Automatica 70 (2016) 295–302

Contents lists available at ScienceDirect

Automatica

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

Brief paper

A jammer's perspective of reachability and LQ optimal control[☆]

Sukumar Srikant, Debasish Chatterjee[1]

Systems & Control Engineering, Indian Institute of Technology Bombay, Powai, Mumbai 400076, India

ARTICLE INFO

Article history:
Received 1 September 2015
Received in revised form
16 December 2015
Accepted 3 March 2016
Available online 11 May 2016

Keywords:
Sparse control
L_0-seminorm
Optimal control
Adaptive control

ABSTRACT

This article treats two problems dealing with control of linear systems in the presence of a jammer that can sporadically turn off the control signal. The first problem treats the standard reachability problem, and the second treats the standard linear quadratic regulator problem under the above class of jamming signals. We provide necessary and sufficient conditions for optimality based on a nonsmooth Pontryagin maximum principle.

1. Introduction

Given a controllable linear system

$$\dot{x}(t) = Ax(t) + Bu_1(t)$$

with $x(t) \in \mathbb{R}^d$ and $u_1(t) \in \mathbb{R}^m$,[2] we let a jammer corrupt the control u_1 with a signal $t \longmapsto u_2(t) \in \{0, 1\}$ that enters multiplicatively, and that can sporadically be "turned off", i.e., set to 0. The effect, therefore, of u_2 turning off is that the control u_1 is deactivated simultaneously, and the system evolves in open-loop. The signal u_2 provides a standard model for denial-of-service attacks for control systems in which the controller communicates with the plant over a network, and such models have been extensively studied in the context of cyberphysical systems; see, e.g., Raymond and Midkiff (2008) and the references therein. In this setting we ask whether it is possible to construct a control $t \longmapsto u_1(t)$ to execute the transfer of states of the resulting system

$$\dot{x}(t) = Ax(t) + Bu_1(t)u_2(t) \tag{1}$$

from given initial to given final states. Or, for instance, whether it is possible to stabilize the resulting system (1) to the origin by suitably designing the control u_1. Since both these problems are trivially impossible to solve if the jammer turns the signal u_2 'off' entirely, to ensure a well-defined problem, in the adaptive control literature typically a persistence of excitation condition, such as, there exist $T, \rho > 0$ such that for all t we have $\frac{1}{T}\int_t^{t+T} u_2(s)\,\mathrm{d}s \geqslant \rho$, is imposed on u_2. Very little, however, is known about either reachability or stabilizability of (1) under the above persistence of excitation condition. In particular, the problem of designing a state feedback $u_1(t) := K(t)x(t)$ such that the closed-loop system is asymptotically stable under the preceding persistence of excitation condition, is open, with partial solutions reported in Mazanti, Chitour, and Sigalotti (2013), Srikant and Akella (2009).

In this article, we study two problems concerning the control system (1). In the first problem, we turn the above-mentioned reachability question around and examine the limits of favourable conditions for the jammer. We ask the question: how long does the jamming signal u_2 need to be set to 'on' or 1 for the aforementioned reachability problem to be solvable? To wit, we are interested in the limiting condition such that if the jamming signal u_2 is set to 'off' or 0 for any longer time, then the standard reachability problem for (1) under the control u_1 would cease to be feasible. More precisely, we study the optimal control problem: given initial time $\bar{t}$ and final time $\hat{t} > \bar{t}$,

$$\begin{aligned} &\underset{u_1,u_2}{\text{minimize}} \quad \|u_2\|_{L_0([\bar{t},\hat{t}])} \\ &\text{subject to} \quad \begin{cases} \dot{z}(t) = Az(t) + Bu_1(t)u_2(t) \quad \text{for a.e. } t \in [\bar{t}, \hat{t}], \\ z(\bar{t}) = \bar{z} \in \mathbb{R}^d, \qquad z(\hat{t}) = \hat{z} \in \mathbb{R}^d, \\ u_1 : [\bar{t}, \hat{t}] \longrightarrow \mathbb{U} \subset \mathbb{R}^m \text{ compact}, \\ u_2 : [\bar{t}, \hat{t}] \longrightarrow \{0, 1\}, \\ u_1, u_2 \text{ Lebesgue measurable.} \end{cases} \end{aligned} \tag{2}$$

☆ S. Srikant was supported in part by the grant 12IRCCSG007 from IRCC, IIT Bombay. D. Chatterjee was supported in part by the grant 12IRCCSG005 from IRCC, IIT Bombay. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Akira Kojima under the direction of Editor Ian R. Petersen.

E-mail addresses: srikant.sukumar@iitb.ac.in (S. Srikant), dchatter@iitb.ac.in (D. Chatterjee).

1 Tel.: +91 22 2576 7879; fax: +91 22 2572 0057.

2 By controllability here we mean that the rank of the matrix $(B \quad AB \quad \cdots \quad A^{d-1}B)$ is equal to d.

http://dx.doi.org/10.1016/j.automatica.2016.03.026
0005-1098/

Here the cost function is the L_0-seminorm of the control u_2, defined to be the Lebesgue measure of the set of times at which the control is non-zero, i.e.,

$$\|u_2\|_{L_0([\bar{t},\hat{t}])} := \mathrm{Leb}\Big(\big\{s \in [\bar{t},\hat{t}] \,\big|\, u_2(s) \neq 0\big\}\Big).$$

We assume that the time difference $\hat{t} - \bar{t}$ is larger than the minimum time required to execute the transfer of the states from $\bar{z}$ to $\hat{z}$ in order to have a well-defined problem, and in addition assume that $0 \in \mathbb{R}^m$ is contained in the interior of $\mathbb{U}$. Notice that while the control u_1 tries to execute the desired manoeuvre, the control u_2 tries to switch to 'on' for the least length of time to enable execution of the aforementioned manoeuvre. We provide necessary conditions for these reachability manoeuvres and in addition provide conditions for optimality in (2).

The second problem that we study in this article is that of the performance of the linear quadratic regulator with respect to the control u_1 in the presence of the jammer u_2. We ask the question: How good is the performance of the standard linear quadratic regulator when the jammer corrupts the u_1 signal by turning it 'off' sporadically? To be precise, given symmetric and non-negative definite matrices $Q_f, Q \in \mathbb{R}^{d\times d}$ and a symmetric and positive definite matrix $R \in \mathbb{R}^{m\times m}$, initial time $\bar{t}$ and final time $\hat{t} > \bar{t}$, we study the following optimal control problem:

$$\begin{aligned}
&\underset{u_1,u_2}{\text{minimize}} && \gamma\,\|u_2\|_{L_0([\bar{t},\hat{t}])} + \tfrac{1}{2}\big\langle z(\hat{t}), Q_f z(\hat{t})\big\rangle \\
& && + \tfrac{1}{2}\int_{\bar{t}}^{\hat{t}} \Big(\langle z(t), Qz(t)\rangle + \langle u_1(t), Ru_1(t)\rangle\Big)\,\mathrm{d}t \\
&\text{subject to} && \begin{cases} \dot{z}(t) = Az(t) + Bu_1(t)u_2(t) \quad \text{for a.e. } t \in [\bar{t},\hat{t}], \\ z(\bar{t}) = \bar{z} \in \mathbb{R}^d, \\ u_1 : [\bar{t},\hat{t}] \longrightarrow \mathbb{R}^m, \\ u_2 : [\bar{t},\hat{t}] \longrightarrow \{0,1\}, \\ u_1, u_2 \text{ Lebesgue measurable,} \end{cases}
\end{aligned} \tag{3}$$

where $\gamma > 0$ is a fixed constant. If u_2 is set to 'off' for the entire duration $[\bar{t},\hat{t}]$, the cost accrued by the quadratic terms corresponding to an $L_2([\bar{t},\hat{t}])$ cost involving the states z and the control u_1 will be high. If u_2 is set to 'on' for the entire duration $[\bar{t},\hat{t}]$, the cost corresponding to $\|u_2\|_{L_0([\bar{t},\hat{t}])}$ will be high. Any solution to the optimal control problem (3) strikes a balance between the two costs: $L_2([\bar{t},\hat{t}])$-costs with respect to u_1 and the states, and the $L_0([\bar{t},\hat{t}])$-cost with respect to u_2. As in the case of (2), we provide necessary conditions for solutions to (3), and in addition provide sufficient conditions for optimality in (3).

It turns out that the optimal control u_1^* corresponding to the optimal control problem (2) is the sparsest control that achieves the steering of the states from $\bar{z}$ to $\hat{z}$ within the allotted time $\hat{t} - \bar{t}$—see Remark 4. The optimal control problem (3) is closely related to the "sparse quadratic regulator" problem treated in Jovanović and Lin (2013); see Remark 8. While the authors of Jovanović and Lin (2013) approached the optimal control problem using approximate methods via L_1 and total variation relaxations, it is possible to tackle the problem directly without any approximations, as we demonstrate in Remark 8. Sparse controls are increasingly becoming popular in the control community with pioneering contributions from Bahavarnia (2015), Fardad, Lin, and Jovanović (2014), Ikeda and Nagahara (2014), Jovanović and Lin (2013), Lin (2013), Lin, Fardad, and Jovanović (2011), Nagahara (2016), Nagahara, Quevedo, and Nešić (2014), Polyak (2014) and Polyak, Khlebnikov, and Shcherbakov (2013). Two distinct threads have emerged in this context: one, dealing with the design of sparse control gains, as in Bahavarnia (2015), Polyak et al. (2013), Polyak (2014), and two, dealing with the design of sparsest control maps as functions of time, as evidenced in the articles (Ikeda & Nagahara, 2014; Jovanović & Lin, 2013; Nagahara, 2016; Nagahara et al., 2014). With respect to Bahavarnia (2015), Polyak et al. (2013), Polyak (2014) our work differs in the sense that we do not design sparse feedback gains, but are interested in the design of sparse control maps that attain certain control objectives. The articles (Ikeda & Nagahara, 2014; Jovanović & Lin, 2013; Nagahara, 2016; Nagahara et al., 2014) deal with L_0-optimal control problems, but none of them treat the precise conditions for L_0-optimality, preferring instead to approximate sparse solutions with the aid of L_1-regularized optimal control problems. To the best of our knowledge, this is the first time that the two optimal control problems (2) and (3) are being studied.

Observe that both the optimal control problems (2) and (3) involve discontinuous instantaneous cost functions, and are consequently difficult to solve. We employ a nonsmooth version of the Pontryagin maximum principle to solve these two problems and study the nature of their solutions. Insofar as the existence of optimal controls is concerned, once again, the discontinuous nature of the instantaneous cost functions lends a nonstandard flavour to the above two problems. We derive our sufficient conditions for optimality with the aid of what is known as an inductive technique. These results are presented in Section 2. We provide detailed numerical experiments in Section 3 and conclude in Section 4.

Our notations are standard; in particular, for a set S we let $\mathbb{1}_S(\cdot)$ denote the standard indicator/characteristic function defined by $\mathbb{1}_S(z) = 1$ if $z \in S$ and 0 otherwise, and we denote by $\langle v, w\rangle = v^\top w$ the standard inner product on Euclidean spaces.

2. Main results

We apply the nonsmooth maximum principle Clarke (2013, Theorem 22.26) to the optimal control problems (2) and (3), for which we first adapt the aforementioned maximum principle from Clarke (2013) to our setting, and refer the reader to Clarke (2013) for related notations, definitions, and generalizations:

Theorem 1. *Let $-\infty < \bar{t} < \hat{t} < +\infty$, and let $\mathbb{U} \subset \mathbb{R}^m$ denote a Borel measurable set. Let a lower semicontinuous instantaneous cost function $\mathbb{R}^d \times \mathbb{U} \ni (\xi,\mu) \longmapsto \Lambda(\xi,\mu) \in \mathbb{R}$, with Λ continuously differentiable in ξ for every fixed μ,[3] and a continuously differentiable terminal cost function $\ell : \mathbb{R}^d \times \mathbb{R}^d \longrightarrow \mathbb{R}$ be given. Consider the optimal control problem*

$$\begin{aligned}
&\underset{u}{\text{minimize}} && \ell\big(x(\bar{t}), x(\hat{t})\big) + \int_{\bar{t}}^{\hat{t}} \Lambda\big(x(t), u(t)\big)\,\mathrm{d}t \\
&\text{subject to} && \begin{cases} \dot{x}(t) = f\big(x(t), u(t)\big) \quad \text{for a.e. } t \in [\bar{t},\hat{t}], \\ u(t) \in \mathbb{U} \quad \text{for a.e. } t \in [\bar{t},\hat{t}], \\ u \text{ Lebesgue measurable,} \\ \big(x(\bar{t}), x(\hat{t})\big) \in E \subset \mathbb{R}^d \times \mathbb{R}^d, \end{cases}
\end{aligned} \tag{4}$$

where $f : \mathbb{R}^d \times \mathbb{R}^m \longrightarrow \mathbb{R}^d$ is continuously differentiable, and E is a closed set. For a real number η, we define the Hamiltonian *H^η by*

$$H^\eta(x,u,p) = \langle p, f(x,u)\rangle - \eta\Lambda(x,u).$$

If $[\bar{t},\hat{t}] \ni t \longmapsto \big(x^*(t), u^*(t)\big)$ is a local minimizer of (4), then there exist an absolutely continuous map $p : [\bar{t},\hat{t}] \longrightarrow \mathbb{R}^d$ together with a scalar η equal to 0 or 1 satisfying the nontriviality condition

$$\big(\eta, p(t)\big) \neq 0 \quad \text{for all } t \in [\bar{t},\hat{t}], \tag{5}$$

the transversality condition

$$\big(p(\bar{t}), -p(\hat{t})\big) \in \eta\partial_x\ell\big(x^*(\bar{t}), x^*(\hat{t})\big) + N_E^L\big(x^*(\bar{t}), x^*(\hat{t})\big), \tag{6}$$

[3] Recall that a map $\varphi : X \longrightarrow \mathbb{R}$ from a topological space X into the real numbers is said to be lower semicontinuous if for every $c \in \mathbb{R}$ the set $\{z \in X \mid \varphi(z) \leqslant c\}$ is closed.