University of Groningen

Networked Systems under Denial-of-Service: Co-located vs. Remote Control Architectures

Feng, Shuai; Tesi, Pietro

Published in:
 Systems & Control Letters

DOI:
[10.1016/j.sysconle.2017.08.003](https://doi.org/10.1016/j.sysconle.2017.08.003)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
 Final author's version (accepted by publisher, after peer review)

Publication date:
 2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Feng, S., & Tesi, P. (2017). Networked Systems under Denial-of-Service: Co-located vs. Remote Control Architectures. *Systems & Control Letters*, 108, 40-47. <https://doi.org/10.1016/j.sysconle.2017.08.003>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Networked Systems under Denial-of-Service: Co-located vs. Remote Control Architectures

Shuai Feng* Pietro Tesi*

* ENTEG, Faculty of Science and Engineering, University of Groningen,
9747 AG Groningen, The Netherlands (e-mail: {s.feng,p.tesi}@rug.nl).

Abstract: In this paper, we study networked systems in the presence of Denial-of-Service (DoS) attacks, namely attacks that prevent transmissions over the communication network. Previous studies have shown that co-located architectures (control unit co-located with the actuators and networked sensor channel) can ensure a high level of robustness against DoS. However, co-location requires a wired or dedicated actuator channel, which could not meet flexibility and cost requirements. In this paper we consider a control architecture that approximates co-location while enabling remote implementation (networked sensor and actuator channels). We analyze closed-loop stability and quantify the robustness “gap” between this architecture and the co-located one.

Keywords: Cyber-physical systems; Denial-of-Service; Networked control systems; Sampled-data control; Control under limited information.

1. INTRODUCTION

The field of cyber-physical systems is becoming more and more important in control engineering and computer science due to its broad spectrum of applications. Especially for safety-critical applications, there is a need for analysis, synthesis and design tools that can guarantee secure and reliable operations despite the presence of malicious attacks (Cheng et al., 2017).

Security of cyber-physical systems involves several research areas, including anomaly detection, verification, and estimation and control in the presence of attacks. Moreover, the problem varies depending on the type of attack one is considering. One usually classifies cyber attacks as either deception attacks, examples being zero-dynamics and bias injection attacks or Denial-of-Service (DoS) attacks. The former affect the data trustworthiness via the manipulation of packets transmitted through the network (Fawzi et al., 2011; Teixeira et al., 2015). DoS attacks are instead intended to affect the timeliness of the information exchange, *i.e.*, to induce packet losses (Amin et al., 2009; Xu et al., 2005). In particular, it is known that DoS attacks are relatively easy to launch (Xu et al., 2005), even without knowledge of the targeted system.

This paper deals with DoS attacks. We consider a control system in which plant-controller communication is networked. An attacker can induce closed-loop instability by denying the communication on sensor and actuator channels. The problem of interest is to design control systems that can tolerate DoS as much as possible. Networked control in the presence of packet losses has been widely investigated (Hespanha et al., 2007). However, it is known that communication failures induced by DoS can exhibit a temporal profile quite different from the one induced by genuine packet losses; in particular, communication failures induced by DoS need not follow a given class of probability distributions (Amin et al., 2009). This raises new theoretical challenges from the perspective of analysis as well as control design.

The literature on networked control under DoS is large and quite diversified. In (Amin et al., 2009; Gupta et al., 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 similar problem is considered in (Ugrinovskii and Langbort, 2014), where the authors study zero-sum games between controllers and strategic jammers. In (Zhang et al., 2016), the authors investigate DoS from the attacker’s viewpoint, and characterize optimal DoS attack scheduling strategies. Shisheh Feroosh and Martínez (2013) consider periodic DoS attacks. The objective is to identify salient features of the DoS signal such as maximum *on/off* cycle in order to de-synchronize the transmission times from the occurrence of DoS. In (De Persis and Tesi, 2014b; De Persis and Tesi, 2015), a framework is introduced where no assumption is made regarding the “structure” of the DoS attack. Instead, a model is considered that constrains DoS only in terms of its *frequency* and *duration*. The contribution is an explicit characterization of DoS frequency and duration for which stability can be preserved through state-feedback control. Extensions have been considered dealing with co-located robust controller design (Feng and Tesi, 2017), nonlinear (De Persis and Tesi, 2014a) and distributed (Senejohnny et al., 2015) systems, as well as with systems where jamming attacks and genuine packet losses coexist (Cetinkaya et al., 2016).

In a recent paper (Feng and Tesi, 2016), we have investigated networked systems under DoS attacks from the perspective of designing “maximally robust” controllers. There, it is shown that co-located architectures (control unit co-located with the actuators and networked sensor channel) can guarantee the highest possible level of robustness against DoS, in the sense that they can guarantee closed-loop 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 control system. Unfortunately, co-location requires wired or dedicated actuator channels, which could not meet flexibility and cost requirements.

In this paper, we consider a control system that approximates co-location while enabling remote implementation (networked sensor and actuator channels). The proposed architecture relies on packet-based control and actuator buffering. The basic idea is to transmit data packets containing plant input predictions whenever DoS is absent, and use the data stored in the actuator buffer during the periods of DoS. This idea has already proved effective to compensate for network delay (Bemporad, 1998; Chaillet and Bicchi, 2008) as well as packet losses (Quevedo and Nešić, 2011; Quevedo et al., 2011). In this paper, however, the peculiarity of the problem leads to a different analysis and design. We follow a “worst-case” type of analysis, accounting for the situation where the network undergoes periods of DoS much larger than the buffer capacity.

The paper contribution is twofold. First, we provide conditions on the prediction horizon under which closed-loop stability is preserved. The analysis is general enough to account for: (i) process disturbances; (ii) measurement and network noise; (iii) non-zero computation times. As a second contribution, we explicitly quantify the robustness “gap” between this control architecture and the co-located one, showing that the ideal bound obtained with co-location is recovered as the prediction horizon increases.

1.1 Notation

We denote by $\mathbb{R}$ the set of reals. For any $\alpha \in \mathbb{R}$, we denote $\mathbb{R}_{\geq \alpha} := \{x \in \mathbb{R} : x \geq \alpha\}$. We let $\mathbb{N}_0$ denote the set of nonnegative integers, $\mathbb{N}_0 := \{0, 1, \dots\}$. For any $\alpha \in \mathbb{N}_0$, we denote $\mathbb{N}_\alpha := \{\alpha, \alpha + 1, \dots\}$. The prime denotes transpose. Given a vector $v \in \mathbb{R}^n$, $\|v\|$ is its Euclidean norm. Given a matrix M , $\|M\|$ is its spectral norm. Given a measurable time 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 time 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. FRAMEWORK AND PAPER OUTLINE

2.1 Process dynamics and network

The process to be controlled is given by

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

where $t \in \mathbb{R}_{\geq 0}$; $x \in \mathbb{R}^m$ is the state, $u \in \mathbb{R}^m$ is the control input and $y \in \mathbb{R}^n$ is measurement vector; A and B are matrices of appropriate size with (A, B) stabilizable; $d \in \mathbb{R}^m$ and $n \in \mathbb{R}^n$ are unknown (bounded) disturbance and noise signals, respectively.

We assume that sensor and actuator channels are networked and subject to Denial-of-Service (DoS) status. The former implies that measurements and control commands are sent only at discrete time instants. Let $\{t_k\}_{k \in \mathbb{N}_0} = \{t_0, t_1, \dots\}$ denote the sequence of transmission attempts. Throughout the paper, we assume for simplicity that the transmission attempts are carried out periodically with period Δ , i.e.,

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

with $t_0 = 0$ by convention. We refer to DoS as the phenomenon for which some transmission attempts may fail.

We shall denote by $\{z_m\}_{m \in \mathbb{N}_0} = \{z_0, z_1, \dots\}$, $z_0 \geq t_0$, the sequence of time instants at which samples of y are successfully transmitted.

2.2 Control objective

The objective is to design a controller $\mathcal{K}$, possibly dynamic, in such a way that the closed-loop stability is maintained despite the occurrence of DoS in measurement and control channels. In this paper, by closed-loop stability we mean that all the signals in the closed-loop system remain bounded for any initial condition x_0 and bounded noise and disturbance signals, and converge to zero in the event that noise and disturbance signals converge to zero.

2.3 Denial-of-Service: Assumptions

Clearly, the problem in question does not have a solution if the DoS amount is allowed to be arbitrary. Following (De Persis and Tesi, 2015), 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 $\{h_n\}_{n \in \mathbb{N}_0}$, $h_0 \geq 0$, denote the sequence of DoS off/on transitions, i.e., the time instants at which DoS changes from zero (transmissions are possible) to one (transmissions are not possible). Hence, $H_n := \{h_n\} \cup [h_n, h_n + \tau_n[$ represents the n -th DoS interval, of 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

$$\Xi(\tau, t) := \bigcup_{n \in \mathbb{N}_0} H_n \cap [\tau, t] \quad (3)$$

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

We make the following assumptions.

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

$$n(\tau, t) \leq \eta + \frac{t - \tau}{\tau_D} \quad (4)$$

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

$$|\Xi(\tau, t)| \leq \kappa + \frac{t - \tau}{T} \quad (5)$$

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. In practice, 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 (5) consistent. Conditions $\tau_D > 0$ and $T > 1$ imply that DoS cannot occur at an infinitely fast rate or be always active. ■

2.4 Previous work and paper outline: co-located and remote control architectures

In Feng and Tesi (2016), we have investigated the problem of designing control architectures that achieve *maximal robustness* against DoS. We briefly summarize the main result of that paper. Consider the control architecture depicted in Figure 1. The control system is *co-located* with the process actuators and is equipped with prediction capabilities so as to reconstruct the missing measurements during DoS periods. Let $\delta := \Delta/b$ be the sampling rate of the control system, where $b \in \mathbb{N}_1$. Notice that choosing $\delta = \Delta/b$ allows to differentiate between controller sampling rate and transmission rate, maintaining Δ possibly large. Finally, let $A_\delta := e^{A\delta}$ and $B_\delta := \int_0^\delta e^{A\tau} B d\tau$. The predictor equations are given by

$$\begin{cases} \xi((q+1)\delta) = A_\delta \alpha(q\delta) + B_\delta u(q\delta) \\ \alpha(q\delta) = \begin{cases} y(q\delta), & \text{if } q\delta = z_m \\ \xi(q\delta), & \text{otherwise} \end{cases} \\ \xi(0) = 0 \end{cases} \quad (6)$$

where $q \in \mathbb{N}_0$, and the control action is given by

$$u(t) = K\alpha(q\delta), \quad t \in [q\delta, (q+1)\delta] \quad (7)$$

Here K is a state-feedback matrix such that all the eigenvalues of $\Phi = A + BK$ have negative real part. The control system consists of a predictor and a state-feedback matrix. The former emulates the process dynamics during the DoS period, and is “reset” whenever new measurements become available.

Theorem 1. (Feng and Tesi (2016)). Consider a process as in (1) under a co-located control system as in (6) and (7). Given any positive definite symmetric matrix M , let P denote the solution of the Lyapunov equation $\Phi'P + P\Phi + M = 0$. Let the controller sampling rate δ be such that

$$\delta \leq \frac{1}{\mu_A} \log \left[\left(\frac{\sigma}{1+\sigma} \right) \frac{\mu_A}{\max\{\|\Phi\|, 1\}} + 1 \right] \quad (8)$$

when $\mu_A > 0$, and

$$\delta \leq \left(\frac{\sigma}{1+\sigma} \right) \frac{1}{\max\{\|\Phi\|, 1\}} \quad (9)$$

when $\mu_A \leq 0$, where μ_A is the logarithmic norm of A and σ is a positive constant satisfying $\gamma_1 - \sigma\gamma_2 > 0$, where γ_1 is equal to the smallest eigenvalue of M and $\gamma_2 := \|2PBK\|$. Then, the closed-loop system is stable for any DoS sequence satisfying Assumptions 1 and 2 with arbitrary η and κ , and with τ_D and T such that

$$\frac{1}{T} + \frac{\Delta}{\tau_D} < 1 \quad (10)$$

The bound (10) is the best bound that one can achieve for the class of DoS signals satisfying Assumptions 1 and 2 in the sense that at or above the threshold “1” (when $1/T + \Delta/\tau_D \geq 1$) one can preserve stability only for *some* DoS signals, not for *all*. In fact, as soon as we reach the threshold “1”, τ_D and T can give rise to DoS signals that disrupt all the transmission attempts. Consider, for example, the DoS signal given by $(h_n, \tau_n) =$

Fig. 1. Co-located control architecture.

Fig. 2. Remote control architecture.

$(n\Delta, 0)$. This signal yields $1/T + \Delta/\tau_D = 1$ and satisfies Assumption 1 and 2 with $\eta = 1$, $\kappa = 0$, $T = \infty$ and $\tau_D = \Delta$. By construction, this DoS signal disrupts all the transmission attempts since it is synchronized with the transmission times. Note that condition $1/T + \Delta/\tau_D < 1$ requires $\tau_D > \Delta$. This means that, on the average, DoS cannot occur at the same rate as (or faster than) Δ .

Despite its robustness, the control architecture in (6)-(7) has a number of practical shortcomings: *co-location* requires a wired or dedicated control channel, which could not meet flexibility and cost requirements. In order to mitigate these shortcomings, we consider an architecture that allows for *remote* control. As shown in Figure 2, we consider a “packetized-and-buffered” architecture in which the controller transmits a sequence of control values containing process input predictions. Whenever communication is available, these values are stored in a buffer and used during the DoS periods. In the remainder of this paper, we formally investigate stability and robustness properties of this architecture. We show that under (10) closed-loop stability is preserved as long as the prediction horizon $h\delta$ satisfies

$$h\delta > \frac{\omega_2}{\omega_1 + \omega_2} (Q + \Delta) \quad (11)$$

where h is the buffer size, δ is the actuator sampling interval, ω_1 and ω_2 are process-dependent constants, and Q is a DoS-dependent constant. The above stability condition can be equivalently rewritten as

$$\frac{1}{T} + \frac{\Delta}{\tau_D} < 1 - \frac{\omega_2(\kappa + \eta\Delta)}{(\omega_1 + \omega_2)h\delta - \omega_2\Delta} \quad (12)$$

This inequality explicitly quantifies the gap between co-located and remote architectures, showing that the latter approaches the ideal bound as $h \rightarrow \infty$.

3. REMOTE CONTROL ARCHITECTURE

The actuator implements a *sample-and-hold* control strategy with sampling period $\delta := \Delta/b$, $b \in \mathbb{N}_1$. The control unit

operates at the time instants z_m at which process measurements are available, returning $h \in \mathbb{N}_1$ controls

$$u_p(z_m) = K\alpha_p(z_m), \quad p = 0, \dots, h-1 \quad (13)$$

where the variable $\alpha_p(z_m)$ defines the prediction of $x(z_m + p\delta)$ at time z_m . In particular, $\alpha_p(z_m)$ satisfies

$$\alpha_0(z_m) = y(z_m), \quad (14)$$

$$\alpha_{p+1}(z_m) = A_\delta \alpha_p(z_m) + B_\delta u_p(z_m) \quad (15)$$

where $p \in \{0, \dots, h-2\}$ and where A_δ and B_δ are as in (6). In practice, the control unit implements a sampled-data version of the process dynamics, which are “reset” whenever a new process measurement becomes available.

Let $I_p(z_m) := [z_m + p\delta, \min\{z_m + (p+1)\delta, z_{m+1}\}]$, where $p \in \{0, \dots, h-2\}$, and $I_{h-1}(z_m) := [z_m + (h-1)\delta, z_{m+1}]$. The control input applied to the process is

$$u(t) = \begin{cases} 0, & t \in [0, z_0[\\ u_p(z_m), & t \in I_p(z_m), \quad p = 0, \dots, h-2 \\ u_{h-1}(z_m), & t \in I_{h-1}(z_m) \end{cases} \quad (16)$$

In words, the control action is kept to zero until the first process measurement is received. Thereafter, as shown in Figure 3, over each interval $[z_m, z_{m+1}[$ the actuator applies the values in the buffer until the buffer is empty and the last value in the buffer $u(z_m + (h-1)\delta)$ is kept until z_{m+1} . On the other hand, if $z_{m+1} < z_m + h\delta$ then all the samples in the buffer are discarded and a new sequence of controls is stored. This renders (16) a “receding horizon” control policy.

Notice that (16) assumes that there is no noise in the actuator channel. This is because, if the actuator receives $u_p + n_u$, where n_u is a noise, the contribution of n_u can be absorbed in the process disturbance d . Also notice that (16) assumes that the time needed to compute all the samples in (13) is zero. This assumption can be relaxed as discussed in Section 4.

4. MAIN RESULT

We first present a lemma which is key to our developments. Then, we analyze the closed-loop behavior within and outside the prediction horizon and provide stability conditions.

4.1 Key lemma

The lemma relates DoS parameters and time elapsing between successful transmissions.

Lemma 1. [Feng and Tesi (2016)] Consider a transmission policy as in (2), along with a DoS signal satisfying Assumptions 1 and 2. If (10) holds true, then the sequence of successful transmissions satisfies $z_0 \leq Q$ and $z_{m+1} - z_m \leq Q + \Delta$ for all $m \in \mathbb{N}_0$, where

$$Q = (\kappa + \eta\Delta) \left(1 - \frac{1}{T} - \frac{\Delta}{\tau_D}\right)^{-1} \quad (17)$$

Fig. 3. Pictorial representation of (16) over a generic interval $[z_m, z_{m+1}[$ where $m \in \mathbb{N}_0$. The grey area represents the prediction horizon $h\delta$ during which the actuator utilizes the samples in the buffer, while the blue area exceeding $z_m + h\delta$ represents the period, say $|D_m|$, over which communication is not possible and there is no further control update.

4.2 Closed-loop behavior within the prediction horizon

Lemma 1 guarantees that the time elapsing between successful transmissions is always bounded under (10). In particular, z_0 is finite. This makes it possible to focus the attention on the closed-loop behavior from z_0 onwards. Consider a generic interval $I_p(z_m)$ and let

$$\phi(t) := \alpha_p(z_m) - x(t), \quad t \in I_p(z_m) \quad (18)$$

Notice that $I_p(z_m)$ might be empty, which happens whenever $z_{m+1} \leq z_m + p\delta$. Thus, in the sequel, we implicitly consider only the intervals $I_p(z_m)$ which are nonempty. Exploiting (18), we can rewrite the process dynamics as

$$\dot{x}(t) = \Phi x(t) + BK\phi(t) + d(t) \quad (19)$$

where $\Phi = A + BK$. Given any symmetric positive definite matrix M , let P be the unique solution of the Lyapunov equation $\Phi'P + P\Phi + M = 0$. Let $V(x) = x'Px$. Its derivative along the solutions to (19), satisfies

$$\begin{aligned} \dot{V}(x(t)) &\leq -\gamma_1 \|x(t)\|^2 + \gamma_2 \|x(t)\| \|\phi(t)\| \\ &\quad + \gamma_3 \|x(t)\| \|d(t)\| \end{aligned} \quad (20)$$

where γ_1 is the smallest eigenvalue of M , $\gamma_2 := \|2PBK\|$ and $\gamma_3 := \|2P\|$. One sees that closed-loop stability depends on ϕ , which, in turn, depends on the actuator sampling period δ . In general, it is not possible to get a dissipation inequality in (20) for an arbitrary δ . Nonetheless, this is possible provided that δ is suitably chosen.

Lemma 2. Consider a dynamical system as in (1) under the control action (16). Let the actuator sampling period δ be chosen as in Theorem 1. Then, there exists a positive constant ρ such that

$$\|\phi(t)\| \leq \sigma \|x(t)\| + \rho \|w_t\|_\infty \quad (21)$$

for all $t \in [z_m, \min\{z_m + h\delta, z_{m+1}\}]$, where $w = [d' \ n']'$.

Proof. See Appendix A. ■

In view of Lemma 2, it is immediate to verify that if actuator sampling period is properly chosen then closed-loop stability is guaranteed since (21) induces an ISS-type of property to the Lyapunov function. In fact, (21) yields

$$\begin{aligned} \dot{V}(x(t)) &\leq -(\gamma_1 - \sigma\gamma_2) \|x(t)\|^2 \\ &\quad + (\gamma_2\rho + \gamma_3) \|x(t)\| \|w_t\|_\infty \\ &= -\gamma_4 \|x(t)\|^2 + \gamma_5 \|x(t)\| \|w_t\|_\infty \end{aligned} \quad (22)$$

where $\gamma_4 := \gamma_1 - \sigma\gamma_2$ and $\gamma_5 := \gamma_2\rho + \gamma_3$. ■

Observe now that for any positive real ε , the Young's inequality (Hardy et al., 1952) yields

$$2\|x(t)\|\|w_t\|_\infty \leq \frac{1}{\varepsilon}\|x(t)\|^2 + \varepsilon\|w_t\|_\infty^2 \quad (23)$$

Using this inequality with $\varepsilon := \gamma_5/\gamma_4$, we obtain

$$\dot{V}(x(t)) \leq -\omega_1 V(x(t)) + \gamma_6\|w_t\|_\infty^2 \quad (24)$$

where $\omega_1 := \gamma_4/(2\alpha_2)$ and $\gamma_6 := \gamma_5^2/(2\gamma_4)$, where α_2 denotes the largest eigenvalue of P .

Accordingly,

$$V(x(t)) \leq e^{-\omega_1(t-z_m)}V(x(z_m)) + \zeta_1\|w_t\|_\infty^2 \quad (25)$$

for all $t \in [z_m, \min\{z_m + h\delta, z_{m+1}\}]$, where $\zeta_1 := \gamma_6/\omega_1$.

4.3 Closed-loop behavior outside the prediction horizon

Lemma 1 guarantees that the time elapsing between successful transmissions is always bounded under (10). Specifically, it ensures that $z_{m+1} - z_m \leq Q + \Delta$ for all $m \in \mathbb{N}_0$, where Q is a DoS-dependent constant. Consider now a transmission instant z_m . At this time, the buffer is full. If $h\delta \geq \Delta$ and there is no DoS at $z_m + \Delta$ then the buffer is filled again at $z_m + \Delta$ and the analysis of Section 4.2 applies. This scenario can be viewed as a co-located one. Thus, there are two critical cases: (i) when $h\delta < \Delta$ meaning that the prediction horizon does not cover one transmission period; and (ii) when $\Delta \leq h\delta < \Delta + Q$, which means that the network is subject to DoS periods that exceed the prediction horizon (cf. Figure 3). We now discuss this scenario in detail.

Let $D_m := [z_m + h\delta, z_{m+1}[$ with $D_m \neq \emptyset$. Recall that the prediction error satisfies

$$\phi(t) = \alpha_{h-1}(z_m) - x(t) \quad (26)$$

for all $I_{h-1}(z_m)$. By hypothesis, $D_m \neq \emptyset$ which means that $z_{m+1} > z_m + h\delta$. This implies that (26) is valid for $t = z_m + h\delta$. Hence, the triangular inequality yields

$$\begin{aligned} \|\phi(t)\| &\leq \|\alpha_{h-1}(z_m)\| + \|x(t)\| \\ &\leq \|x(z_m + h\delta)\| + \|\phi(z_m + h\delta)\| + \|x(t)\| \end{aligned} \quad (27)$$

for all $t \in D_m$. We first look at $\phi(z_m + h\delta)$. Note that ϕ is continuous on $I_{h-1}(z_m)$. This, along with Lemma 2, implies that

$$\|\phi(z_m + h\delta)\| \leq \sigma\|x(z_m + h\delta)\| + \rho\|w_{z_m+h\delta}\|_\infty \quad (28)$$

Combining this inequality with (27), we get

$$\|\phi(t)\| \leq (1 + \sigma)\|x(z_m + h\delta)\| + \|x(t)\| + \rho\|w_t\|_\infty \quad (29)$$

for all $t \in D_m$. Then, combining (29) and (20) yields

$$\begin{aligned} \dot{V}(x(t)) &\leq (\gamma_2 - \gamma_1)\|x(t)\|^2 \\ &\quad + \gamma_2(1 + \sigma)\|x(t)\|\|x(z_m + h\delta)\| \\ &\quad + (\gamma_3 + \rho\gamma_2)\|x(t)\|\|w_t\|_\infty \end{aligned} \quad (30)$$

By applying again the Young's inequality with $\varepsilon := \frac{\gamma_3 + \rho\gamma_2}{\gamma_1 - \sigma\gamma_2}$, we obtain

$$\begin{aligned} \dot{V}(x(t)) &\leq \gamma_2\|x(t)\|^2 \\ &\quad + \gamma_2(1 + \sigma)\|x(t)\|\|x(z_m + h\delta)\| + \gamma_7\|w_t\|_\infty^2 \end{aligned} \quad (31)$$

where $\gamma_7 := \frac{(\gamma_3 + \rho\gamma_2)^2}{2(\gamma_1 - \sigma\gamma_2)}$.

Hence, for every $t \in D_m$ such that $\|x(t)\| \geq \|x(z_m + h\delta)\|$ we have

$$\dot{V}(x(t)) \leq \omega_2 V(x(t)) + \gamma_7\|w_t\|_\infty^2 \quad (32)$$

where $\omega_2 := \gamma_2(2 + \sigma)/\alpha_1$. On the other hand, for every $t \in D_m$ such that $\|x(t)\| \leq \|x(z_m + h\delta)\|$ we have

$$\dot{V}(x(t)) \leq \omega_2 V(x(z_m + h\delta)) + \gamma_7\|w_t\|_\infty^2 \quad (33)$$

From these expressions, we conclude that during each interval D_m exceeding the prediction horizon the Lyapunov function satisfies

$$\begin{aligned} V(x(t)) &\leq e^{\omega_2(t-z_m-h\delta)}V(x(z_m + h\delta)) \\ &\quad + \zeta_2 e^{\omega_2(t-z_m-h\delta)}\|w_t\|_\infty^2 \end{aligned} \quad (34)$$

where $\zeta_2 := \gamma_7/\omega_2$.

4.4 Stability analysis

The foregoing analysis indicates that the overall behavior of the system can be regarded as the one of a switched system that behaves stably over the intervals $[z_m, z_m + h\delta]$ and unstably over the intervals D_m . We have the following result.

Theorem 2. Consider a dynamical system as in (1) under the control action (16). Given any positive definite symmetric matrix Q , let P denote the solution of the Lyapunov equation $\Phi'P + P\Phi + Q = 0$ with $\Phi = A + BK$ stable. Let the actuator sampling period δ be as in Theorem 1. Consider any DoS pattern satisfying (10). Then, the closed-loop system is stable if the prediction horizon $h\delta$ satisfies

$$h\delta > \frac{\omega_2}{\omega_1 + \omega_2}(Q + \Delta) \quad (35)$$

where Q is as in Lemma 1, $\omega_1 = \frac{\gamma_1 - \sigma\gamma_2}{2\alpha_2}$ and $\omega_2 = \frac{\gamma_2(2+\sigma)}{\alpha_1}$ where γ_1 , γ_2 and σ are as in Theorem 1 and where α_1 and α_2 are the smallest and largest eigenvalues of P , respectively.

Proof. Recall that z_0 exists finite in view of Lemma 1. Hence, we can restrict ourselves to study the closed-loop behavior from z_0 onwards. We also assume without loss of generality that $h\delta < Q + \Delta$ otherwise the result follows immediately from the analysis of Section 4.2. We show that

$$V(x(z_m)) \leq \lambda e^{-\beta(z_m - z_0)}V(x(z_0)) + 2\zeta(z_m) \sum_{k=0}^m L^k \quad (36)$$

where β is defines as

$$\beta := \frac{\omega_1 h \delta - \omega_2(Q + \Delta - h \delta)}{Q + \Delta} \quad (37)$$

and where

$$\lambda := e^{(\omega_1 + \omega_2)Q} \quad (38)$$

$$L := e^{-\beta \Delta} \quad (39)$$

$$\zeta(t) := 2 \max\{\zeta_1, \zeta_2\} e^{\omega_2(Q + \Delta - h \delta)} \|w_t\|_\infty^2 \quad (40)$$

Note that $\beta > 0$ in view of (35).

We prove this claim through an induction argument. The claim trivially holds for $m = 0$ since $\lambda \geq 1$. Assume next that the claim is true up to z_m with $m \in \mathbb{N}_1$. First recall that (25) implies

$$\begin{aligned} V(x(t)) &\leq e^{-\omega_1(t-z_m)} V(x(z_m)) + \zeta(t) \\ &\leq \lambda e^{-\beta(t-z_0)} V(x(z_0)) + 2e^{-\beta(t-z_m)} \zeta(t) \sum_{k=0}^m L^k + \zeta(t) \end{aligned} \quad (41)$$

for all $t \in [z_m, \min\{z_m + h\delta, z_{m+1}\}]$, where the second inequality follows from $\omega_1 > \beta$. Then, we have two cases. Assume $D_m = \emptyset$, which implies that $z_{m+1} \leq z_m + h\delta$. Note that $z_{m+1} - z_m \geq \Delta$ by construction. Hence, at time z_{m+1} we have $e^{-\beta(z_{m+1}-z_m)} \leq e^{-\beta \Delta} = L$. This shows that (36) holds true at z_{m+1} . Assume next $D_m \neq \emptyset$. In view of (25) and since V is continuous, we get

$$V(x(z_m + h\delta)) \leq e^{-\omega_1 h \delta} V(x(z_m)) + \zeta_1 \|w_{z_m+h\delta}\|_\infty^2 \quad (42)$$

Combining this inequality with (34) yields

$$\begin{aligned} V(x(t)) &\leq e^{\omega_2(t-z_m-h\delta)} V(x(z_m + h\delta)) + \zeta(t) \\ &\leq e^{\omega_2(t-z_m-h\delta)} [e^{-\omega_1 h \delta} V(x(z_m)) + \zeta(t)] + \zeta(t) \\ &\leq e^{-\beta(Q+\Delta)} V(x(z_m)) + 2\zeta(t) \\ &\leq \lambda e^{-\beta(t-z_0)} V(x(z_0)) + 2\zeta(t) \sum_{k=0}^{m+1} L^k \end{aligned} \quad (43)$$

for all $t \in D_m$, where the third inequality follows from the fact that $t \leq z_m + Q + \Delta$ and the definition of β , while the last inequality follows from the fact that $e^{-\beta(Q+\Delta)} \leq L$ and the fact that $Q + \Delta + z_m \geq z_{m+1} \geq t$ for all $t \in D_m$. Thus, (36) holds true at z_{m+1} .

Using this property, the proof can be easily finalized. In fact, (41) and (43) yield

$$V(t) \leq \lambda e^{-\beta(t-z_0)} V(x(z_0)) + 2\zeta(t) \sum_{k=0}^{m+1} L^k + \zeta(t) \quad (44)$$

for all $[z_m, z_{m+1}]$, and the sum term is bounded for any m since $L < 1$. This concludes the proof. ■

Remark 2. One sees from (35) that in order to get stability it is not necessary that the prediction horizon covers the maximum period of DoS. In fact, also large DoS periods can be tolerated. In particular, one sees that robustness increases with: (i) small values of ω_2 , which corresponds to mild open-loop unstable

dynamics; and (ii) large values of ω_1 , which corresponds to a large decay rate of the DoS-free closed-loop dynamics. The latter should be therefore taken into account when designing the state-feedback matrix K . ■

Remark 3. Taking into account the expression of Q from Lemma 1, it is immediate to see that (35) can be equivalently rewritten as (12). This shows that the considered architecture approaches the ideal bound (10) as $h \rightarrow \infty$. Inequality (12) does also quantify the “gap” between remote and co-located architectures for given any finite h . ■

Remark 4. In this paper, we have assumed that the time needed to compute h control values is zero. This assumption can be restrictive if h is large. As discussed in Feng and Tesi (2016), Lemma 1 follows from the fact that every time interval of duration $Q + \Delta$ contains at least a DoS-free interval of duration Δ . This result can be generalized. In fact, one can show that if we replace (10) with

$$\frac{1}{T} + \frac{\mu \Delta}{\tau_D} < 1, \quad \mu \in \mathbb{N}_1 \quad (45)$$

then every time interval of duration

$$Q_\mu := (\kappa + \eta \mu \Delta) \left(1 - \frac{1}{T} - \frac{\mu \Delta}{\tau_D}\right)^{-1} \quad (46)$$

contains at least a DoS-free interval of duration $\mu \Delta$. This property is very useful to account for non-zero computation times T_c . In fact, this property makes it possible to regard T_c as an extended DoS interval. The only modification is that the sequence of control values should start from $\alpha_{\lceil T_c/\Delta \rceil \Delta}$. ■

5. A NUMERICAL EXAMPLE

The numerical example is taken from (Forni et al., 2010). The system to be controlled is open-loop unstable and is characterized by the matrices

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (47)$$

The state-feedback matrix is given by

$$K = \begin{bmatrix} -2.1961 & -0.7545 \\ -0.7545 & -2.7146 \end{bmatrix} \quad (48)$$

The control system parameters are $\gamma_1 = 1$, $\gamma_2 = 2.1080$, $\alpha_1 = 0.2779$, $\alpha_2 = 0.4497$, $\|\Phi\| = 1.9021$, $\omega_1 = 0.5025$, $\omega_2 = 15.1709$ and $\mu_A = 1.5$. Disturbance d and noise n are random signals with uniform distribution in $[-0.01, 0.01]$.

The network transmission rate is given by $\Delta = 0.1$ s. As for the actuator sampling period, Theorem 1 yields $\delta < 0.1508$. We select $\delta = 0.1$ s in order to synchronize the controller sampling rate with Δ .

Figure 4 shows simulation results comparing the co-located architecture in (Feng and Tesi, 2016) with the one proposed in this paper. We consider a sustained DoS attack with variable period and duty cycle, generated randomly. Over a simulation horizon of 50s, the DoS signal yields $|\Xi(0, 50)| = 34.65$ s and $n(0, 50) = 39$. This corresponds to values (averaged over 50s) of $\eta \leq 3.1$, $\kappa \leq 0.8442$, $\tau_D \approx 1.2821$ and $T \approx 1.4430$, and $\sim 70\%$ of transmission failures. Moreover,

Fig. 4. Simulation results for the example in case disturbance and noise are random signals with uniform distribution in $[-0.01, 0.01]$. Top: Co-located architecture; Middle: Remote architecture without buffering; Bottom: Remote architecture with buffering ($h = 5$).

$$\frac{\Delta}{\tau_D} + \frac{1}{T} \approx 0.7710 \quad (49)$$

For the co-located architecture, the requirement for closed-loop stability is clearly satisfied.

Consider next the remote architecture. In agreement with (35), to get stability one needs a prediction horizon $h\delta \geq 4.9153$, which corresponds to $h \geq 50$. One sees from Figure 4 that without buffering, namely when $h = 1$, the closed-loop system is unstable. On the other hand, $h = 5$ is sufficient to obtain a satisfactory closed-loop response. The results indicate that in some cases, like the one discussed here, small values of h can approximate well the co-located architecture without actually requiring large h or co-location. In this sense, the proposed control architecture provides an effective solution to trade-off ease of implementation and robustness against DoS.

The conservativeness of the theoretical bound on h is somewhat implicit to the “worst-case” type of analysis pursued here.

6. CONCLUDING REMARKS

In this paper, we have investigated networked systems in the presence of Denial-of-Service attacks, comparing co-located and remote control architectures. While the former achieves the highest level of robustness for a very large class of DoS signals, it need not be always feasible from a practical point of view. We introduced a control architecture that approximates co-location while enabling remote implementation. This architecture relies on “packet-based” control: the control unit transmits a sequence of controls containing process input predictions. Whenever communication is allowed, the controls are stored in a buffer and used during the DoS periods. We studied closed-loop stability, quantifying the robustness “gap” between this architecture and the co-located one. As it emerges from both the analysis and the numerical simulations, the proposed architecture provides an effective solution for trading-off ease of implementation and robustness against DoS.

The present results can be extended in many directions. The case of partial state measurements represents an interesting research venue, where the results in (Feng and Tesi, 2017) may prove relevant in this regard. Another interesting research line pertains the study of remote control architecture when sensor and actuator channels are independent. In this case, it is interesting to study closed-loop stability in connection with asynchronous DoS.

Appendix A. PROOF OF LEMMA 2

Consider any interval $[z_m, \min\{z_m + h\delta, z_{m+1}\}]$, $m \in \mathbb{N}_0$. The proof is divided into two steps. First, we provide an upper bound on the error dynamics ϕ at the sampling times $z_m + p\delta$. Second, we provide an upper bound on the error dynamics ϕ between inter-samplings. This provides an upper bound on ϕ over the whole interval $[z_m, \min\{z_m + h\delta, z_{m+1}\}]$.

We start by deriving an upper bound on $\|\phi(z_m + p\delta)\|$. It is simple to verify that the dynamics of the variable α related to the predictor equation satisfies

$$\alpha_p(z_m) = A_\delta^p \alpha_0(z_m) + \sum_{k=0}^{p-1} A_\delta^{p-k-1} B_\delta K \alpha_k(z_m) \quad (A.1)$$

for all $p \in \{0, 1, \dots, h-1\}$ such that $z_m + p\delta < z_{m+1}$. Moreover, the process dynamics satisfies

$$\begin{aligned} x(t) &= e^{A(t-z_m)} x(z_m) + \int_{z_m}^t e^{A(t-\tau)} B u(\tau) d\tau \\ &\quad + \int_{z_m}^t e^{A(t-\tau)} d(\tau) d\tau \end{aligned} \quad (A.2)$$

for all $t \in \mathbb{R}_{\geq z_m}$. Combining these two expressions, we get

$$\begin{aligned} \phi(z_m + p\delta) &= \alpha_p(z_m) - x(z_m + p\delta) \\ &= e^{A p\delta} n(z_m) - \int_{z_m}^{z_m + p\delta} e^{A(z_m + p\delta - \tau)} d(\tau) d\tau \end{aligned} \quad (A.3)$$

for all $p \in \{0, 1, \dots, h-1\}$ such that $z_m + p\delta < z_{m+1}$. Here, we exploited the relation $A_\delta^p = e^{A p\delta}$ and the fact that

$$\begin{aligned}
& \int_{z_m}^{z_m+p\delta} e^{A(z_m+p\delta-\tau)} B u(\tau) d\tau \\
&= \sum_{k=0}^{p-1} \left[\int_{z_m+k\delta}^{z_m+(k+1)\delta} e^{A(z_m+p\delta-\tau)} B d\tau \right] K \alpha_k(z_m) \\
&= \sum_{k=0}^{p-1} e^{A\delta(p-k-1)} \left[\int_0^\delta e^{As} B ds \right] K \alpha_k(z_m) \\
&= \sum_{k=0}^{p-1} A_\delta^{p-k-1} B_\delta K \alpha_k(z_m) \tag{A.4}
\end{aligned}$$

where the second equality is obtained using the change of variable $s = z_m + (k+1)\delta - \tau$. We then have

$$\|\phi(z_m + p\delta)\| \leq \rho_1 \|w_t\|_\infty \tag{A.5}$$

where $\rho_1 := (1 + \frac{1}{\mu_A}) e^{\mu_A(h-1)\delta}$.

We can now provide an upper bound on the prediction error ϕ between inter-sampling instants. To this end, observe that the dynamics of ϕ satisfies

$$\begin{aligned}
\dot{\phi}(t) &= -\dot{x}(t) \\
&= A\phi(t) - \Phi\alpha_p(z_m) - d(t), \quad t \in I_p(z_m) \tag{A.6}
\end{aligned}$$

where $\phi(z_m) = n(z_m)$. Let now

$$f(t - z_m - p\delta) := \int_{z_m+p\delta}^t e^{\mu_A(t-\tau)} d\tau \tag{A.7}$$

where $t \in \mathbb{R}_{z_m+p\delta}$. Then, the solution of (A.6) satisfies

$$\begin{aligned}
\|\phi(t)\| &\leq e^{\mu_A(t-z_m-p\delta)} \|\phi(z_m + p\delta)\| \\
&\quad + \kappa_1 f(t - z_m - p\delta) (\|d_t\|_\infty + \|\alpha_p(z_m)\|) \\
&\leq \rho_1 \rho_2 \|w_t\|_\infty + \kappa_1 f(t - z_m - p\delta) \|d_t\|_\infty \\
&\quad + \kappa_1 f(t - z_m - p\delta) (\|\phi(t)\| + \|x(t)\|) \tag{A.8}
\end{aligned}$$

for all $t \in I_p(z_m)$, where we defined $\rho_2 := \max\{e^{\mu_A\delta}, 1\}$ and $\kappa_1 := \max\{\|\Phi\|, 1\}$.

Observe that $f(0) = 0$ and that $f(t - z_m - p\delta)$ is monotonically increasing with t . Thus, any positive real δ such that

$$f(\delta) \leq \frac{1}{\kappa_1} \frac{\sigma}{(1 + \sigma)} \tag{A.9}$$

ensures (21) with $\rho := \sigma + \rho_1 \rho_2 (1 + \sigma)$.

The explicit expression for δ follows from Theorem 1. $\blacksquare$

REFERENCES

Amin, S., Cárdenas, A., and Sastry, S. (2009). Safe and secure networked control systems under denial-of-service attacks. *Hybrid systems: Computation and Control*, 31–45.

Bemporad, A. (1998). Predictive control of teleoperated constrained systems with unbounded communication delays. In *37th IEEE Conference on Decision and Control, Tampa, FL, USA*.

Cetinkaya, A., Ishii, H., and Hayakawa, T. (2016). Networked control under random and malicious packet losses. *IEEE Transactions on Automatic Control*, (99), 1–16.

Chaillet, A. and Bicchi, A. (2008). Delay compensation in packet-switching networked controlled systems. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, 3620–3625. IEEE.

Cheng, P., Datta, A., Shi, L., and Sinopoli, B. (eds.) (2017). Special issue on secure control of cyber physical systems. *IEEE Transactions on Control of Network Systems*.

De Persis, C. and Tesi, P. (2014a). On resilient control of nonlinear systems under denial-of-service. In *Proc. of the IEEE Conference on Decision and Control, Los Angeles, CA, USA*.

De Persis, C. and Tesi, P. (2014b). Resilient control under denial-of-service. In *Proc. of the 19th IFAC World Conference, Cape Town, South Africa*, 134–139.

De Persis, C. and Tesi, P. (2015). Input-to-state stabilizing control under denial-of-service. *IEEE Transactions on Automatic Control*, 60(11), 2930–2944.

Fawzi, H., Tabuada, P., and Diggavi, S. (2011). Secure state-estimation for dynamical systems under active adversaries. *Annual Allerton Conference on Communication, Control, and Computing*.

Feng, S. and Tesi, P. (2016). Resilient control under denial-of-service: Robust design. In *2016 American Control Conference, Boston, MA, USA (extended version available at arXiv:1603.02564)*.

Feng, S. and Tesi, P. (2017). Resilient control under denial-of-service: Robust design. *Automatica*, 79, 42 – 51.

Forni, F., Galeani, S., Nešić, D., and Zaccarian, L. (2010). Lazy sensors for the scheduling of measurement samples transmission in linear closed loops over networks. In *IEEE Conference on Decision and Control and European Control Conference, Atlanta, USA*.

Gupta, A., Langbort, C., and Başar, T. (2010). Optimal control in the presence of an intelligent jammer with limited actions. In *Proc. of the 49th IEEE Conference on Decision and Control, Atlanta, GA, USA*, 1096–1101.

Hardy, G., Littlewood, J., and Polya, G. (1952). *Inequalities*. Cambridge, UK. Cambridge University Press.

Hespanha, J., Naghshtabrizi, P., and Xu, Y. (2007). A survey of recent results in networked control systems. *Proceedings of the IEEE*, 95, 138–162.

Quevedo, D.E. and Nešić, D. (2011). Input-to-state stability of packetized predictive control over unreliable networks affected by packet-dropouts. *IEEE Transactions on Automatic Control*, 56, 370–375.

Quevedo, D.E., Østergaard, J., and Nešić, D. (2011). Packetized predictive control of stochastic systems over bit-rate limited channels with packet loss. *IEEE Transactions on Automatic Control*, 56(12), 2854–2868.

Senejohnny, D., Tesi, P., and De Persis, C. (2015). Self-triggered coordination over a shared network under denial-of-service. In *Proc. of the 54th IEEE Conference on Decision and Control, Osaka, Japan*.

Shisheh Feroosh, H. and Martínez, S. (2013). On multi-input controllable linear systems under unknown periodic dos jamming attacks. In *2013 SIAM Conference on Control and its Applications, San Diego, CA*.

Teixeira, A., Shames, I., Sandberg, H., and Johansson, K. (2015). A secure control framework for resource-limited adversaries. *Automatica*, 51, 135–148.

- Ugrinovskii, V. and Langbort, C. (2014). Control over adversarial packet-dropping communication networks revisited. *arXiv:1403.5641*.
- Xu, W., Trappe, W., Zhang, Y., and Wood, T. (2005). The feasibility of launching and detecting jamming attacks in wireless networks. In *Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, 46–57. ACM.
- Zhang, H., Cheng, P., Shi, L., and Chen, J. (2016). Optimal dos attack scheduling in wireless networked control system. *IEEE Transactions on Control Systems Technology*, 24(3), 843–852.