

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, Cárdenas, 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 Froush 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

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

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 $u_{\text{ideal}}(t) = Kx(t_k)$ 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,

$$H_n := [h_n, h_n + \tau_n[\quad (2)$$

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.

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

where \setminus means relative complement.

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

$$u(t) = Kx(t_{k(t)}) \quad (4)$$

where

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

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

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

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

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

$$\Xi(t) := \bigcup_{n \in \mathbb{N}} H_n \cap [0, t] \quad (7)$$

denote the total interval of DoS up to the current time. Given an interval I , let $|I|$ denote its length.

Assumption 1. The DoS sequence $\{h_n\}$, $n \in \mathbb{N}$, is such that $\inf_{n \in \mathbb{N}} \tau_n > 0$. Moreover, there exist constants $\kappa \in \mathbb{R}_{\geq 0}$ and $\tau \in \mathbb{R}_{> 0}$ such that

$$|\Xi(t)| \leq \kappa + \frac{t}{\tau} \quad (8)$$

for all $t \in \mathbb{R}_{\geq 0}$. \square

Remark 1. Condition $\inf_{n \in \mathbb{N}} \tau_n > 0$ ensures that $\{h_n\}$ is *non-Zeno* and that infinitely many DoS intervals always have strictly positive Lebesgue measure. Inequality (8) expresses the property that the DoS satisfies a *slow-on-the-average* type condition, as introduced by Hespanha and Morse (1999) for hybrid systems analysis. In the present context, the rationale behind (8) is that if $\kappa = 0$ then the average time interval of DoS is at least τ . On the other hand, $\kappa > 0$ allows for consideration of DoS at the process start-up, *i.e.* when $h_0 = 0$. \square

3. MAIN RESULTS

In this section, a simple control update rule is considered, which is capable of preserving GES for any DoS signal satisfying Assumption 1 with τ sufficiently large. A discussion on the results along with implementation aspects is deferred to the next section.

Let

$$e(t) := x(t_{k(t)}) - x(t) \quad (9)$$

where $t \in \mathbb{R}_{\geq 0}$, represent the error between the value of the process state at the last successful control update and the value of the process state at the current time. Consistent with the comments made right after (4), if $h_0 = 0$ then $e(t) = -x(t)$ for all $t \in H_0$. The closed-loop system composed of (1) and (4) can be then rewritten as

$$\dot{x}(t) = \Phi x(t) + BK e(t) \quad (10)$$

where $\Phi := A + BK$. Consider now the following control update rule

$$\|e(t)\| \leq \sigma \|x(t)\|, \quad \forall t \notin \Xi(t) \quad (11)$$

where $\sigma \in \mathbb{R}_{> 0}$ is a free design parameter. As shown hereafter, such an update rule is capable of preserving GES for any DoS signal satisfying Assumption 1 with τ sufficiently large.

Condition (11) was first introduced in Tabuada (2007) in the context of event-based control. The difference here is that, due to the presence of DoS, one cannot enforce this condition for all $t \geq 0$, but only over those time-intervals where communication is indeed possible.

To fix the ideas, it is convenient to briefly comment on a possible implementation of condition (11), referring the reader to Section 4 for a thorough discussion and possible variations. The simplest architecture one can think of for implementing (11) is as in Figure 2(a). The ‘‘Logic’’ block measures continuously the state x , computes the error signal e and detects the instants t_k at which (11) holds with

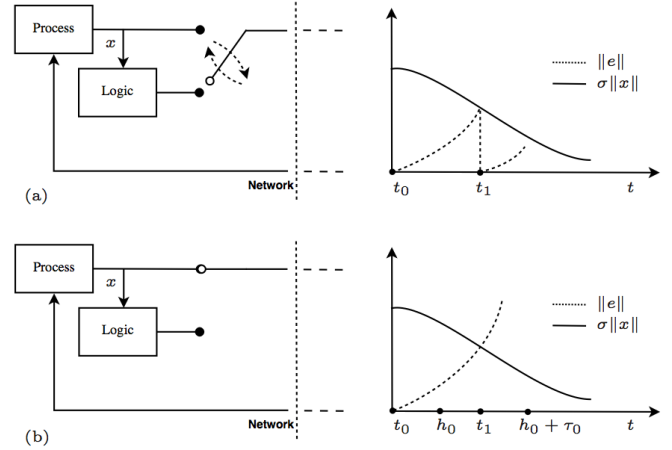


Fig. 2. Ideal mechanism for the fulfillment of (11): (a) absence of DoS; (b) presence of DoS.

the equality relation. At these instants, the logic samples the state and attempt to transmit it to the controller. In accordance with (9), if the control update is successful then e is reset to zero. Under DoS, the logic turns to a different operating mode by continuously attempting to update the control action, as depicted in Figure 2(b). In this way, at time $h_n + \tau_n$ when communication is restored, the logic is able to transmit immediately the sampled measurement so that (11) is enforced.

In the following subsection, for ease of exposition, we assume that this is indeed the case. In practice, when implementing (11) on a digital platform, due to the finite sampling rate, a time interval will necessarily elapse from the time $h_n + \tau_n$ at which DoS is over, to the time at which the logic successfully samples and transmits. As anticipated, this case will be addressed in Section 4.

3.1 Stability analysis

We now study the trajectories of the closed-loop system composed of (1) and (4) with control update law (11). An alternative approach to stability analysis, based on Lyapunov functions, is discussed in Appendix A.

Observe first that Φ is a stability matrix by hypothesis. Then there exist $\mu \in \mathbb{R}_{\geq 1}$ and $\lambda \in \mathbb{R}_{> 0}$ such that $\|e^{\Phi t}\| \leq \mu e^{-\lambda t}$ for all $t \in \mathbb{R}_{\geq 0}$, where μ and λ can be easily computed using algebraic matrix theory. This, in turns, implies

$$\|x(t)\| \leq \omega_1 e^{-\lambda t} + \int_{\Theta(t)} \omega_2 e^{-\lambda(t-s)} \|e(s)\| ds + \int_{\Xi(t)} \omega_2 e^{-\lambda(t-s)} \|e(s)\| ds \quad (12)$$

having defined $\omega_1 := \mu \|x(0)\|$ and $\omega_2 := \mu \|BK\|$ where, given a matrix M , $\|M\|$ denotes its spectral norm. We now evaluate the two integral terms in the above formula.

Consider first the set $\Theta(t)$, over which (11) holds by construction. The corresponding integral term can be therefore upper bounded as

$$\int_{\Theta(t)} \omega_2 e^{-\lambda(t-s)} \|e(s)\| ds \leq \int_0^t \omega_3 e^{-\lambda(t-s)} \|x(s)\| ds \quad (13)$$

where $\omega_3 := \sigma\omega_2$.

Consider next the set $\Xi(t)$ and, in particular, the n -th DoS interval H_n . Over each H_n , the process dynamics are governed by

$$\dot{x}(t) = Ax(t) + BKx(t_k(h_n)) \quad (14)$$

In addition, there exist $\theta \in \mathbb{R}_{\geq 1}$ and $\rho \in \mathbb{R}_{\geq 0}$ such that $\|e^{At}\| \leq \theta e^{\rho t}$ for all $t \in \mathbb{R}_{\geq 0}$. It is not difficult to verify that, over H_n , the closed-loop dynamics can be upper bounded as $\|x(t)\| \leq \bar{\theta} e^{\rho(t-h_n)} \|x(h_n)\|$ where $\bar{\theta} =: \theta + \theta(1 + \sigma)\|BK\|/\rho$.

Let now

$$n(t) = \begin{cases} -1, & \text{if } t < h_0 \\ \sup \{n \in \mathbb{N} \mid h_n < t\}, & \text{otherwise} \end{cases} \quad (15)$$

denote the last (up to the current time) DoS positive edge-triggering. In addition, let $\tau_n(t) := \min\{\tau_n, t - h_n\}$ denote the length of the last DoS interval up to time t . Simple calculations then yield

$$\begin{aligned} & \int_{h_n}^{h_n + \tau_n(t)} \omega_2 e^{-\lambda(t-s)} \|e(s)\| ds \\ & \leq \omega_*(\rho) e^{-\lambda(t-h_n)} \left[e^{(\lambda+\rho)\tau_n(t)} - 1 \right] \|x(h_n)\| \end{aligned} \quad (16)$$

for all $n \in \mathbb{N}$ with $n \leq n(t)$, where $\omega_*(\rho) =: \bar{\omega}/(\lambda + \rho)$ and $\bar{\omega} =: \omega_2(1 + \sigma) + \omega_2\bar{\theta}$.

By increasing ρ if necessary, we can always assume that $\omega_*(\rho) \leq 1$. Specifically, let ρ be any positive scalar such that $\|e^{At}\| \leq \theta e^{\rho t}$ where $\theta \in \mathbb{R}_{\geq 1}$. Hence, by defining $\rho_* := \inf \{\zeta \in \mathbb{R}_{\geq \rho} \mid \omega_*(\zeta) \leq 1\}$ and $\delta_n(t) := e^{(\lambda+\rho_*)\tau_n(t)} - 1$ one can always rewrite (16) as

$$\int_{h_n}^{h_n + \tau_n(t)} \omega_2 e^{-\lambda(t-s)} \|e(s)\| ds \leq \delta_n(t) e^{-\lambda(t-h_n)} \|x(h_n)\|$$

Hence, the last integral term of (12) can be upper bounded as

$$\int_{\Xi(t)} \omega_2 e^{-\lambda(t-s)} \|e(s)\| ds \leq \sum_{n=0}^{n(t)} \delta_n(t) e^{-\lambda(t-h_n)} \|x(h_n)\| \quad (17)$$

Using (13) and (17), the following results can be obtained; see (Bainov and Simeonov, 1992, Theorem 16.4).

Theorem 1. Consider the system Σ composed of (1) under a state-feedback control as in (4). Let $\Phi = A + BK$, with $\mu \in \mathbb{R}_{\geq 1}$ and $\lambda \in \mathbb{R}_{>0}$ satisfying $\|e^{\Phi t}\| \leq \mu e^{-\lambda t}$ for all $t \in \mathbb{R}_{\geq 0}$. Let the control update rule satisfy (11) with $\lambda - \sigma\mu\|BK\| > 0$. Then, Σ is GES for any DoS sequence satisfying Assumption 1 with

$$\tau > \frac{\lambda + \rho_*}{\lambda - \sigma\mu\|BK\|} \quad (18)$$

In particular, (6) holds true with constants $\alpha = \mu e^{\kappa(\lambda+\rho_*)}$ and $\beta = \lambda - \sigma\mu\|BK\| - (\lambda + \rho_*)/\tau$. \square

Remark 2. Condition $\lambda - \sigma\mu\|BK\| > 0$ must be satisfied even in the absence of DoS. It reflects the fact that, even when communication is always possible, in order to achieve stability, the control action must be updated frequently enough. On the other hand, (18) imposes constraints on the admissible DoS signals. In this respect, notice that τ must always be greater than one. This is consistent with intuition, and reflects the fact that, to achieve stability, the total length of DoS intervals must be a suitable *fraction* of the time (in fact, $|\Xi(t)| \leq t/\tau$ when $\kappa = 0$). \square

4. IMPLEMENTATION AND RESILIENT CONTROL LOGICS

The analysis of Section 3 hinges upon the fulfillment of condition (11). Such a condition cannot be implemented on digital platforms in that, in order to be fulfilled, it would require continuous transmission attempts upon DoS detection, *i.e.* an *infinite* sampling rate. Motivated by this, we first discuss how Theorem 1 can be generalized so as to possibly account for finite sampling rate constraints. We then consider a number of implementation possibilities that can be used to trade-off performance vs. communication resources within the proposed framework.

4.1 Stability under finite sampling rate

We first consider the following definition.

Definition 2. A control update sequence $\{t_k\}$ is said to occur at a *finite sampling rate* if there exist an $\varepsilon \in \mathbb{R}_{>0}$ such that

$$\Delta_k := t_{k+1} - t_k \geq \varepsilon \quad (19)$$

for all $k \in \mathbb{N}$. \square

Consider now a control update sequence $\{t_k\}$ along with a DoS sequence $\{h_n\}$, and let $\mathbb{S}_n := \{k \in \mathbb{N} \mid t_k \in H_n\}$ denote the set of integers associated with an attempt to update the control action during H_n . Accordingly, by defining

$$\Delta_{\mathbb{S}_n} := \sup_{k \in \mathbb{S}_n} \Delta_k \quad (20)$$

then $\bar{H}_n := [h_n, h_n + \tau_n + \Delta_{\mathbb{S}_n}[$ will provide an upper bound on the n -th time interval over which the control action is not updated, while

$$\bar{\Xi}(t) := \bigcup_{n \in \mathbb{N}} \bar{H}_n \cap [0, t] \quad (21)$$

will provide an upper bound on the total interval up to the current time over which the control action is not updated. Each $\Delta_{\mathbb{S}_n}$ essentially models the additional delay in the control update that may arise under finite sampling rate. In fact, under (19), $\Delta_{\mathbb{S}_n}$ will be greater than or equal to ε so that $|\bar{H}_n|$ will be strictly greater than $|H_n|$. Notice that \bar{H}_n may be exactly equivalent to the n -th time interval over which the control action is not updated. One may in fact have situations where a control update is requested just before the time $h_n + \tau_n$ at which the n -th DoS interval is over and the next sampling time is scheduled at $h_n +$

$\tau_n + \Delta_{S_n}$. Such a case cannot be ruled out being h_n and τ_n unknown.

The following result can be stated which extends the conclusions of Theorem 1 to control update sequences possibly occurring at a finite sampling rate.

Theorem 2. Consider the system Σ composed of (1) under a state-feedback control as in (4). Let $\Phi = A + BK$, with $\mu \in \mathbb{R}_{\geq 1}$ and $\lambda \in \mathbb{R}_{> 0}$ satisfying $\|e^{\Phi t}\| \leq \mu e^{-\lambda t}$ for all $t \in \mathbb{R}_{\geq 0}$. Let the control update rule satisfy (11) with $\lambda - \sigma\mu\|BK\| > 0$ and $\Xi(t)$ replaced by $\bar{\Xi}(t)$. Then, Σ is GES for any DoS sequence $\{h_n\}$ satisfying Assumption 1 with

$$\tau > \left(\frac{\lambda + \rho_*}{\lambda - \sigma\mu\|BK\|} \right) \left(1 + \frac{\Delta_*}{\tau_*} \right) \quad (22)$$

where $\Delta_* := \sup_{n \in \mathbb{N}} \Delta_{S_n}$ and $\tau_* := \inf_{n \in \mathbb{N}} \tau_n$. In particular, under the stated conditions, (6) holds true with constants $\alpha = \mu e^{(\lambda + \rho_*)(1 + \Delta_*/\tau_*)\kappa}$ and $\beta = \lambda - \sigma\mu\|BK\| - (\lambda + \rho_*)(1 + \Delta_*/\tau_*)/\tau$. \square

Remark 3. Theorem 2 differs from Theorem 1 due to the presence of Δ_* and τ_* . This has an intuitive explanation. In fact, in the ideal case considered in Theorem 1, $\Delta_* = 0$ since a control update can always occur as soon as DoS is over. Under finite sampling rate, each DoS interval will instead possibly introduce an additional delay in the control update. \square

4.2 Implementation and resilient control logics

The framework introduced with Theorem 2 is flexible enough so as to allow the designer to choose from several implementation options, some of which are described in the following. Although these solutions originate from fundamentally different approaches, they exhibit the common feature of *resilience*, by which we mean not only to ensure a certain degree of robustness against DoS, but also the ability to counteract it by changing the control update rule upon communication loss.

Event/Time-driven logics. As discussed in Section 3, the simplest architecture one can think of consists in using a “Logic” block that measures continuously the state x , computes the error signal e and detects the instants (events) at which

$$\|e(t)\| = \sigma\|x(t)\| \quad (23)$$

At these instants, the logic updates the control action. In the presence of DoS, the logic turns instead to a periodic operating mode with communication attempts occurring at a higher frequency rate ¹.

Proposition 1. Let Δ_1 be a positive scalar less than or equal to Δ_2 , with Δ_2 given by $\phi(\Delta_2) = \sigma$, the latter being the unique solution at Δ_2 of the scalar Riccati equation $\dot{\phi}(t) = \|\Phi\| + (\|\Phi\| + \|BK\|)\phi(t) + \|BK\|\phi^2(t)$ initialized at $\phi(0) = 0$. Then, the control update rule

$$t_{k+1} = \begin{cases} t_k + \Delta_1, & \text{if } t_k \in \Xi(t) \\ & \text{or } x(t_k) = 0 \end{cases} \quad (24)$$

$$\inf \{ t \in \mathbb{R}_{> t_k} : \|e(t)\| = \sigma\|x(t)\| \}, \quad \text{otherwise}$$

¹ A periodic update is also enforced when $x(t_k)=0$. This is because application of the second of (24) for $x(t_k)=0$ would result in a continuous control update.

satisfies the conditions of Theorem 2 with $\Delta_* = \Delta_1$ and $\Delta_k \geq \Delta_1$ for all $k \in \mathbb{N}$. \square

Purely time-driven logics. The rationale behind (24) is that, upon DoS, transmission is attempted at the sampling rate specified by Δ_1 , while, in the absence of DoS, less frequent control updates are allowed. This scheme has the positive feature of saving communication resources but requires continuous process state monitoring. If dedicated hardware is not available for this purpose, alternative logics like the next one may prove relevant.

Proposition 2. Let Δ_1 and Δ_2 be positive scalars with $\Delta_1 \leq \Delta_2$ and Δ_2 as in Proposition 1. Then, the control update rule

$$t_{k+1} = \begin{cases} t_k + \Delta_1, & \text{if } t_k \in \Xi(t) \\ t_k + \Delta_2, & \text{otherwise} \end{cases} \quad (25)$$

satisfies the conditions of Theorem 2 with $\Delta_* = \Delta_1$ and $\Delta_k \geq \Delta_1$ for all $k \in \mathbb{N}$. \square

Self-triggering logics. As a final option, we note that purely time-driven logics can be relaxed to more flexible aperiodic implementations by letting Δ_k to take values based on the available data. Logics of this kind are typically referred to as “self-triggering” in that the next update instant is computed directly by the control unit. Let $t_1, t_2 \in \mathbb{R}_{\geq 0}$ with $t_2 \geq t_1 \geq 0$ and define

$$\chi(t_2, t_1) := \left[e^{\Phi(t_2-t_1)} + \int_{t_1}^{t_2} e^{\Phi(t_2-s)} BK ds \right] x(t_1) \quad (26)$$

Thus $\chi(t_k, t_{k(t)})$ provides a prediction of $x(t_k)$ based on the last successful measurement $x(t_{k(t)})$. Thus, one can set Δ_k depending on the magnitude of $\|\chi(t_k, t_{k(t)})\|$: the larger $\|\chi(t_k, t_{k(t)})\|$ the smaller Δ_k and viceversa, which corresponds to increasing the sampling rate as the distance of the process state from the origin gets larger.

Proposition 3. Let Δ_1 and Δ_2 be positive scalars with $\Delta_1 \leq \Delta_2$ and Δ_2 as in Proposition 1. Let $\varphi : \mathbb{R}_{\geq 0} \mapsto [0, 1]$, be a class \mathcal{K} function. Then, the control update rule

$$t_{k+1} = t_k + \Delta_2 - (\Delta_2 - \Delta_1)\varphi(\|\chi(t_k, t_{k(t)})\|) \quad (27)$$

satisfies the conditions of Theorem 2 with $\Delta_* = \Delta_1$ and $\Delta_k \geq \Delta_1$ for all $k \in \mathbb{N}$. \square

5. CONCLUSIONS

We have studied resilient control strategies for linear systems under DoS. We have shown that to conclude asymptotic stability, DoS signals must not be active for more than a certain percentage of time on the average. The resilient nature of the proposed control strategy descends from its ability to adapt the sampling rate to the state of the process and to the occurrence of DoS attacks. The results lend themselves to be extended in various directions. We have not investigated the effect of possible limitations on the information, such as disturbances, quantization and delays, and leave the topic for future investigation. We envision the use of similar techniques to handle the problem in the presence of output feedback

and for nonlinear systems. Regarding the latter extension, the alternative Lyapunov-based analysis of the problem presented in Appendix A suits well our purpose. One of the motivations to consider control problems over networks descends from problems of distributed coordination and control of large-scale systems. Investigating our approach to resilient control under DoS for event-based coordination problems such as those in De Persis and Frasca (2013) represents another interesting research venue.

APPENDIX

Appendix A. LYAPUNOV-BASED APPROACH

Lyapunov arguments provide an alternative analysis of the problem that can be useful in some cases, such as when we deal with nonlinear control systems (De Persis and Tesi (2014)). Consider again the control system composed of (1) under a state-feedback control as in (4) with control update rule (11). Given any positive definite matrix $Q = Q^\top \in \mathbb{R}^{n_x \times n_x}$, let P be the unique solution of the Lyapunov equation $\Phi^\top P + P\Phi + Q = 0$. Then, by taking $V(x) = x^\top P x$ as a Lyapunov function, and computing it along the solution to (10), it is immediate to see that

$$\alpha_1 \|x(t)\|^2 \leq V(x(t)) \leq \alpha_2 \|x(t)\|^2 \quad (\text{A.1})$$

$$\dot{V}(x(t)) \leq -\gamma_1 \|x(t)\|^2 + \gamma_2 \|x(t)\| \|e(t)\| \quad (\text{A.2})$$

hold for all $t \in \mathbb{R}_{\geq 0}$, with α_1 and α_2 equal to the smallest and largest eigenvalue of P , respectively; γ_1 equal to the smallest eigenvalue of Q ; $\gamma_2 := \|K^\top B^\top P + PBK\|$.

Consider first $\Theta(t)$, over which (11) holds by construction. In this case, simple calculations yield

$$\dot{V}(x(t)) \leq -\omega_1 V(x(t)) \quad (\text{A.3})$$

where $\omega_1 := (\gamma_1 - \gamma_2\sigma)/\alpha_2$.

Consider next $\Xi(t)$. Very simple calculations show that $\|e(t)\| \leq (1 + \sigma)\|x(h_n)\| + \|x(t)\|$ for all $t \in H_n$. Thus, for all $t \in H_n$ such that $\|x(h_n)\| \leq \|x(t)\|$, one has

$$\begin{aligned} \dot{V}(x(t)) &\leq -\gamma_1 \|x(t)\|^2 + \gamma_2(2 + \sigma)\|x(t)\|^2 \\ &< \omega_2 V(x(t)) \end{aligned} \quad (\text{A.4})$$

where $\omega_2 := \gamma_2(2 + \sigma)/\alpha_1$. Conversely, for all $t \in H_n$ such that $\|x(h_n)\| > \|x(t)\|$, one has

$$\dot{V}(x(t)) < \omega_2 V(x(h_n)) \quad (\text{A.5})$$

Combining the last two inequalities with (A.3), the following result can be established.

Theorem 3. Consider the system Σ composed of (1) under a state-feedback control as in (4). Given any positive definite matrix $Q = Q^\top \in \mathbb{R}^{n_x \times n_x}$, let P be the unique solution of the Lyapunov equation $\Phi^\top P + P\Phi + Q = 0$ with $\Phi = A + BK$. Let $V(x) = x^\top P x$, and let the control update parameter σ in (11) be such that $\gamma_1 - \sigma\gamma_2 > 0$, with γ_1 and γ_2 as in (A.2). Then, Σ is GES for any DoS sequence $\{h_n\}$ satisfying Assumption 1 with

$$\tau > \frac{\omega_1 + \omega_2}{\omega_1} \quad (\text{A.6})$$

where $\omega_1 = (\gamma_1 - \gamma_2\sigma)/\alpha_2$ and $\omega_2 = \gamma_2(2 + \sigma)/\alpha_1$, and α_1 and α_2 as in (A.1). In particular, (6) holds true with $\alpha = \sqrt{e^{\kappa(\omega_1 + \omega_2)}}\alpha_2/\alpha_1$ and $\beta = [\omega_1 - (\omega_1 + \omega_2)/\tau]/2$. \square

The results of Section 4 can be then applied to the present case simply with

$$\tau > \left(\frac{\omega_1 + \omega_2}{\omega_1} \right) \left(1 + \frac{\Delta_*}{\tau_*} \right) \quad (\text{A.7})$$

in place of (22).

REFERENCES

- S. Amin, A. Cárdenas, and S.S. Sastry. Safe and secure networked control systems under denial-of-service attacks. *In Hybrid systems: Computation and Control*, pages 31–45, 2009.
- D.D. Bainov and P.S. Simeonov. *Integral Inequalities and Applications*. Kluwer Academic Publishers, Dordrecht, 1992.
- E. Byres and J. Lowe. The myths and facts behind cyber security risks for industrial control systems. *Proceedings of the VDE Congress, Berlin*, 2004.
- C. De Persis and P. Frasca. Robust self-triggered coordination with ternary controllers. *IEEE Transactions on Automatic Control*, 58(12):3024–3038, 2013.
- C. De Persis and P. Tesi. Resilient control under Denial-of-Service. *ArXiv, 0850369*, 2013.
- C. De Persis and P. Tesi. On resilient control of nonlinear systems under Denial-of-Service. *Submitted to 53rd IEEE Conference on Decision and Control, Los Angeles, CA*, 2014.
- C. De Persis, R. Sailer, and F. Wirth. Parsimonious event-triggered distributed control: A Zeno free approach. *Automatica*, 49:2116–2124, 2013.
- M.C.F. Donkers and W.P.M.H. Heemels. Output-based event-triggered control with guaranteed \mathcal{L}_2 -gain and Improved Event-triggering. *Proc. of the 49th IEEE CDC, Atlanta, GA, USA*, 2010.
- H. Shisheh Foroush and S. Martínez. On triggering control of single-input linear systems under pulse-width modulated dos jamming attacks. *International Journal of Robust and Nonlinear Control*, 2013. Submitted.
- A. Gupta, C. Langbort, and T. Başar. Optimal control in the presence of an intelligent jammer with limited actions. *Proc. of the 49th IEEE CDC, Atlanta, GA, USA*, 2010.
- W.P.M.H. Heemels, K.H. Johansson, and P. Tabuada. An introduction to event-triggered and self-triggered control. *Proc. of the 51th IEEE CDC, Maui, Hawaii USA*, 2012.
- J. P. Hespanha and A.S. Morse. Stability of switched systems with average dwell-time. *Proc. of the 38th IEEE CDC, Orlando, Florida USA*, 1999.
- M. Mazo Jr, A. Anta, and P. Tabuada. An ISS self-triggered implementation of linear controllers. *Automatica*, 46(8):1310–1314, 2010.
- L. Schenato, B. Sinopoli, M. Franceschetti, K. Poolla, and S.S. Sastry. Foundations of control and estimation over lossy networks. *Proceedings of the IEEE*, 95:163–185, 2007.
- G. Seyboth, D. Dimarogonas, and K. Johansson. Event-based broadcasting for multi-agent average consensus. *Automatica*, 49:245–252, 2013.
- P. Tabuada. Event-triggered real-time scheduling of stabilizing control tasks. *IEEE Trans. Aut. Contr.*, 52:1680–1685, 2007.
- X. Wang and M. Lemmon. Event-triggering in distributed networked control systems. 56(3):586–601, 2011.
- G. Zhai, B. Hu, K. Yasuda, and A.N. Michel. Stability analysis of switched systems with stable and unstable subsystems: An average dwell time approach. *In Proc. of the ACC, Chicago, Illinois, USA*, 2000.