Automatica 79 (2017) 42–51

Contents lists available at ScienceDirect

Automatica

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

Resilient control under Denial-of-Service: Robust design☆

Shuai Feng, Pietro Tesi

ENTEG, Faculty of Science and Engineering, University of Groningen, 9747 AG Groningen, The Netherlands

ARTICLE INFO

Article history:
Received 15 January 2016
Received in revised form
27 July 2016
Accepted 9 January 2017
Available online 2 March 2017

Keywords:
Cyber-physical systems
Denial-of-Service
Sampled-data control
Networked control systems

ABSTRACT

In this paper, we study networked control systems in the presence of Denial-of-Service (DoS) attacks, namely attacks that prevent transmissions over the communication network. The control objective is to maximize frequency and duration of the DoS attacks under which closed-loop stability is not destroyed. A family of impulsive controllers is proposed, which achieve the considered control objective for a general class of DoS signals. An example is given to illustrate the proposed solution approach.

1. Introduction

Owing to advances in computing and communication technologies, recent years have witnessed a growing interest towards cyber–physical systems (CPSs), *i.e.*, systems where physical processes are monitored/controlled via embedded computers and networks (Lee, 2008; Sha, Gopalakrishnan, Xue, & Qixin, 2008). The concept of CPSs is extremely appealing for automation but it raises many theoretical and practical challenges. In particular, CPSs have triggered the attention towards networked control in the presence of cyber attacks. In fact, unlike general-purpose computing systems where attacks limit their impact to the cyber realm, attacks to CPSs can impact the physical world (Teixeira, Shames, Sandberg, & Johansson, 2015). There are varieties of cyber attacks such as *Denial-of-Service (DoS)*, *bias injection* and *zero-dynamics attacks* (Teixeira et al., 2015), to name a few. The last two are examples of attacks that affect the integrity of data, while DoS attacks are meant to compromise the availability of data.

This paper is concerned with DoS attacks. We consider a control system in which the measurement channel is networked; the attacker objective is to induce closed-loop instability by blocking the plant-controller communication. In wireless networks, this can be realized by injecting an interference signal (*jamming* signal) in the communication channel, examples being *constant*, *random* and *protocol-aware jamming* (DeBruhl & Tague, 2011; Pelechrinis, Iliofotou, & Krishnamurthy, 2011; Tague, Li, & Poovendran, 2009). It is known that communication failures induced by DoS can have a profile quite different from genuine packet losses, the latter being the case considered in the majority of studies on networked control; in particular, failures induced by DoS need not follow a class of probability distributions (Amin, Càrdenas, & Sastry, 2009). This raises many new theoretical challenges from the perspective of analysis and control design.

In Amin et al. (2009) and Gupta, Langbort, and Başar (2010), the authors address the problem of finding optimal control and DoS attack strategies assuming a maximum number of jamming actions over a prescribed control horizon. A quite similar problem is considered in Ugrinovskii and Langbort (2014), where the authors study zero-sum games between controllers and strategic jammers. In Shisheh Foroush and Martínez (2012, 2013), the authors study DoS attacks in the form of pulse-width modulated signals. The objective is to identify salient features of the DoS signal such as maximum *on/off* cycle in order to de-synchronize the transmissions from the occurrence of DoS. In De Persis and Tesi (2014, 2015), a framework is introduced where no assumption is made regarding the "structure" of the DoS attack signal. A general model is considered that constrains DoS only in terms of its *frequency* and *duration*. The main contribution is an explicit characterization of DoS frequency and duration for which closed-loop stability can be preserved by means of state-feedback controllers. Building on this framework, extensions have been considered dealing with dynamic controllers (Dolk, Tesi, De Persis, & Heemels, 2015), nonlinear (De Persis & Tesi, 2016) and distributed (Senejohnny, Tesi, & De Persis, 2015) systems, as well as with systems where jamming attacks and genuine packet losses coexist (Cetinkaya, Ishii, & Hayakawa, 2015).

In this paper, we study networked systems subject to DoS attacks from the perspective of designing maximally robust

☆ The material in this paper was partially presented at the 2016 American Control Conference, July 6–8, 2016, Boston, MA, USA. This paper was recommended for publication in revised form by Associate Editor Hideaki Ishii under the direction of Editor Christos G. Cassandras.

E-mail addresses: s.feng@rug.nl (S. Feng), p.tesi@rug.nl (P. Tesi).

http://dx.doi.org/10.1016/j.automatica.2017.01.031
0005-1098/

controllers. To this end, in Section 2, we introduce a measure of robustness against DoS, which is related to the average percentage of transmission failures, or *jamming rate* (Anantharamu, Chlebus, Kowalski, & Rokicki, 2011), that the closed-loop system can tolerate before instability can occur. In Section 3, we then focus on control design. From the point of view of robustness, static feedback has inherent limitations. In fact, static feedback generates control updates only when new process measurements become available. Intuitively, such a limitation can be overcome via dynamic controllers. In particular, a natural approach is to equip the control system with prediction capabilities so as to reconstruct the missing measurements during DoS. Inspired by recent results on finite-time observers (Ferrante, Gouaisbaut, Sanfelice, & Tarbouriech, 2014; Raff & Allgöwer, 2008), we focus the attention on *impulsive* controllers, which make use of dynamical observers with measurements-triggered state resetting. We show that in case of full-state measurements, this class of controllers is indeed *maximally robust* in the sense that it guarantees stability for all the DoS signals with frequency and duration below a certain critical threshold beyond which stability can be lost irrespective of the adopted controller. For the case of partial-state measurements, the gap from the optimal bound is explicitly quantified as a function of process observability index, packets transmission rate and DoS parameters. In addition, we show that the optimal bound can be recovered if the sensor system is equipped with computation capabilities and the communication protocol is acknowledgment-based. Both continuous and sampled-data implementations are discussed. As for the latter, in order to preserve stability along with the same optimality bound, one has to constrain the controller sampling rate. Such constraints are explicitly characterized in Section 4. In Section 5, an example is discussed, while Section 6 ends the paper with concluding remarks. A note on the case of delays is given in Appendix B. A preliminary version of this paper appeared in Feng and Tesi (2016).

Notation. Given a vector $v \in \mathbb{R}^n$, $\|v\|$ is its Euclidean norm. Given a matrix M, $\|M\|$ is its spectral norm. Given two sets A and B, we denote by $B \setminus A$ the relative complement of A in B, *i.e.*, the set of all elements belonging to B, but not to A. Given an interval $\mathcal{I}$, $|\mathcal{I}|$ denotes its length, and given a set $\mathcal{S} = \bigcup_k \mathcal{I}_k$ consisting of a countable union of intervals $\mathcal{I}_k$, $|\mathcal{S}|$ denotes its Lebesgue measure. Given a measurable function $f : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^n$ and a time interval $[0, t]$ we denote the $\mathcal{L}_\infty$ norm of $f(\cdot)$ on $[0, t]$ by $\|f_t\|_\infty := \sup_{s\in[0,t]} \|f(s)\|$. Given a measurable function $f : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^n$ we say that f is bounded if its $\mathcal{L}_\infty$ norm is finite.

2. The framework

2.1. Process dynamics and network

Consider the control architecture in Fig. 1. The process to be controlled is given by

$$\begin{cases} \dot{x}(t) = Ax(t) + Bu(t) + d(t) \\ y(t) = Cx(t) + n(t) \\ x(0) = x_0 \end{cases} \tag{1}$$

where $t \in \mathbb{R}_{\geq 0}$; $x \in \mathbb{R}^{n_x}$ is the state, $u \in \mathbb{R}^{n_u}$ is the control input, and $y \in \mathbb{R}^{n_y}$ is the measurement vector; (A, B) is stabilizable; $d \in \mathbb{R}^{n_x}$ is an unknown bounded disturbance, while $n \in \mathbb{R}^{n_y}$ accounts for bounded measurement and network-induced noises.

We assume that the measurement channel is networked and subject to DoS. Let t_k denote the kth transmission attempt. We shall assume that the transmission attempts are carried out periodically with period Δ, *i.e.*,

$$t_{k+1} - t_k = \Delta, \quad k \in \mathbb{N}_0 \tag{2}$$

with $t_0 := 0$.

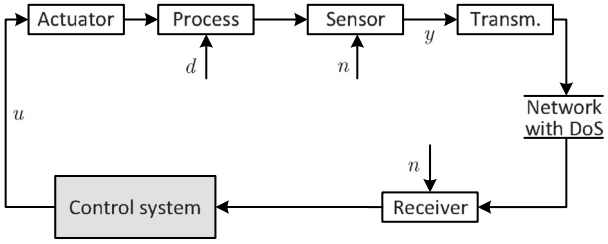

Fig. 1. Network control framework.

2.2. Denial-of-Service

We refer to DoS as the phenomenon for which some transmission attempts may fail. We consider a general DoS model that constrains the attacker action in time by only posing limitations on the frequency of DoS attacks and their duration. Let $\{h_n\}_{n\in\mathbb{N}_0}$ with $h_0 \geq 0$ denote the sequence of DoS *off/on* transitions, that is, the time instants at which DoS exhibits a transition from zero (transmissions are successful) to one (transmissions are not successful). Hence,

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

represents the nth DoS time-interval, of a length $\tau_n \in \mathbb{R}_{\geq 0}$, over which the network is in DoS status. If $\tau_n = 0$, then H_n takes the form of a single pulse at h_n. Given $\tau, t \in \mathbb{R}_{\geq 0}$ with $t \geq \tau$, let $n(\tau, t)$ denote the number of DoS *off/on* transitions over $[\tau, t]$, and let

$$\Xi(\tau, t) := \bigcup_{n\in\mathbb{N}_0} H_n \bigcap [\tau, t] \tag{4}$$

be the subset of $[\tau, t]$ where the network is in DoS status.

Assumption 1 (*DoS Frequency*)**.** There exist constants $\eta \in \mathbb{R}_{\geq 0}$ and $\tau_D \in \mathbb{R}_{>0}$ such that

$$n(\tau, t) \leq \eta + \frac{t - \tau}{\tau_D} \tag{5}$$

for all $\tau, t \in \mathbb{R}_{\geq 0}$ with $t \geq \tau$. ■

Assumption 2 (*DoS Duration*)**.** There exist constants $\kappa \in \mathbb{R}_{\geq 0}$ and $T \in \mathbb{R}_{>1}$ such that

$$|\Xi(\tau, t)| \leq \kappa + \frac{t - \tau}{T} \tag{6}$$

for all $\tau, t \in \mathbb{R}_{\geq 0}$ with $t \geq \tau$. ■

Remark 1. Assumptions 1 and 2 do only constrain a given DoS signal in terms of its *average* frequency and duration. Following Hespanha and Morse (1999), τ_D can be defined as the average dwell-time between consecutive DoS off/on transitions, while η is the chattering bound. Assumption 2 expresses a similar requirement with respect to the duration of DoS. It expresses the property that, on the average, the total duration over which communication is interrupted does not exceed a certain *fraction* of time, as specified by $1/T$. Like η, the constant κ plays the role of a regularization term. It is needed because during a DoS interval, one has $|\Xi(h_n, h_n + \tau_n)| = \tau_n > \tau_n/T$. Thus κ serves to make (6) consistent. Conditions $\tau_D > 0$ and $T > 1$ imply that DoS cannot occur at an infinitely fast rate or be always active. ■

2.3. A robustness measure against DoS

Suppose that a transmission attempt t_k falls within H_n. Due to the finite transmission rate $1/\Delta$, the first successful attempt after t_k need not occur exactly when H_n is over. Thus, $\bar{H}_n := H_n \cup [h_n + \tau_n, h_n + \tau_n + \Delta[$ yields an upper bound on the nth time interval