Automatica 78 (2017) 297–308

Contents lists available at ScienceDirect

Automatica

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

Planning for optimal control and performance certification in nonlinear systems with controlled or uncontrolled switches[☆]

Lucian Buşoniu [a], Jamal Daafouz [b], Marcos Cesar Bragagnolo [c], Irinel-Constantin Morărescu [b]

[a] Automation Department, Technical University of Cluj-Napoca, Memorandumului 28, 400114 Cluj-Napoca, Romania

[b] Université de Lorraine, CRAN, UMR 7039 and CNRS, CRAN, UMR 7039, 2 av. Forêt de Haye, Vandœuvre-lès-Nancy, France

[c] Université de Reims Champagne-Ardenne, CReSTIC, BP 1039, 51687 Reims Cedex 2, France

ARTICLE INFO

Article history:
Received 15 January 2016
Received in revised form
1 August 2016
Accepted 9 December 2016

Keywords:
Switched systems
Optimal control
Planning
Nonlinear systems

ABSTRACT

We consider three problems for discrete-time switched systems with autonomous, general nonlinear modes. The first is optimal control of the switching rule so as to optimize the infinite-horizon discounted cost. The second and third problems occur when the switching rule is uncontrolled, and we seek either the worst-case cost when the rule is unknown, or respectively the expected cost when the rule is stochastic. We use optimistic planning (OP) algorithms that can solve general optimal control with discrete inputs such as switches. We extend the analysis of OP to provide certification (upper and lower) bounds on the optimal, worst-case, or expected costs, as well as to design switching sequences that achieve these bounds in the deterministic case. In this case, since a minimum dwell time between switching instants is often required, we introduce a new OP variant to handle this constraint, and analyze its convergence rate. We provide consistency and closed-loop performance guarantees for the sequences designed, and illustrate that the approach works well in simulations.

1. Introduction

Switched systems consist of a set of linear or nonlinear dynamics called modes, together with a rule for switching between these modes (Liberzon, 2003). They are employed to model real-world systems that are subject to known or unknown abrupt parameter changes such as faults (Du, Jiang, & Shi, 2015; Li, Gao, Shi, & Lam, 2016), including for instance embedded systems the automotive industry, aerospace, and energy management. This important class of hybrid systems is therefore heavily studied, with a main focus on stability and stabilization, see surveys (Lin & Antsaklis, 2009; Shorten, Wirth, Mason, Wulff, & King, 2007) and papers (Branicky, 1998; Daafouz, Riedinger, & Iung, 2002; Geromel & Colaneri, 2006; Lee & Dullerud, 2007; Pettersson & Lennartson, 1997). Performance optimization for switched systems has also been investigated, see e.g. the survey (Zhu & Antsaklis, 2015) and papers (Bengea & DeCarlo, 2005; Claeys, Daafouz, & Henrion, 2016; Riedinger, Iung, & Kratz, 2003; Seatzu, Corona, Giua, & Bemporad, 2006; Shaikh & Caines, 2007; Xu & Antsaklis, 2003). Hybrid versions of the Pontryagin Maximum Principle or dynamic programming have been proposed (Riedinger et al., 2003; Shaikh & Caines, 2007), with the drawback of lacking efficient numerical algorithms. Suboptimal solutions with guaranteed performance are given by Geromel, Deaecto, and Daafouz (2013) and Zhang and Abate (2012). The former efficiently represents the approximate value function using relaxations. The latter proves that the so-called min-switching strategies are consistent, i.e. that they improve performance with respect to non-switching strategies. Certification bounds (Geromel & Korogui, 2008) (lower and upper bounds on performance) are provided for linear switched systems with a dwell time assumption by Jungers and Daafouz (2013). Claeys et al. (2016) treat the problem by introducing modal occupation measures, which allow relaxation to a primal linear programming (LP) formulation. Overall, however, optimal control remains unsolved for general switched systems.

Motivated by this, our paper makes the following contributions. We propose an approach inspired from the field of planning in artificial intelligence, to either design switching sequences with near-optimal performance when switching is controllable, or to

[☆] The material in this paper was partially presented at the 54th IEEE Conference on Decision and Control, December 15–18, 2015, Osaka, Japan. This paper was recommended for publication in revised form by Associate Editor Huijun Gao under the direction of Editor Ian R. Petersen.

E-mail addresses: lucian@busoniu.net (L. Buşoniu), jamal.daafouz@univ-lorraine.fr (J. Daafouz), marcoscesarbragagnolo@gmail.com (M.C. Bragagnolo), constantin.morarescu@univ-lorraine.fr (I.-C. Morărescu).

http://dx.doi.org/10.1016/j.automatica.2016.12.027
0005-1098/

evaluate the performance when switching acts as a disturbance. We call the first problem PO, and the second either PW when the switching rule is unknown, in which case we estimate the worst-case performance; or PS when the switches evolve stochastically along a known Markov chain, in which case we evaluate the expected performance. Throughout, we consider a set of autonomous, general *nonlinear* modes, and a performance index consisting of the discounted infinite-horizon sum of general, nonquadratic stage costs. Optimistic planning (Buşoniu & Munos, 2012; Hren & Munos, 2008; Munos, 2014) is used to search the space of possible switching sequences. In all cases, our approach guarantees certification, lower and upper bounds on the (expected) performance.

When it makes sense to do so, namely in PO and PW, the method also designs a switching sequence that achieves the certification bounds. Since a minimum dwell time δ between switching instants must often be ensured, we introduce a new optimistic planner called OPδ that handles this constraint, and analyze its convergence rate. The analysis provides consistency and closed-loop performance guarantees for the sequences designed. Different from typical results, consistency shows improvement with respect to any suboptimal sequences, not only stationary ones. Finally, we illustrate the practical performance of the approach in simulations for several linear examples and a nonlinear one.

Compared to the optimal control methods reviewed above, the advantages of our approach include: a characterization of the certification bounds, a procedure to design a worst-case sequence, a design method with minimum dwell time, improved consistency results, and the ability to handle very general nonlinear modes. While a high computational complexity is unavoidable due to this generality, our analysis is focused precisely on characterizing the relation between computation and quality of the bounds.

An important remark is that much of the literature focuses on stability (Lin & Antsaklis, 2009; Shorten et al., 2007), whereas our aim is to provide near-optimality guarantees. Stability is a separate, difficult problem for discounted costs (Cardoso De Castro, Canudas De Wit, & Garin, 2012; Kiumarsi, Lewis, Modares, Karimpour, & Naghibi-Sistani, 2014; Postoyan, Buşoniu, Nešić, & Daafouz, in press). Nevertheless, in some cases our approach can exploit existing stability conditions: e.g. for some types of linear modes stability may be guaranteed under a dwell time constraint using the approach of Geromel and Colaneri (2006), in which case OPδ can enforce this constraint and thereby ensure stability.

The stochastic switching in PS leads to a Markov jump system, and there is a large body of literature dealing with such systems, again with a focus on linear modes (Boukas, 2006; Costa, Fragoso, & Marques, 2005), see e.g. Vargas, Costa, and do Val (2006) for optimal control. A recent nonlinear result is given by Zhong, He, Zhang, and Wang (2014), who analyze the stability properties of optimal mode inputs for Markov jump systems with nonlinear controlled modes. The practical implementation of Zhong et al. (2014) works for unknown mode dynamics, but without error guarantees, whereas all our methods provide tightly characterized bounds.

In the context of existing planning methods, solving PO and PW without dwell-time is a straightforward application of optimistic planning (Hren & Munos, 2008). In contrast, enforcing a minimum dwell-time requires deriving a novel algorithm and its accompanying analysis. Finally, solving PS can be seen as a special case of optimistic planning for stochastic systems (Buşoniu & Munos, 2012), but the nature of this special case allows us to derive a streamlined analysis. Compared to the preliminary version of this work (Buşoniu, Bragagnolo, Daafouz, & Morarescu, 2015), here we handle the new case of stochastic switching, provide consistency and closed-loop guarantees, and study two additional examples; in addition to including more technical discussion at several points in the paper.

Next, Section 2 formalizes the problem and Section 3 gives the necessary background. The approach is described in Section 4 for the optimal and worst-case problems PO and PW, and in Section 5 for the stochastic switching problem PS. Section 6 evaluates the planners in simulation examples of all these problems. Section 7 concludes.

List of symbols and notations

x, X, σ, S	State, state space, mode, set of modes
M	Number of modes
f_σ, p	Dynamics in mode σ, mode probabilities
$d, \boldsymbol{\sigma}_d$	Depth, mode sequence of length/depth d
γ, g, G	Discount factor, stage cost, cost bound
$J; \underline{J}, \bar{J}, \tilde{J}$	Cost; optimal, worst-case, expected cost
$\rho, v, \tilde{v}$	Reward function, value, expected value
r	Reward value
n	Computation budget
$\mathcal{T}, \mathcal{T}^*, \mathcal{L}(\mathcal{T})$	Tree, near-optimal tree, leaves of $\mathcal{T}$
l, b	Lower, upper bound on deterministic value
L, B	Lower, upper bound on expected value
l^*, b^*, L^*, B^*	Best bounds found by the algorithms
d^*	Largest depth found by the algorithms
ε	Near-optimality or sub-optimality
κ	Branching factor of near-optimal tree
K	Complexity of dwell-time problem
β	Complexity of stochastic problem
δ, Δ	Minimum dwell time, dwell time
e, λ	Leaf contribution, contribution cutoff
C, a, b, c	Constants
$\underline{\cdot}, \bar{\cdot}$	Quantity $\cdot$ in optimal, worst-case problem
$\cdot_\delta$	Quantity $\cdot$ for minimum dwell-time δ
$O(\cdot), \Omega(\cdot)$	Bounded above, below by $\cdot$ up to constants
$\tilde{O}(\cdot)$	Bounded above by $\cdot$ up to logarithmic terms
$[\cdot, \cdot]$	Concatenation of two mode sequences

2. Problem statement

Consider a discrete-time nonlinear switched system with states $x \in X$. The system can be at each step k in one of M modes $\sigma \in S = \{\sigma^1, \ldots, \sigma^M\}$, where each mode is autonomous:

$$x_{k+1} = f_{\sigma_k}(x_k). \tag{1}$$

The dwell time is defined as the number of steps during which the mode remains unchanged after a switch. A function $g(x_k, \sigma_k)$ assigns a numerical stage cost to each state–mode pair, e.g. quadratic in x_k up to saturation limits, see Example 1. Given a fixed initial state x_0, define an infinitely-long switching sequence $\boldsymbol{\sigma}_\infty = (\sigma_0, \sigma_1, \ldots)$ and the infinite-horizon discounted cost of this sequence:

$$J(\boldsymbol{\sigma}_\infty) = \sum_{k=0}^{\infty} \gamma^k g(x_k, \sigma_k) \tag{2}$$

where $\gamma \in (0, 1)$ is the discount factor and $x_{k+1} = f_{\sigma_k}(x_k)$. The dynamics f can be very general and a closed-form mathematical expression may not be available for them; the only requirement is that f can be simulated numerically.

To start with, we define two different problems:

PO. **Optimal control**: Find the optimal value $\underline{J} = \inf_{\boldsymbol{\sigma}_\infty} J(\boldsymbol{\sigma}_\infty)$ and a corresponding switching sequence that achieves it.