Proceedings of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Resilient Control under Denial-of-Service

C. De Persis * **P. Tesi** *

* *ITM, Faculty of Mathematics and Natural Sciences, University of Groningen, 9747 AG Groningen, The Netherlands*

Abstract: We investigate resilient control strategies for linear systems under Denial-of-Service (DoS) attacks. By DoS attacks we mean interruptions of communication on measurement (sensor-to-controller) and control (controller-to-actuator) channels carried out by an intelligent adversary. We characterize the duration of these interruptions under which stability of the closed-loop system is preserved. The resilient nature of the control descends from its ability to adapt the sampling rate to the occurrence of the DoS.

Keywords: Cyber-physical systems; Digital control; Control under limited information; Resilient control.

1. INTRODUCTION

In recent years there has been a growing interest concerning feedback control systems that are implemented over communication networks. These networks impose that measurements are acquired at discrete times, transmitted and received by the controller. The latter processes the received information and computes the control signal. This can in turn be sampled and transmitted to the actuators. Common limitations on these signals that travel over a network are quantization, delays and loss of information. Due to the limited bandwidth of the communication channel, as well as possible constraints on the available computational power, much research has been devoted to reduce the use of the communication line, by designing the sampling sequence based on current status of the process to control. This has given raise to a very active line of research in the context of *event/self-triggering* control; see Heemels, Johansson, and Tabuada (2012) for a recent comprehensive overview of the topic.

In the literature, several aspects of event/self-triggering control have been investigated, including output-feedback (Donkers and Heemels (2010)), robustness against additive disturbances (Mazo Jr, Anta, and Tabuada (2010)), large-scale systems (Wang and Lemmon (2011); De Persis, Sailer, and Wirth (2013)) and distributed coordinated control (Seyboth, Dimarogonas, and Johansson (2013); De Persis and Frasca (2013)), to name a few. On the other hand, an aspect of primary importance for which less results are available is the robustness of such schemes against malicious attacks.

Attacks to computer networks have become ever more prevalent over the last years. In this respect, one of the most common type of attack is the so-called *Denial-of-Service* (DoS); see Byres and Lowe (2004). While networked control formulations have previously considered sensor/control packet losses (Schenato, Sinopoli, Franceschetti, Poolla, and Sastry (2007)), dealing with DoS phenomena requires fundamentally different analysis tools. In fact, in contrast with classical networked control systems where packet losses can be reasonably modeled as random events, assuming a stochastic characterization of the DoS attacks would be inherently limiting in that it would fail to capture the malicious and intelligent nature of an attacker.

Prompted by these considerations, this paper discusses the problem of controlling networked systems subject to DoS attacks, whose underlying strategy is *unknown*. More specifically, we consider a classical *sampled-data* control scheme consisting of a continuous-time linear process in feedback loop with a digital controller. An attacker, according to some unknown strategy, can interrupt both sensor and control communication channels. Under such circumstances, the process evolves under out-of-date control. Within this context, we address the question of designing control update rules that are robust against the occurrence of DoS. In this respect, the main contribution of this paper is to show that suitable control update rules do exist whenever the ratio between the "active" and "sleeping" periods of jamming is small enough on the average. This somehow reminds of stability problems for systems that switch between stable and unstable modes; see e.g. Zhai, Hu, Yasuda, and Michel (2000). In our paper, however, the peculiarity of the problem under study leads to a different analysis and results. We also show that the results here introduced are flexible enough so as to allow the designer to choose from several implementation options that can be used for trading-off performance vs. communication resources. Although these solutions originate from different approaches, they exhibit the common feature of *resilience*, by which we mean the possibility to adapt the sampling rate to the DoS occurrence.

Previous contributions to this research line have been reported in Amin, Càrdenas, and Sastry (2009); Gupta, Langbort, and Başar (2010). In these papers, however, the framework is substantially different. They consider a pure discrete-time setting and the goal is to find optimal control and attack strategies assuming a maximum number of jamming actions over a prescribed (finite) control horizon. Here, we do not formulate the problem as an optimal control design problem. The controller can be designed according to any suitable design method, robustness and

resilience against DoS attacks being achieved thanks to the design of the control update rule. Perhaps, the closest reference to our research is Foroush and Martínez (2013). In that paper, the authors consider a situation where the attack strategy is known to be *periodic*, though of unknown period and duration. The goal is then to identify period and duration of the jamming activity so as to determine the time-intervals where communication is possible. Their framework should be therefore looked at as complementary more than alternative to the present one, since the former deals with cases where one can adjust the control updates so that they never fall into the jamming activity periods. Such a feature is conceptually impossible to achieve in scenarios such as the one considered in this paper, where the jamming strategy is neither known nor prefixed (the attacker can modify on-line the attack strategy).

Due to lack of space, proofs have been omitted. They can be found in De Persis and Tesi (2013).

2. FRAMEWORK AND PROBLEM OVERVIEW

The framework of interest is represented in Figure 1. We consider a remote plant-operator setup, in which the process to be controlled is described by the differential equation

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

where $t \in \mathbb{R}_{\geq 0}$; $x \in \mathbb{R}^{n_x}$ is the state and $u \in \mathbb{R}^{n_u}$ is the control input; We assume that a state-feedback matrix K has been designed such that all the eigenvalues of $A+BK$ have negative real part.

The control action is implemented via a *sample-and-hold* device. Let $\{t_k\}$, $k \in \mathbb{N}$, $t_0 := 0$, represent the sequence of time instants at which it is desired to update the control action. At the present stage, for simplicity of exposition, we simply refer to the "Logic" block as the device responsible for generating $\{t_k\}$. Thus, whatever the logic underlying this block, in the ideal situation where data can be sent and received at any desired instant of time, the control input applied to the process would be $u_{\text{ideal}}(t) = K\,x(t_k)$ for all $t \in [t_k, t_{k+1}[$.

We shall refer to *Denial-of-Service* (DoS, for short) as the phenomenon that prevents u_{ideal} from being executed at each desired t_k. In this paper, we consider the case of a DoS simultaneously affecting both control and measurement channels. This amounts to assuming that, in the presence of DoS, data can be *neither sent nor received*. Let $\{h_n\}$, $n \in \mathbb{N}$, $h_0 \geq 0$, represent the sequence of DoS positive edge-triggering, *i.e.* the time instants at which the DoS exhibits a transition from, say, zero (communication is possible) to, say, one (communication is interrupted). Accordingly,

$$H_n := [h_n, h_n + \tau_n[\tag{2}$$

will denote the n-th DoS time-interval, of a length τ_n, over which communication is not possible. We then assume that, in the presence of DoS, the actuator generates an input that is based on the *most recently received* control signal. Specifically, denote the set of time-instants up to time t where communication is possible by

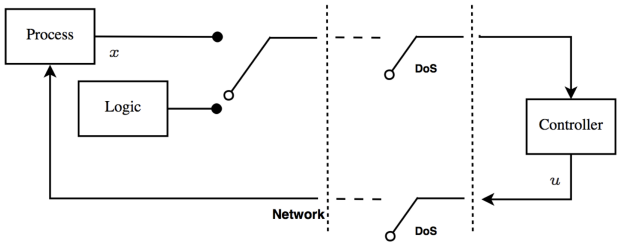

Fig. 1. Block diagram of the closed-loop system under DoS on the communication channels.

$$\Theta(t) := [0, t] \setminus \bigcup_{n \in \mathbb{N}} H_n \tag{3}$$

where $\setminus$ means relative complement.

Accordingly, the control input applied to the process can be expressed as

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

where

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

denote the last (up to the current time) successful control update. Notice that $h_0 = 0$ implies $k(0) = -1$, which raises the question of assigning a value to the control input when communication is not possible at the process start-up. In this respect, we assume that when $h_0 = 0$ then $u(0) = 0$, and we let $x(t_{-1}) := 0$ for notational consistency.

2.1 Problem overview

To begin with, we introduce the following definition.

Definition 1. Consider the control system Σ composed of (1) under a state-feedback control as in (4). Σ is said to be *globally exponentially stable* (GES) if there exist $\alpha, \beta \in \mathbb{R}_{>0}$ such that

$$\|x(t)\| \leq \alpha e^{-\beta t} \|x(0)\| \tag{6}$$

for all $t \in \mathbb{R}_{\geq 0}$ and for all $x(0) \in \mathbb{R}^{n_x}$, where $\|\cdot\|$ stands for Euclidean norm. □

Various approaches have been considered assuring GES to the control system in the absence of DoS; *e.g.*, see Heemels et al. (2012) for recent results and a discussion on questions related to periodic vs aperiodic implementations. A natural question then arises on whether mechanisms do exist that are capable of preserving GES under DoS.

In this respect, some preliminary considerations are in order. Whatever the rule generating the $\{t_k\}$-sequence, ultimate goal of the "Logic" block is to update the control action frequently enough so that stability is not destroyed. While in principle this is always possible in the absence of DoS, the same conclusions do not hold if DoS is allowed to be arbitrary. For instance, for open-loop unstable systems, one immediately sees that if $\tau_0 = \infty$ then stability is lost irrespective of how $\{t_k\}$ is chosen. These points motivate the following restriction on the admissible DoS signals considered throughout the paper.

Given a sequence $\{h_n\}$, let