Systems & Control Letters 96 (2016) 124–131

Contents lists available at ScienceDirect

Systems & Control Letters

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

Networked control of nonlinear systems under Denial-of-Service☆

C. De Persis *, P. Tesi

Engineering and Technology Institute Groningen (ENTEG) and the Jan Willems Center for Systems and Control, Faculty of Mathematics and Natural Sciences, University of Groningen, The Netherlands

ARTICLE INFO

Article history:
Received 26 April 2015
Received in revised form
4 April 2016
Accepted 28 July 2016

Keywords:
Cyber-physical systems
Networked control systems
Nonlinear control systems
Jamming
Stability analysis

ABSTRACT

We investigate the analysis and design of a control strategy for nonlinear systems under Denial-of-Service attacks. Based on an ISS-Lyapunov function analysis, we provide a characterization of the maximal percentage of time that feedback information can be lost without resulting in instability of the system. Motivated by the presence of a digital channel we consider event-based controllers for which a minimal inter-sampling time is explicitly characterized.

1. Introduction

Motivated by interest in the analysis and control of critical infrastructures such as power networks, supply chains and transportation systems, recent years have witnessed increasing research interests in large-scale engineered systems. To achieve the prescribed control goal, these systems require exchange of information that often occurs in digital form. In turn this has triggered interest in control over communication channels. One of the topics that has stimulated broad interest is the so-called event-based control [1] in which sampling times are designed in real-time with the ultimate goal of saving communication resources while still guaranteeing the control goal. Event-based control has found fertile ground also in the area of cooperative control; *e.g.*, see [2,3].

A natural research question raises when dealing with control over a communication channel: whether or not stability properties and performance are preserved in the presence of loss of feedback information. This loss of information could be due to not only malfunctioning but also malicious actions by an adversarial entity [4,5]. In the latter case, the assumption on the kind of information loss should be kept to a minimum since intelligent adversaries might not follow e.g. any statistical pattern. This aspect is in contrast with other work where the loss of information is mainly due to the unreliability of the communication channel [6].

Several contributions to the topic of stability/stabilization in the presence of adversarial entities have been reported in the last few years, with main emphasis on the so-called *Denial-of-Service* (DoS), a class of attack strategies primarily intended to affect the timeliness of information exchange [7]. In [4,8], the authors address the problem of security constrained optimal control for discrete-time linear systems in which packets may be jammed by a malicious adversary, and the goal is to find optimal control and attack strategies assuming a maximum number of jamming actions over a prescribed (finite) control horizon. Similar scenarios are considered in [9], where the problem of stabilizing a discrete-time linear system under DoS is casted as a dynamic zero-sum game, and in [10] where the authors investigate the problem of designing optimal attack schedules to maximize the expected average estimation error at the remote estimator.

An alternative scenario is addressed in [11], where the authors consider the problem of stability under *periodic* DoS for linear sampled-data systems under state-feedback. The idea there is to identify the jamming signal so as to restrict the information exchange to the time intervals where no DoS occurs. This approach has been then extended in [12] by considering energy-constrained, but otherwise *unknown* DoS attacks.

In [13], we addressed afresh the problem of stability under energy-constrained, but unknown, DoS attacks within the framework of linear sampled-data systems under state-feedback. Our

☆ This work is partially supported by the Dutch Organization for Scientific Research (NWO) under the auspices of the project QUICK (QUantized Information Control for formation Keeping) and by the Research programme Robust Design of Cyber-physical Systems, financed by the Dutch Technology Foundation STW (grant no. 12696).

* Corresponding author.
E-mail addresses: c.de.persis@rug.nl (C. De Persis), p.tesi@rug.nl (P. Tesi).

http://dx.doi.org/10.1016/j.sysconle.2016.07.007
0167-6911/

work differs in many aspects from the aforementioned papers: in [4,8–10], the authors consider a pure discrete-time setting, while here we deal with sampled-data networked systems and the performance analysis is concerned with the continuous-time process state. Second, we do not formulate the problem as an optimal control design problem. The controller can be designed according to any suitable design method, robustness against DoS attacks being achieved thanks to the design of the network transmission times. Finally, we focus on nonlinear systems. Our work also differs from the one in [11,12] since the goal is not to identify the jamming signal; rather, the goal is to determine if stabilization is possible assuming only a bound on the fraction of the time the jammer is active. The considered approach, inspired by [1], consists in a suitable logic that determines in real-time the frequency of controller updates (the sampling times) depending on the DoS occurrence. In particular, the controller in [13] enjoys the following features:

(i) It ensures *global* exponential stability of the closed-loop system whenever the intervals over which communication is possible are predominant with respect to the intervals over which communication is denied;
(ii) It allows for the state-feedback control to be designed in accordance with any control design method, robustness against DoS being achieved thanks to the sampling logic;
(iii) It is *resilient* since the sampling rate varies depending on the DoS occurrence;
(iv) It allows for an explicit characterization of convergence rate, minimal inter-sampling time, and ratios between the "active" and "sleeping" periods of DoS which do not destroy closed-loop stability;
(v) It is flexible enough so as to allow the designer to choose from several implementation options that can be used to trade-off performance vs. communication resources.

The objective of this paper is to put forward the investigation of similar ideas for nonlinear systems. Although we follow the line of arguments of [13], a few of the steps we take are very peculiar to nonlinear systems, making the extension far from straightforward and deserving attention on its own right. It is shown that under certain additional conditions, which are needed to avoid finite-escape times phenomena during DoS, *asymptotic stability* can be still ensured. The analysis combines elements from event-based control and ISS control Lyapunov functions.

A preliminary version of the paper was presented in [14]. Compared with the latter, this paper has undergone a major reorganization of the arguments and provides a complete version of all the proofs. One of the ideas on which the results are derived, namely taking into account in the analysis the switching between stable and unstable modes is not new and has been studied for both linear [15] and nonlinear systems [16]. We stress that the main novelty of our results, which makes our contribution profoundly different from previous work, lies in the design of an event-based resilient control strategy and in the explicit characterization of the intervals during which stable and unstable modes are active as a consequence of the DoS status.

The remainder of this paper is organized as follows. In Section 2 we introduce the framework of interest and provide an overview of the problem. In Section 3, we describe the considered class of DoS attacks and provide some preliminary stability results. The main result with a characterization of the class of DoS signals under which stability is preserved is given in Section 4. In Section 5, we discuss the theoretical results as well as their practical implementation. An example is given in Section 6. Section 7 provides concluding remarks and outlines future research directions.

Notation: The notation for this paper is in the main standard. For a vector $x \in \mathbb{R}^n$, $\|x\|$ denotes Euclidean norm. Given a continuously differentiable function f, we denote by ∇f its gradient. A function $\alpha : [0, \infty) \to [0, \infty)$ is said to be of class $\mathcal{K}$ if it is continuous, strictly increasing, and $\alpha(0) = 0$. In addition, it is said to be of class $\mathcal{K}_\infty$ if $\alpha(s) \to \infty$ as $s \to \infty$. A function $\beta : [0, \infty) \times [0, \infty) \to [0, \infty)$ is said to be of class $\mathcal{KL}$ if $\beta(\cdot, t)$ is of class $\mathcal{K}$ for each fixed $t \geq 0$ and $\beta(r, t)$ decreases to 0 as $t \to \infty$ for each fixed $r \geq 0$. Given two functions f and g, we denote by $f \circ g$ the composite function $f(g)$.

2. Framework

We consider nonlinear systems of the form

$$\dot{x} = f(x, u), \tag{1}$$

where $x \in \mathbb{R}^n$ is the state and $u \in \mathbb{R}^m$ is the control input. We assume that $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is Lipschitz continuous on compacts and satisfies $f(0, 0) = 0$. We also assume that there exists a function $\psi : \mathbb{R}^n \to \mathbb{R}^m$, which is Lipschitz continuous on compacts and satisfies $\psi(0) = 0$, such that $u = \psi(x)$ renders the closed-loop system input-to-state stable (ISS) with respect to measurement errors e in the sense that there exist a $\mathcal{C}^1$ function $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ and class $\mathcal{K}_\infty$ functions $\alpha_1, \alpha_2, \gamma$ such that

$$\begin{aligned} &\alpha_1(\|x\|) \leq V(x) \leq \alpha_2(\|x\|) \\ &\nabla V(x) f(x, \psi(x + e)) \leq -\lambda V(x) + \gamma(\|e\|), \end{aligned} \tag{2}$$

with $\lambda > 0$.

As shown in [17], the ISS property (2) does always hold whenever the closed-loop system is ISS in the classical sense [18].

The control action is implemented via a *sample-and-hold* device. In a *nominal* situation, given a sequence of times $\{t_k\}_{k \in \mathbb{N}_0}$, $t_0 := 0$, the control signal u is given by $\psi(x(t_k))$ for all $t \in [t_k, t_{k+1})$, where $k \in \mathbb{N}_0$. The mechanism that generates this sequence of times will be specified in the sequel. By nominal situation is meant that at each time t_k at which the actuator needs to update the control value, it correctly receives the sampled value $\psi(x(t_k))$. The focus of this paper is on a scenario that is different from the nominal one, namely one in which there might be times in the sequence $\{t_k\}_{k \in \mathbb{N}_0}$ at which the control signal cannot be updated since no information regarding $\psi(x(t_k))$ is received by the actuator. This loss of information can be caused by several factors, such as a defective communication channel or as a consequence of the action of an adversarial entity. Throughout the remainder of this note, we will refer to such a phenomenon as *Denial-of-Service* (DoS). Practical considerations will be discussed in some detail in Section 5.

Let $\{h_n\}_{n \in \mathbb{N}_0}$, $h_0 \geq 0$, represent the sequence of DoS *off/on* transitions, *i.e.*, the sequence of time instants where the network changes from nominal to DoS status. Along with $\{h_n\}_{n \in \mathbb{N}_0}$, we consider a sequence $\{\tau_n\}_{n \in \mathbb{N}_0}$, $\tau_n \geq 0$, which specifies the duration of the nth DoS status. Accordingly, we let

$$H_n = \{h_n\} \cup [h_n, h_n + \tau_n[\tag{3}$$

represent the nth DoS time-interval. We assume that, during DoS, the actuator generates an input that is based on the *most recently received* control signal. A very similar analysis can be carried out in case the zero-input strategy is considered [6]. Given $t \in \mathbb{R}_{\geq 0}$, we denote by $\Theta(t)$ the set of all successful transmissions over the interval $[0, t]$. The control signal applied to the process at each time can be therefore expressed in compact form as

$$u(t) = \psi(x(t_{k(t)})) \tag{4}$$

where

$$k(t) := \begin{cases} -1, & \text{if } \Theta(t) = \emptyset \\ \sup\{k \in \mathbb{N} | t_k \in \Theta(t)\}, & \text{otherwise.} \end{cases} \tag{5}$$