$. The dynamic programming principle holds (Bertsekas (1995)) in the form: For $t = N-1, ..., 0$,

$$V(y^t, t) = \min_{u_t} \mathbb{E}\Big[L_t(x_t, u_t) + V(\mathcal{Y}^{t+1}, t+1)\Big|\mathcal{Y}^t = y^t\Big] \tag{3}$$

with $V(y^N, N) = \mathbb{E}[F_N(x_N)|\mathcal{Y}^N = y^N]$. The optimal value for (2) is then $\mathbb{E}V(\mathcal{Y}^0, 0)$. Solving these equations is prohibitively expensive in general. Hence, approximation methods should be employed. The simplest one is to implement a *single* open loop control in which all random vectors x_0, y_k, v_k are replaced by their nominal values. Other sophisticated strategies include certainty equivalence control, open-loop feedback control (Bertsekas (1995)) and dual control (La et al. (2016)). They are presented in the following.

Certainty equivalence (CE) control. Set $t = 0$, $\mathcal{Y}^0 = \{x_0\}$.

Step 1: Compute $\bar{x}_t = \mathbb{E}[x_t|\mathcal{Y}^t]$.

Step 2: Find $u_t^*, u_{t+1}^*, ..., u_{N-1}^* \in \mathbb{R}^{n_u}$ that solve the deterministic problem $\min_{u_t, ..., u_{N-1}} J_N^{\mathrm{ce}}(\bar{x}_t, u)$ with

$$J_N^{\mathrm{ce}}(\bar{x}_t, u) = F_N(x_N) + \sum_{k=t}^{N-1} L_k(x_k, u_k)$$

in which all random vectors are replaced by their nominal values.

Step 3: Apply the control $\mu_t^*(\mathcal{Y}^t) = u_t^*$ to the system. If $t < N-1$, take measurements y_{t+1}; set $\mathcal{Y}^{t+1} = \{\mathcal{Y}^t, y_{t+1}\}$, $t = t+1$ and go to Step 1. If $t = N$, stop.

Note that CE ignores the uncertainties in the estimate $\bar{x}_t$.

Open-loop feedback. The open-loop feedback differs from the certainty equivalence control in that the objective function at each step is the expected value with respect to the random vectors involved, i.e., $\min_{u_t, ..., u_{N-1}} J_N^{\mathrm{of}}(\bar{x}_t, u)$ with

$$J_N^{\mathrm{of}}(\bar{x}_t, u) = \mathbb{E}\Big[F_N(x_N) + \sum_{k=t}^{N-1} L_k(x_k, u_k)\Big|\mathcal{Y}^t = y^t\Big].$$

This strategy takes care of the uncertainties in $\bar{x}_t$ but the computation of expected values is demanding.

Dual control. Choose $\alpha \geq 0$ as a weight and N_d as a length of the future horizon. Set $t = 0$, $\mathcal{Y}^0 = \{x_0\}$.

Step 1: Compute $\bar{x}_t = \mathbb{E}[x_t|\mathcal{Y}^t]$.

Step 2: Find $u_t^*, u_{t+1}^*, ..., u_{N-1}^* \in \mathbb{R}^{n_u}$ that solve the deterministic problem $\min_{u_t, ..., u_{N-1}} J_N^d(\bar{x}_t, u)$ where

$$J_N^d(\bar{x}_t, u) = J_N(\bar{x}_t, u) + \alpha\sqrt{\delta x_t^T C^t(\bar{x}_t, u)\delta x_t}. \tag{4}$$

Here $\delta x_t = \frac{\partial J_N^{\mathrm{ce}*}}{\partial \bar{x}_t}(\bar{x}_t)$ with $J_N^{\mathrm{ce}*}(\bar{x}_t)$ is the optimal value of $J_N^{\mathrm{ce}}(\bar{x}_t, u)$ and $C^t(\bar{x}_t, u) = (M^t)^{-1}$ with

$$M^t = \sum_{k=1}^{\max\{t+N_d, N\}} \left(\frac{\partial \eta}{\partial x_0}(x_k)\right)\left(\frac{\partial \eta}{\partial x_0}(x_k)\right)^T.$$

Step 3: Apply the control $\mu_t^*(\mathcal{Y}^t) = u_t^*$ to the system. If $t < N-1$, take measurements y_{t+1}; set $\mathcal{Y}^{t+1} = \{\mathcal{Y}^t, y_{t+1}\}$, $t = t+1$ and go to Step 1. Else stop.

We note that M^t is called the *Fisher information matrix* and $C^t(\cdot)$ is called the variance-covariance matrix. The second term in (4) can be interpreted as the variance of the objective function caused by uncertainty in the initial states.

The schemes above present NMPC for a fixed time horizon (batch NMPC). Similar procedures can be developed for receding horizon control and problems with constraints, e.g., Kirches et al. (2012). We also call CE nominal NMPC.

3. A ROCKET CAR: CONFLICT AND AGREEMENT OF INFORMATION GAIN AND PERFORMANCE CONTROL

Consider an object that can accelerate and brake. We aim to steer it on a straight line from point A to point B in minimal time. The model can be described as

$$\begin{cases} \dot{x}_1(t) = x_2(t), \\ \dot{x}_2(t) = a u_1(t) - b u_2(t), \end{cases} \tag{5}$$