University of Groningen

Resilient Control Under Denial-of-Service

De Persis, Claudio; Tesi, Pietro

Published in:
Lecture Notes in Control and Information Sciences

DOI:
[10.1007/978-3-030-65048-3_3](https://doi.org/10.1007/978-3-030-65048-3_3)

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
Publisher's PDF, also known as Version of record

Publication date:
2021

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

Citation for published version (APA):

De Persis, C., & Tesi, P. (2021). Resilient Control Under Denial-of-Service: Results and Research Directions. In R. M. G. Ferrari, & A. M. H. Teixeira (Eds.), *Lecture Notes in Control and Information Sciences* (pp. 41-60). (Lecture Notes in Control and Information Sciences; Vol. 486). Springer Science and Business Media Deutschland GmbH. https://doi.org/10.1007/978-3-030-65048-3_3

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.

Chapter 3

Resilient Control Under Denial-of-Service: Results and Research Directions

Claudio De Persis and Pietro Tesi

Abstract The question of security is becoming central for the current generation of engineering systems which more and more rely on networks to support monitoring and control tasks. This chapter addresses the question of designing network control systems that are resilient to Denial-of-Service, that is to phenomena which render a communication network unavailable to use. We review recent results in this area and discuss some of the research challenges.

3.1 Introduction

Security is becoming central for modern engineering systems which more and more rely on networks to support monitoring and control tasks [1]. The main concern is that networks, especially wireless networks, can exhibit unreliable behavior as well as security vulnerabilities, and their malfunctioning can severely affect the systems which our society crucially relies on [2].

Denial-of-Service (DoS) is one of the most common, yet severe, malfunctions that a network can exhibit. By DoS, one usually refers to the phenomenon by which a communication network becomes unavailable to use, meaning that data exchange cannot take place. It is a general term incorporating different types of malfunctions (all causing network unavailability) such as congestion, devices de-authentication and jamming interference [3, 4], and it can be generated by unintentional or intentional sources in which case, the latter, one often uses the term *DoS attacks*. Due to its disruptive effects and common occurrence, DoS has become a central research theme in the context of networked control systems [5].

C. De Persis

University of Groningen, Nijenborgh 4, 9747 AG Groningen, Groningen, Netherlands

e-mail: c.de.persis@rug.nl

P. Tesi (✉)

University of Florence, Via di Santa Marta 3, 50139 Florence, Italy

e-mail: pietro.tesi@unifi.it

© Springer Nature Switzerland AG 2021

R. M. G. Ferrari et al. (eds.), *Safety, Security and Privacy for Cyber-Physical Systems*,

Lecture Notes in Control and Information Sciences 486,

https://doi.org/10.1007/978-3-030-65048-3_3

This chapter addresses the question of designing DoS-resilient networked control systems. The literature on this topic is vast and diverse, and covers linear [6–9], nonlinear [10–12] and distributed systems [13–15]. In this chapter, we will review some of the results in this area, building on the framework developed in [8]. The objective here is not to provide a comprehensive literature review, which is an almost impossible task. Rather, the objective is to discuss, for three macro areas, basic results, and research challenges.

In Sect. 3.2, we consider a centralized framework where controller and plant exchange data through a network which can undergo DoS. For a given controller, we characterize frequency and duration of DoS under which closed-loop stability is preserved. Related to this problem, we review other DoS models considered in the literature, and discuss two open problems in this area: optimality and the role of transmission scheduling. In Sect. 3.3, we focus the attention on the problem of designing control systems that maximize robustness against DoS. We show that this problem has clear connections with the problem of designing finite time observers, and discuss challenges that arise when the control unit is placed remotely from the plant actuators. In Sect. 3.4, we finally consider distributed systems, the area which currently poses most of the research challenges. We first discuss DoS-resilient consensus, which is as a prototypical distributed control problem [16]. Subsequently, we discuss some of the challenges that arise when dealing with networks having more complex dynamics, as well as the problem of identifying critical links in networks with peer-to-peer architecture. The chapter ends with some concluding remarks in Sect. 3.5.

The main focus of this chapter is on linear systems. In fact, while nonlinear systems have their own peculiarities [10–12], most of the issues arising with DoS are shared by linear systems as well. In this chapter, we will mostly consider control problems. Yet, a large and fruitful research line has been developed also for *remote* estimation problems, that is problems in which the objective is to reconstruct the process state through network measurements [17, 18]. In this chapter, the focus is on methods to achieve resilience against DoS. In the context of DoS attacks, research has been carried out also to determine optimal attack scheduling [19, 20]. Albeit not central to our discussion, we will further elaborate on this point in Sect. 3.2 when discussing optimality issues.

We finally point out that DoS is only one of the aspects that affect the security of networked control systems. In the last years, a large amount of research has been carried out on this topic, mostly in connection with security against attacks, which include for instance, *bias injection*, *zero dynamics*, and *replay* attacks. We refer the interested reader to [2, 5, 21] for a general overview of security issues in networked cyber-physical systems.

Fig. 3.1 Schematic representation of the networked control system

3.2 Stability Under Denial-of-Service

3.2.1 Basic Framework

Consider a dynamical system given by

$$\dot{x}(t) = Ax(t) + Bu(t) + w(t), \quad (3.1)$$

where $t \geq 0$ is the time; $x \in \mathbb{R}^{n_x}$ is the state, $u \in \mathbb{R}^{n_u}$ is the control signal, and $w \in \mathbb{R}^{n_x}$ is a disturbance; A and B are matrices with (A, B) stabilizable. The control action is implemented over a communication network, which renders the overall control system a *networked control system*.

Let K be a controller designed in such a way that $\Phi := A + BK$ is Hurwitz (all the eigenvalues of Φ have negative real part), and let $\Delta > 0$ be a constant specifying the desired update rate for the control signal. Ideally, the control signal is then $u_*(t) := Kx(k\Delta)$ for all $t \in [k\Delta, (k+1)\Delta)$ with $k = 0, 1, \dots$, as in classic sampled-data control. Throughout this chapter, the sequence $\{k\Delta\}_{k=0,1,\dots}$ will be referred to as the sequence of transmission times or more simply *transmissions*, that is t is a transmission time if and only if $t = k\Delta$ for some k . Due to the presence of a communication network some of the transmissions can fail, that is $u(t) \neq u_*(t)$. Whenever transmissions fail, we say that the network is under *Denial-of-Service* (DoS). In general, DoS can have a genuine or malicious nature, in which case we refer to DoS *attacks*. A schematic representation of the networked control system is reported in Fig. 3.1.

Remark 3.1 Throughout this chapter, we do not distinguish whether transmissions fail because communication is not possible (for instance, when transmission devices are disconnected from the network) or because data are corrupted (and discarded) due to interference signals [3, 4]. In fact, from a control perspective it is sufficient to interpret DoS as a mechanism inducing packet losses (*cf.* Sect. 3.2.2.1). $\square$

3.2.1.1 Stability in DoS-Free Networks

Even in the ideal situation in which the network is DoS-free, the transmission times must be carefully chosen. The following result addresses this point and is key for the results of Sect. 3.2.2.

Given any positive definite matrix $Q = Q^\top$, let P be the unique solution to the Lyapunov equation

$$\Phi^\top P + P \Phi + Q = 0, \quad (3.2)$$

Let α_1 and α_2 be the smallest and largest eigenvalue of P , respectively. Let γ_1 be the smallest eigenvalue of Q and let $\gamma_2 := \|2PBK\|$. Given a square matrix M , let μ_M be its logarithmic norm, that is $\mu_M := \max\{\lambda \mid \lambda \in \text{spec}\{(M + M^\top)/2\}\}$, and let

$$\bar{\Delta} := \begin{cases} \left(\frac{\sigma}{1+\sigma} \right) \frac{1}{\max\{\|\Phi\|, 1\}}, & \mu_A \leq 0 \\ \frac{1}{\mu_A} \log \left[\left(\frac{\sigma}{1+\sigma} \right) \frac{1}{\max\{\|\Phi\|, 1\}} \mu_A + 1 \right], & \mu_A > 0 \end{cases} \quad (3.3)$$

where $\sigma \in (0, \gamma_1/\gamma_2)$.

Definition 3.1 (cf. [22]) Consider a dynamical system $\dot{x} = f(x, w)$, and let $\mathcal{L}_\infty$ denote the set of measurable locally essentially bounded functions. We say that the system is *input-to-state stable* (ISS) if there exist a $\mathcal{KL}$ -function β and a $\mathcal{K}_\infty$ -function γ such that, for all $x(0)$ and $w \in \mathcal{L}_\infty$,

$$\|x(t)\| \leq \beta(\|x(0)\|, t) + \gamma(\|w\|_\infty) \quad (3.4)$$

for all $t \geq 0$, where $\|w\|_\infty := \sup_{s \geq 0} \|w(s)\|$. If (3.4) holds when $w \equiv 0$, then the system is said to be *globally asymptotically stable* (GAS). $\square$

Lemma 3.1 ([8]) Consider the system Σ given by (3.1) with $u(t) = Kx(k\Delta)$ for all $t \in [k\Delta, (k+1)\Delta)$, $k = 0, 1, \dots$, and where K is such that $\Phi = A + BK$ is Hurwitz. Suppose that the network is DoS-free. Then, Σ is ISS with respect to w for every $\Delta \leq \bar{\Delta}$. $\square$

According to Lemma 3.1, $\bar{\Delta}$ should be then interpreted as an upper bound on the transmission times under which ISS is guaranteed.

3.2.2 Input-to-State Stability Under DoS

We now turn the attention to the question of stability in the presence of DoS.

Let $\{h_n\}$ with $n = 0, 1, \dots$ and $h_0 \geq 0$ be the sequence of DoS *off/on* transitions, that is the time instants at which DoS changes from zero (transmissions succeed)

to one (transmissions fail). Then $H_n := \{h_n\} \cup [h_n, h_n + \tau_n)$ represents the n th DoS time interval, of duration $\tau_n \geq 0$, during which all the transmissions fail. Given non-negative reals τ and t with $t \geq \tau$, the symbol

$$\mathcal{E}(\tau, t) := \bigcup_n H_n \cap [\tau, t] \quad (3.5)$$

represents the subset of the interval $[\tau, t]$, where transmissions fail. Accordingly, $\Theta(\tau, t) := [\tau, t] \setminus \mathcal{E}(\tau, t)$ represents the subset of $[\tau, t]$, where transmissions succeed. Let now $\{s_r\}$ with $r = 0, 1, \dots$ denote the sequence of successful transmissions, that is $t = s_r$ for some r if and only if $t = k\Delta$ for some k and $t \in \Theta(0, t)$. Then the control signal is given by

$$\begin{cases} u(t) = Kz(t) \\ \dot{z}(t) = 0, & t \neq s_r \\ z(t) = x(t), & t = s_r \end{cases} \quad (3.6)$$

with $z(0^-) = 0$. In simple terms, the control signal behaves in a sample-and-hold fashion according to the last successful transmission. Here, the notation $z(0^-) = 0$ implies that the controller has zero initial conditions if no data is received at $t = 0$, that is $z(0) = 0$ if $s_0 > 0$.

The main question now is to determine the amount of DoS that the control system can tolerate before undergoing instability. Such an amount is obviously not arbitrary (the extreme case is when the network is constantly under DoS and no transmission can succeed). The result which follows stipulates that, in order to get stability, both DoS *frequency* and *duration* should be sufficiently small.

Given $\tau, t \in \mathbb{R}_{\geq 0}$ with $t \geq \tau$, let $v(\tau, t)$ be the number of DoS *off/on* transitions occurring on the interval $[\tau, t)$.

Assumption 3.1 (*DoS frequency*). There exist constants $\eta \geq 0$ and $\tau_D > 0$ such that

$$v(\tau, t) \leq \eta + \frac{t - \tau}{\tau_D} \quad (3.7)$$

for all τ and t with $t \geq \tau$. □

Assumption 3.2 (*DoS duration*). There exist $\kappa \geq 0$ and $T > 0$ such that

$$|\mathcal{E}(\tau, t)| \leq \kappa + \frac{t - \tau}{T} \quad (3.8)$$

for all τ and t with $t \geq \tau$. □

Theorem 3.1 ([8]) Consider the system Σ given by (3.1) with control signal (3.6), where K is such that $\Phi = A + BK$ is Hurwitz. Let the inter-transmission time Δ be

chosen as in Lemma 3.1. Then, Σ is ISS for every DoS pattern satisfying Assumption 3.1 and 3.2 with arbitrary η and κ , and with τ_D and T such that

$$\frac{\Delta}{\tau_D} + \frac{1}{T} < \frac{\omega_1}{\omega_1 + \omega_2} \quad (3.9)$$

where $\omega_1 := (\gamma_1 - \gamma_2\sigma)/2\alpha_2$ and $\omega_2 := 2\gamma_2/\alpha_1$, and where $\alpha_1, \alpha_2, \gamma_1$ and γ_2 are as in Lemma 3.1. $\square$

Limiting the DoS frequency and duration is necessary in order to render stability a feasible task. We note in particular that (3.9) requires $T > 1$, otherwise, the network would be always in a DoS status. Limiting the duration of DoS is not sufficient since stability may be destroyed also by DoS patterns with short duration but having high frequency. In this respect, condition $\tau_D > \Delta$ in (3.9) captures the fact that DoS cannot occur at the same rate as the transmission rate. We will further elaborate on the role of the transmission rate in Sect. 3.2.3.1.

3.2.2.1 Models of DoS

Other DoS models have been proposed in the literature, mostly in connection with discrete-time formulations. While these models sometimes originate from different approaches, they all stipulate that, in order to get stability, the DoS action must be constrained in time.

In discrete-time setting, a natural counterpart of Assumptions 3.1 and 3.2 is to require that there exist positive constants $c \geq 0$ and $\lambda > 0$ such that

$$\sum_{k=k_0}^{k_1-1} (1 - \theta_k) \leq c + \frac{k_1 - k_0}{\lambda} \quad (3.10)$$

for all integers k_0 and k_1 with $k_1 > k_0$, where $\theta_k = 0$ when there is DoS at time k and $\theta_k = 1$ otherwise. Similar to (3.10), other formulations focusing on *finite-horizon* control problems [6] consider constraints of the type

$$\sum_{k=0}^T (1 - \theta_k) \leq c \quad (3.11)$$

where T is the control horizon of interest, while *probabilistic* variants of (3.10) have been proposed in [23, 24].

All these models are high-level models in the sense that they abstract away the rule according to which the network undergoes DoS. This approach is useful when there is little knowledge regarding the type of DoS and the network characteristics. When more information is available other models can be used. For instance, [17] considers DoS in wireless networks caused by jamming signals. For this setting, a transmission at time k is successful with probability

$$1 - 2Q\left(\sqrt{\alpha \frac{p_k}{\omega_k + \sigma^2}}\right), \quad (3.12)$$

where $Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\eta^2/2} d\eta$, and α is a parameter. This model dictates that the probability that a transmission at time k is successful depends on the ratio between the transmission power p_k and the interference power ω_k (the source of DoS) which is added to the noise power σ^2 of the channel. Constraints on DoS are expressed in a similar way as (3.11), for example, by imposing $\sum_{k=0}^T \omega_k \leq c$.

A detailed comparison among these several other models has been recently reported in [25, 26], to which the interested reader is referred to. In this respect, it is worth noting that the majority of the DoS models considered in the literature differ from the packet-loss models, for instance, *Bernoulli* models, considered in the classic literature on networked control [27]. The latter, in fact, are more effective in characterizing the quality of the network in normal operating condition, while DoS models account for abnormal situations such as prolonged periods of time where no transmission can succeed.

3.2.3 Research Directions: Scheduling Design and Min–Max Problems

Theorem 3.1 is a prototypical result which shows that network control systems with suitably designed transmission rates enjoy some level of robustness against DoS. The result discussed here has been extended in several directions, which include for instance nonlinear systems and output feedback [10–12, 29], as well as robustness to transmission delay [28] and quantization [29]. While much remains to be done also in these areas, especially for nonlinear systems, in the sequel, we will focus on other aspects which we perceive as much less explored.

3.2.3.1 Transmission Scheduling

The preceding analysis rests on the assumption that the transmission rate is constant. This assumption can be easily relaxed by replacing the constraint $\Delta \leq \bar{\Delta}$ with the constraint $\Delta_k \leq \bar{\Delta}$ for all $k \geq 0$. In this case, Theorem 3.1 continues to hold provided that in (3.9) we replace Δ with $\sup_k \Delta_k$ [8].

This opens the way to the use of a more sophisticated transmission Policies, for instance, *event-triggered* policies [30]. The event-triggered paradigm in particular advocates the idea that transmissions should take place only when strictly needed, and this can play an important role in the context of DoS. In fact, especially for distributed systems (*cf.* Sect. 3.4), aggressive transmission policies can exacerbate DoS by inducing congestion phenomena. Limiting the amount of transmissions can therefore help to maintain a satisfactory network throughput.

It might be argued that low-rate transmissions render stability much more fragile in the sense that with low-rate transmissions stability can be destroyed by low-rate DoS. This fact is captured in Theorem 3.1 where the fulfillment of (3.9) becomes more difficult to satisfy as Δ increases. Yet, at least in the context of DoS *attacks*, the implication “low-rate transmissions $\implies$ more vulnerability” need not apply. In fact, unless an attacker can access the sensor logic, event-triggered logics can render difficult for an attacker to predict when transmissions will take place, thus to learn the transmission policy. The idea of rendering the transmissions less predictable has been explored in [31], where the transmission times are *randomized*, but not in the context of event-triggered control. Understanding the amount of information needed to predict the transmission times associated with an event-triggered logic could lead to the development of control schemes that ensure low-rate transmissions along with low predictability of the transmission times.

3.2.3.2 Optimality

The preceding analysis does not take DoS into account at the stage of designing the controller. In the next section, we will focus on the question of robustness. Hereafter, we make some considerations on the design of optimal control laws.

The design of control laws that are optimal in the presence of DoS is for sure one of the most challenging problems. A simple instance of this problem is as follows. Consider a finite-horizon optimal control problem where the goal is to minimize the desired cost function, say $\sum_{k=0}^T f(x(k), u(k))$. In the presence of DoS, a classic minimization problem of this type turns out to be a “min max” problem in which the objective is to minimize the cost function overall possible DoS patterns within a certain class $\mathcal{C}$ (for instance all DoS patterns such that $\sum_{k=0}^T (1 - \theta_k) \leq c$, where $\theta_k = 0$ when there is DoS at time k and $\theta_k = 1$ otherwise), that is

$$\min_u \max_{\mathcal{C}} \sum_{k=0}^T f(x(k), u(k))$$

Only a few papers have addressed this or similar problems; see for instance [6, 17, 32, 33] for DoS-resilient state estimation. Problems of this type are naturally cast in a game-theoretic framework. The main difficulty is that, depending on the objective function, pure Nash equilibria may not exist.

3.3 Robust Control Design

In the preceding section, we considered a basic formulation in which the problem is to determine the amount of DoS that a given control system can tolerate. In this section, we consider the question of *designing* the control system so as to maximize

robustness against DoS. Later in Sect. 3.3.1, we will discuss some open problems in this research area.

In connection with the model of DoS considered in Sect. 3.2.2, the question of designing robust control systems amounts to searching for control laws that ensure stability for all DoS patterns satisfying

$$\frac{\Delta}{\tau_D} + \frac{1}{T} < \alpha \quad (3.13)$$

with α as closest as possible to 1. We notice that $\alpha = 1$ is the best possible bound since for $\alpha > 1$ there would exist DoS patterns that satisfy (3.13) but for which no control system can guarantee stability. As an example, for $\alpha > 1$ the DoS pattern characterized by $(\tau_D, T) = (\infty, 1)$ satisfies (3.13) but causes the network to be always in a DoS status; as another example, for $\alpha > 1$ the DoS pattern characterized by $(\tau_D, T) = (\Delta, \infty)$ with $h_n = n\Delta, n = 0, 1, \dots$, satisfies condition (3.13) but destroys all the transmissions since the occurrence of DoS is exactly synchronized with the transmission times $t_k = k\Delta, k = 0, 1, \dots$.

3.3.1 Control Schemes Based on Finite-Time Observers

A natural way for increasing robustness against DoS is to equip the controller with a “copy” of the system dynamics so as to compensate for the lack of data, that is to use *observer-based* controllers. To fix the ideas, suppose that the whole state of the system is available for measurements, and consider the following controller (recall that $\{s_r\}$ denotes the sequence of successful transmissions):

$$\begin{cases} u(t) = Kz(t) \\ \dot{z}(t) = Az(t) + Bu(t), & t \neq s_r \\ z(t) = x(t), & t = s_r \end{cases} \quad (3.14)$$

with $z(0^-) = 0$. In simple terms, this controller runs a copy of the system dynamics and its state is reset whenever a new measurement becomes available. Intuitively, in the ideal case where there are no process disturbances a single measurement $x(s_r)$ is sufficient to get stability since, starting from s_r , one has $z(t) \equiv x(t)$. In the sequel, we consider the general case where one only has partial state measurements, assuming that the system to control is disturbance-free. The case of disturbances is discussed later in Sect. 3.3.1.

Consider a stabilizable and observable system

$$\begin{cases} \dot{x}(t) = Ax(t) + Bu(t) \\ y(t) = Cx(t) \end{cases} \quad (3.15)$$

where $y \in \mathbb{R}^{n_y}$ is the output (the measurement signal). Let $\mathcal{S}$ denote the set of successful transmissions, and let $\{v_m\}_{m=0,1,\dots}$ be the sequence of successful transmissions preceded by $\mu - 1$ consecutive successful transmissions, that is such that

$$\{v_m, v_m - \Delta, \dots, v_m - (\mu - 1)\Delta\} \in \mathcal{S} \quad (3.16)$$

where μ is a positive integer. The following result establishes an important property related to the frequency at which consecutive successful transmissions occur, and is independent of the system and the controller.

Lemma 3.2 ([28]) *Consider any DoS pattern satisfying Assumptions 3.1 and 3.2 with*

$$\frac{\Delta}{\tau_D} + \frac{1}{T} < 1 - (\mu - 1)\frac{\Delta}{\tau_D}, \quad (3.17)$$

where Δ is the inter-transmission time and μ is an arbitrary positive integer. Then, $v_0 \leq Q + (\mu - 1)\Delta$ and $v_{m+1} - v_m \leq Q + \Delta$ for all m , where

$$Q := (\kappa + \mu\eta\Delta) \left(1 - \frac{1}{T} - \frac{\mu\Delta}{\tau_D}\right)^{-1}, \quad (3.18)$$

with κ and μ as in Assumptions 3.1 and 3.2. $\square$

Lemma 3.2 essentially says that, for any positive integer μ , if (3.17) holds then we always have μ consecutive successful transmissions. The idea then is to equip the controller with a *finite-time* observer which is able to reconstruct the state of the system in μ steps; in turn, condition (3.17) ensures that the process state will be reconstructed in finite time, enabling the control unit to apply correct control signals even if the network subsequently undergoes large periods of DoS. We will now formalize these considerations.

Let μ denote the observability index of $(C, e^{A\Delta})$ (note that if (C, A) is observable then also $(C, e^{A\Delta})$ is observable for generic choices of Δ), and consider the following controller:

$$\begin{cases} u(t) = Kz(t) \\ \dot{z}(t) = Az(t) + Bu(t), & t \neq v_m \\ z(t) = \zeta(t), & t = v_m \end{cases} \quad (3.19)$$

where

$$\begin{cases} \dot{\zeta}(t) = A\zeta(t) + Bu(t), & t \neq s_r \\ \zeta(t) = \zeta(t^-) + M(y(t) - C\zeta(t^-)), & t = s_r \end{cases} \quad (3.20)$$

with $z(0^-) = \zeta(0^-) = 0$, and where M is selected in such a way that $R^\mu = 0$ with $R := (I - MC)e^{A\Delta}$.¹

¹Note that M always exists if (C, A) is observable. In fact, $R^\mu = 0$ amounts to requiring that

The functioning of (3.19)–(3.20) is as follows. System (3.19) runs a copy of (3.15), and its state is reset to ζ whenever μ consecutive successful transmissions take place. In turn, (3.20) gives a finite-time estimate of x , which is correct after μ consecutive successful transmissions. In case of full state measurements $\mu = 1$, $\{v_m\} = \{s_r\}$ and the controller (3.19)–(3.20) reduces to (3.14).

Theorem 3.2 ([28]) *Consider a stabilizable and observable system as in (3.15) with the controller (3.19)–(3.20). Let Δ be the inter-transmission time. Then, the closed-loop system is GAS for any DoS pattern satisfying Assumptions 3.1 and 3.2 with arbitrary η and κ , and with τ_D and T satisfying*

$$\frac{\Delta}{\tau_D} + \frac{1}{T} < 1 - (\mu - 1)\frac{\Delta}{\tau_D} \quad (3.22)$$

where μ is the observability index of $(C, e^{A\Delta})$. □

When $\mu = 1$, which holds in case of full state measurements and can be enforced if C is a design parameter, (3.22) reduces to the ideal bound $1/T + \Delta/\tau_D < 1$. Notice that in this case, by Lemma 3.2 at least one successful transmission is guaranteed to occur. On the other hand, for any given $\mu > 1$ (C is a problem constraint), one can get close to $1/T + \Delta/\tau_D < 1$ by decreasing Δ , and the limit is only dictated by the maximum transmission rate allowed by the network.

3.3.2 Performant Observers and Packetized Control

The control scheme (3.19)–(3.20) relies on the use of a finite-time observer with state resetting. In this section, we discuss two peculiarities of this control scheme which deserve special attention.

3.3.2.1 Robustness of Finite-Time Observers

Theorem 3.2 relies on a finite-time observer which ensures *fast* state reconstruction. Interestingly, to the best of our knowledge, it is not possible to obtain similar results by means of *asymptotic* observers. This suggests that, as far as stability is concerned, estimation speed is the primary factor.

$$\text{rank} \begin{bmatrix} Ce^{A\Delta} \\ Ce^{2A\Delta} \\ \vdots \\ Ce^{\mu A\Delta} \end{bmatrix} = n_x \quad (3.21)$$

Since $e^{A\Delta}$ is regular, this is equivalent to the fact that $(C, e^{A\Delta})$ is μ -steps observable. The detailed procedure for constructing M can be found for instance in [34, Sect. 5].

In the presence of disturbance or measurement noise, however, using finite-time observers can negatively affect the control system performance. In order to illustrate this point, consider a variant of system (3.15) given by

$$\begin{cases} \dot{x}(t) = Ax(t) + Bu(t) + d(t) \\ y(t) = Cx(t) + n(t) \end{cases} \quad (3.23)$$

where d and n represent disturbance and measurement noise, respectively. One can extend the analysis also to such situation [28, Theorem 1], but the estimation error $e(v_m) = z(v_m) - x(v_m)$ at the times v_m turns out to be

$$e(v_m) = \sum_{k=0}^{\mu-1} R^k M n(z_m - k\Delta) + \sum_{k=0}^{\mu-2} R^k v(z_m - k\Delta),$$

where

$$v(t) := -(I - MC) \int_{t-\Delta}^t e^{A(t-s)} d(s) ds$$

Albeit stability in an ISS sense is preserved, this implies that one can have a large noise amplification, which is a well-known fact for deadbeat observers.

An important investigation in this area concerns the development of observers that guarantee robustness to noise and disturbance while preserving the properties of finite-time observers in the ideal case where noise and disturbance are zero. This research line has recently attracted an independent renewed interest in the context of hybrid systems [35]. Achievements in this area could contribute not only to control problems but also to estimation problems, another extremely active research area in the context of DoS [17, 18, 35].

3.3.2.2 Robustness in Remote Control Architectures

Another peculiarity of the control system (3.19)–(3.20) is that it requires that the control unit is *co-located* with the process actuators, which is needed to continuously update the control signal (Fig. 3.2). In case the control unit is instead placed *remotely* the situation is inevitably more complex.

For remote systems, a possible approach is to emulate co-located architectures through *buffering/packetized control* [36, 37]. In simple terms, the basic idea is that at the transmission times the control unit should transmit not only the current control update but also the predictions of future control updates to be stored at the process side and to be used during the periods of DoS. In [38], for the case of full state measurements, it was shown that the ideal bound $1/T + \Delta/\tau_D < 1$ achievable through co-location becomes

$$\frac{\Delta}{\tau_D} + \frac{1}{T} < 1 - \frac{\omega_2(\kappa + \eta\Delta)}{(\omega_1 + \omega_2)h\Delta - \omega_2\Delta} \quad (3.24)$$

Fig. 3.2 (Left) Remote control framework. (Right) Co-located control framework

where $h \geq 1$ is the buffer size, and all the other quantities are as in Theorem 3.1. The ideal bound is thus recovered as $h \rightarrow \infty$.

Depending on the problem at hand, large values of h might be needed to get close enough to the ideal bound $1/T + \Delta/\tau_D < 1$. This brings the issue of burdening the network traffic due to sending a potentially large amount of data per transmission, which might exacerbate congestion phenomena. To mitigate this issue, it becomes imperative to develop transmission mechanisms able to make a parsimonious use of the network (*cf.* Sect. 3.2.3.1). In this context, *self-triggered* control² represents a promising approach since it enables to jointly computing control input predictions and transmission times at the controller side.

3.4 Distributed Systems

Research on DoS-resilient control systems have originally developed for centralized architectures. The last couple of years have instead witnessed tremendous efforts to extend analysis and design methodologies to distributed systems which are the quintessential form of network systems. Most of the research in this area has been developed for *consensus*-like problems [13–15, 40–43]. In this section, we first present a distributed consensus algorithm which is resilient to DoS, and then we discuss some of the (many) open problems in this area.

²In self-triggered control [39], the update times are not selected on the basis of a continuous monitoring of the process state. Rather, they are based on predictions using previously received data and knowledge of the plant dynamics.

Fig. 3.3 Schematic representation of a distributed consensus algorithm

3.4.1 DoS-Resilient Distributed Consensus

Consider a connected undirected graph $G = (I, E)$, where $I = \{1, 2, \dots, n\}$ is the set of nodes and $E \subseteq I \times I$ is the set of edges. For each node $i \in I$, we denote by $\mathcal{N}_i$ the set of its neighbors, and by d_i its degree, that is, the cardinality of $\mathcal{N}_i$. The consensus problem consists in developing distributed control algorithms with which the nodes, each starting from some initial value, say x_i , eventually converge to a common value by exchanging data with their neighbors (Fig. 3.3).

3.4.1.1 Consensus Algorithm

In this section, we describe the consensus algorithm. We consider the same model of DoS considered in Sect. 3.2. In this respect, we recall that given nonnegative reals τ and t with $t \geq \tau$, we denote by $\mathcal{E}(\tau, t)$ the subset of the interval $[\tau, t]$ during which the network is under DoS, and by $\Theta(\tau, t) := [\tau, t] \setminus \mathcal{E}(\tau, t)$ the subset of $[\tau, t]$ during which the network is DoS-free.

Each node is modelled as a hybrid dynamical system Σ_i that obeys continuous evolution and discrete updates, which occur when the node communicates with its neighbors. The continuous evolution of node i is given by

$$\begin{cases} \dot{x}_i(t) = u_i(t) \\ \dot{u}_i(t) = 0 \\ \dot{\theta}_i(t) = -1 \end{cases} \quad (3.25)$$

where $x_i, u_i, \theta_i \in \mathbb{R}, i \in I$. When the clock variable θ_i reaches zero, a discrete update occurs: node i polls its neighbors and update its control signal according to

$$\begin{cases} x_i(t) = x_i(t^-) \\ u_i(t) = \begin{cases} \text{sign}_\varepsilon(\text{ave}_i(t)), & t \in \Theta(0, t) \\ 0, & \text{otherwise} \end{cases} \\ \theta_i(t) = \begin{cases} f_i(t), & t \in \Theta(0, t) \\ \frac{\varepsilon}{4d_i}, & \text{otherwise} \end{cases} \end{cases} \quad (3.26)$$

Here, the function $\text{sign}_\varepsilon : \mathbb{R} \rightarrow \{-1, 0, +1\}$ is given by $\text{sign}(z)$ if $|z| \geq \varepsilon$ and zero otherwise, where $\varepsilon > 0$ is a sensitivity parameter which is used at the design stage to trade-off frequency of the control updates vs. accuracy of the consensus region; the function $\text{ave}_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is given by

$$\text{ave}_i(t) := \sum_{j \in \mathcal{N}_i} (x_j(t) - x_i(t)). \quad (3.27)$$

and represents the local average that node i forms with its neighbors; finally, the function $f_i : \mathbb{R}^n \rightarrow \mathbb{R}_{>0}$ is given by

$$f_i(x(t)) := \begin{cases} \frac{|\text{ave}_i(t)|}{4d_i}, & |\text{ave}_i(t)| \geq \varepsilon \\ \frac{\varepsilon}{4d_i}, & \text{otherwise} \end{cases}$$

In simple terms, at each update time node i polls its neighbors. If the transmission succeeds (there is no DoS) then node i computes its own local average that provides information on the nodes disagreement, and the control law is updated accordingly; otherwise, the control signal is set to zero, meaning that node i remains at its current value. At the same time, node i computes, through the function f_i , the next time instant at which an update will occur. For this reason, the control logic is referred to as *self-triggered*. The use of this logic for consensus in DoS-free networks has been first proposed in [44].

3.4.1.2 Resilience Against DoS

The result which follows characterizes the robustness properties of (3.25)–(3.26) in the presence of DoS. Let

$$\mathcal{E} := \{x \in \mathbb{R}^n : |\sum_{j \in \mathcal{N}_i} (x_j - x_i)| < \varepsilon \ \forall i \in I\} \quad (3.28)$$

Theorem 3.3 ([13]) *Consider a connected undirected graph where the nodes follow the logic (3.25)–(3.26). Consider any DoS pattern satisfying Assumptions 3.1 and 3.2 with η and κ arbitrary, and with τ_D and T satisfying*

$$\frac{\Delta}{\tau_D} + \frac{1}{T} < 1 \quad (3.29)$$

where $\Delta := \varepsilon/(4d_{\min})$ with $d_{\min} := \min_{i \in I} d_i$. Then, for every initial condition, the nodes converge in finite time to a point belonging to the set $\mathcal{E}$ in (3.28). $\square$

This consensus algorithm has some interesting features that make it appealing in distributed control: (i) it is fully distributed, also with respect to the clocks of the nodes which do not have to be synchronized; (ii) it achieves finite-time convergence, where the accuracy of consensus depends on a design parameter ε which can be used to trade-off frequency of the control updates and consensus accuracy; (iii) it relies on a transmission logic in which the updates take place only when strictly needed. Like in event-based control, this feature helps to reduce the communication burden which is especially important in distributed settings.

Other features of the algorithm are discussed in the section which follows.

3.4.2 Complex Network Systems and Critical Links

Developing DoS-resilient distributed control algorithms is probably the topic where most of the research challenges are concentrated. In the sequel, we will discuss two important topics where results are lacking.

3.4.2.1 Networks with Complex Dynamics

As mentioned at the beginning of Sect. 3.4, most of the research in this area has been developed for consensus-like problems [13–15, 40–43]. Problems of this type are somehow “manageable” in the sense that they involve systems with *stable* or *neutrally stable* dynamics (like integrators in the context of consensus), which considerably simplifies analysis and design. For instance, in the consensus algorithm discussed in the previous section, one takes advantage of the fact that the dynamics are integrators. This makes it possible to “stop” the state evolution in the presence of DoS, which is instrumental to prevent the nodes from drifting away, simply by zeroing the control input. This is in general not possible with more complex network dynamics.

Networks having more complex (even linear time-invariant) dynamics arise in many other distributed control problems, for instance, in the context of distributed stabilization of large-scale systems where the dynamics are those associated to the *physical* systems to control (rather than deriving from the control algorithm). As an example, consider a network of physically coupled systems with dynamics

$$\dot{x}_i(t) = A_i x_i(t) + B_i u_i(t) + \sum_{j \in \mathcal{N}_i} H_{ij} x_j(t),$$

where x_i is state of subsystem i , u_i is its local control signal, and where H_{ij} defines how subsystem i physically interacts with neighboring processes. A communication network is used to enable the design of distributed control laws

$$u_i(t) = K_i x_i(t_k^i) + \sum_{j \in \mathcal{N}_i} L_{ij} x_j(t_k^j)$$

that should regulate the state of each subsystem to zero, where t_k^i is the k th update time of the control law for subsystem i .

The problem of designing DoS-resilient control schemes for this class of systems has been preliminary studied in [45]. Compared with consensus problems, however, the analysis becomes more complex and the stability conditions more conservative, requiring the subsystems to satisfy suitable *small-gain* properties on their couplings. The reason is that during DoS it is in general not possible to “stop” the evolution of x_i (as done in consensus) since the evolution of x_i does not depend solely from u_i but also on the various x_j . As a consequence, one needs strong conditions on the coupling matrices H_{ij} in order to ensure that the subsystems do not get far from the origin during DoS. This is an example of distributed control problems where even analysis tools are largely lacking.

3.4.2.2 Determining Critical Links

The consensus problem considered in Sect. 3.4.1, assumes that DoS simultaneously affects all the network links. This assumption is reasonable for networks where the data exchange is carried through a single access point. For *peer-to-peer* networks, this assumption need not be realistic. Concerning the consensus problem, it is possible to show that a result analogous to Theorem 3.3 holds provided that condition (3.29) is replaced with

$$\delta^{ij} := \frac{\Delta}{\tau_D^{ij}} + \frac{1}{T^{ij}} < 1 \quad (3.30)$$

where τ_D^{ij} and T^{ij} characterize DoS frequency and duration affecting the link (i, j) . This result was proven in [13]. Even more, the same conclusions continue to hold even when $\delta^{ij} \geq 1$ for some network links (meaning that communication over the link (i, j) is never possible). This happens whenever removing such links does not cause the graph to be disconnected. Specifically, if X is any set of links such that $G_X := (I, E \setminus X)$ remains connected, then consensus is preserved whenever $\delta^{ij} < 1$ for all $(i, j) \in E \setminus X$; see [13].

For consensus-like networks, one can introduce a simple notion of “critical” links as the links (or the minimum number of links) causing the network to disconnect, and one can identify such links by using classic tools like the Stoer–Wagner *mincut* algorithm [46]. For networks involving more complex dynamics such as the one

discussed in the previous subsection, the situation is instead much more involved. In particular, the loss of a link may render the network unstable even if the underlying graph remains connected. Developing efficient methods to identify and minimize the number of critical links through topology and control design is another key aspect to achieve DoS resilience.

3.5 Conclusions

In this chapter, we reviewed some recent results on DoS-resilient networked control. While much has been done in this area, there remain several problems of paramount importance which are yet not fully understood. We mention in particular the design of DoS-resilient optimal control laws, the design of robust control laws for remote control architectures and the design of DoS-resilient distributed control algorithms, the latter being the area where most of the results are lacking.

The present discussion is by no means exhaustive. We refer the interested reader to [2, 5, 21] for additional references on this topic, as well as for a more general overview of security issues in networked systems.

References

1. Lee, E.: Cyber physical systems: design challenges. In: IEEE International Symposium on Object Oriented Real-Time Distributed Computing (ISORC) (2008)
2. Teixeira, A., Shames, I., Sandberg, H., Johansson, K.H.: A secure control framework for resource-limited adversaries. *Automatica* **51**, 135–148 (2015)
3. Bicakci, K., Tavli, B.: Denial-of-service attacks and counter-measures in IEEE 802.11 wireless networks. *Comput. Stand. Interf.* **31**, 931–941 (2009)
4. Pelechrinis, K., Iliofotou, M., Krishnamurthy, S.: Denial of service attacks in wireless networks: the case of jammers. *IEEE Commun. Surv. Tutor.* **13**, 245–257 (2011)
5. Lun, Y., D’Innocenzo, A., Smarra, F., Malavolta, I., Di Benedetto, M.: State of the art of cyber-physical systems security: an automatic control perspective. *J. Syst. Softw.* **149**, 174–216 (2019)
6. Amin, S., Cárdenas, A., Sastry, S.: Safe and secure networked control systems under denial of-service attacks. In: *Hybrid Systems: Computation and Control*, pp. 31–45 (2009)
7. Shisheh Feroosh, H., Martínez, S.: On event-triggered control of linear systems under periodic Denial-of-Service attacks. In: *IEEE Conference on Decision and Control*, Maui, HI, USA (2012)
8. De Persis, C., Tesi, P.: Input-to-state stabilizing control under denial-of-service. *IEEE Trans. Autom. Control* **60**, 2930–2944 (2015)
9. Lu, A., Yang, G.: Input-to-state stabilizing control for cyber-physical systems with multiple transmission channels under denial-of-service. *IEEE Trans. Autom. Control* **63**, 1813–1820 (2018)
10. De Persis, C., Tesi, P.: Networked control of nonlinear systems under denial-of-service. *Syst. Control Lett.* **96**, 124–131 (2016)
11. Dolk, V., Tesi, P., De Persis, C., Heemels, W.: Event-triggered control systems under denial-of-service attacks. *IEEE Trans. Control Netw. Syst.* **4**, 93–105 (2016)

12. Kato, R., Cetinkaya, A., Ishii, H.: Stabilization of nonlinear networked control systems under denial-of-service attacks: a linearization approach 2019 American Control Conference, Philadelphia, PA, USA (2019)
13. Senejohnny, D., Tesi, P., De Persis, C.: A jamming-resilient algorithm for self-triggered network coordination. *IEEE Trans. Control Netw. Syst.* **5**, 981–990 (2017)
14. Lu, A., Yang, G.: Distributed consensus control for multi-agent systems under denial-of-service. *Inf. Sci.* **439**, 95–107 (2018)
15. Yang, H., Li, Y., Dai, L., Xia, Y.: MPC-based defense strategy for distributed networked control systems under DoS attacks. *Syst. Control Lett.* **128**, 9–18 (2019)
16. Nowzari, C., Garcia, E., Cortés, J.: Event-triggered communication and control of networked systems for multi-agent consensus. *Automatica* **105**, 1–27 (2019)
17. Li, Y., Shi, L., Cheng, P., Chen, J., Quevedo, D.: Jamming attacks on remote state estimation in cyber-physical systems: a game-theoretic approach. *IEEE Trans. Autom. Control* **60**, 2831–2836 (2015)
18. Li, Y., Quevedo, D., Dey, S., Shi, L.: SINR-based DoS attack on remote state estimation: a game-theoretic approach. *IEEE Trans. Control Netw. Syst.* **4**, 632–642 (2017)
19. Zhang, H., Cheng, P., Shi, L., Chen, J.: Optimal DoS attack scheduling in wireless networked control system. *IEEE Trans. Control Syst. Technol.* **24**, 843–852 (2016)
20. Zhang, H., Cheng, P., Shi, L., Chen, J.: Optimal denial-of-service attack scheduling with energy constraint. *IEEE Trans. Autom. Control* **60**, 3023–3028 (2015)
21. Mehran Dibaji, S., Pirani, M., Bezalel Flamholz, D., Annaswamy, A., Johansson, K., Chakraborty, A.: A systems and control perspective of CPS security. *Ann. Rev. Control* **47**, 394–411 (2019)
22. Sontag, E.: Input to state stability: basic concepts and results. *Nonlinear Optim. Control Theory Lect. Notes Math.* **163–220**, 2008 (1932)
23. Cetinkaya, A., Ishii, H., Hayakawa, T.: Event-triggered control over unreliable networks subject to jamming attacks. In: *IEEE Conference on Decision and Control*, Osaka, Japan (2015)
24. Cetinkaya, A., Ishii, H., Hayakawa, T.: A probabilistic characterization of random and malicious communication failures in multi-hop networked control. *SIAM J. Control Optim.* **56**, 3320–3350 (2018)
25. De Persis, C., Tesi, P.: A comparison among deterministic packet-dropouts models in networked control systems. *IEEE Control Syst. Lett.* **2**, 109–114 (2017)
26. Cetinkaya, A., Ishii, H., Hayakawa, T.: An overview on denial-of-service attacks in control systems: attack models and security analyses. *Entropy* **21**(210), 1–29 (2019)
27. Hespanha, J., Naghshtabrizi, P., Xu, Y.: A survey of recent results in networked control systems. *Proc. IEEE* **95**, 138–162 (2007)
28. Feng, S., Tesi, P.: Resilient control under denial-of-service: robust design. *Automatica* **79**, 42–51 (2017)
29. Wakaiki, M., Cetinkaya, A., Ishii, H.: Quantized output feedback stabilization under DoS attacks. In: *2018 American Control Conference*, Milwaukee, WI, USA (2018)
30. Tabuada, P.: Event-triggered real-time scheduling of stabilizing control tasks. *IEEE Trans. Autom. Control* **52**, 1680–1685 (2007)
31. Cetinkaya, A., Kikuchi, K., Hayakawa, T., Ishii, H.: Randomized transmission protocols for protection against jamming attacks in multi-agent consensus. [arXiv:1802.01281](https://arxiv.org/abs/1802.01281)
32. Gupta, A., Langbort, C., Başar, T.: Optimal control in the presence of an intelligent jammer with limited actions. In: *IEEE Conference on Decision and Control*, Atlanta, GA, USA (2010)
33. Befekadu, G., Gupta, V., Antsaklis, P.: Risk-sensitive control under a class of denial-of-service attack models. In: *American Control Conference*, San Francisco, CA, USA (2011)
34. O'Reilly, J.: *Observers for Linear Systems*. Academic Press: Mathematics in Science & Engineering, London (1983)
35. Li, Y., Sanfelice, R.: A finite-time convergent observer with robustness to piecewise-constant measurement noise. *Automatica* **57**, 222–230 (2015)
36. Chaillet, A., Bicchi, A.: Delay compensation in packet-switching networked controlled systems. In: *IEEE Conference on Decision and Control*, Cancun, Mexico (2008)

37. Quevedo, D., Ørstergaard, J., Nešić, D.: Packetized predictive control of stochastic systems over bit-rate limited channels with packet loss. *IEEE Trans. Autom. Control* **56**, 2854–2868 (2011)
38. Feng, S., Tesi, P.: Networked control systems under denial-of-service: co-located versus remote architectures. *Syst. Control Lett.* **108**, 40–47 (2017)
39. Heemels, W., Johansson, K., Tabuada, P.: An introduction to event-triggered and self-triggered control. In: *IEEE Conference on Decision and Control*, Maui, Hawaii, USA (2012)
40. Feng, Z., Hu, G.: Distributed secure average consensus for linear multi-agent systems under DoS attacks. In: *American Control Conference*, Seattle, WA, USA (2017)
41. Kikuchi, K., Cetinkaya, A., Hayakawa, T., Ishii, H.: Stochastic communication protocols for multi-agent consensus under jamming attacks. In: *IEEE Conference on Decision and Control*, Melbourne, Australia (2017)
42. Senejohnny, D., Tesi, P., De Persis, C.: Resilient self-triggered network synchronization. In: Tarbouriech, S., Girard, A., Hetel, L. (eds), *Control Subject to Computational and Communication Constraints*. *Lecture Notes in Control and Information Sciences*, vol. 475. Springer, Cham (2018)
43. Amini, A., Azarbahram, A., Mohammadi, A., Asif, A.: Resilient event-triggered average consensus under denial of service attack and uncertain network. In: *6th International Conference on Control, Decision and Information Technologies (CoDIT)*, Paris, France (2019)
44. De Persis, C., Frasca, P.: Robust self-triggered coordination with ternary controllers. *IEEE Trans. Autom. Control* **58**, 3024–3038 (2013)
45. Feng, S., Tesi, P., De Persis, C.: Towards stabilization of distributed systems under denial-of-service. In: *IEEE Conference on Decision and Control*, Melbourne, Australia (2017)
46. Stoer, M., Wagner, F.: A simple min-cut algorithm. *J. ACM* **44**, 585–591 (1997)