Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Denial of Service Attacks on Control Systems with Packet Loss

William Casbolt * **Iñaki Esnaola** *,** **Bryn Jones** *

* *Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield S1 3JD, UK, (e-mail: {wgcasbolt1, esnaola, b.l.jones}@sheffield.ac.uk).*
** *Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA.*

Abstract:
The performance of control systems with packet loss as a result of an attack over the actuation communication channel is analysed. The operator is assumed to monitor the state of the channel by measuring the average number of packet losses and an attack detection criteria is established based on the statistic. The performance of the attacker is measured in terms of the increase of the linear quadratic cost function of the operator subject to a given detection constraint. Within that setting, the optimal denial of service (DoS) attack strategy is formulated for UDP-like and TCP-like communication protocols. For both communication protocols, DoS attack constructions that are independent and identically distributed (IID) are compared to those that are non-stationary. The main contributions of this paper are (i) explicit characterisation of the expected cost increase of the optimal attack constructions and the associated packet loss parameter for the IID case, (ii) proof, by example, that non-stationary random attacks outperform IID attacks in the presence of detection constraints.

Keywords: Secure networked control systems; control and estimation with data loss; control under communication constraints.

1. INTRODUCTION

The introduction of advanced sensing and communication capabilities to control systems gives rise to vulnerabilities that can be exploited with a malicious intent by an attacker (Colbert and Kott (2016)). While the security challenges that control systems face are multifaceted and of diverse nature, the simplicity of implementation of Denial of Service (DoS) attacks in field zones and control zones makes them particularly suitable to exploit software and hardware faults. In this paper we study DoS attacks over control systems that experience packet losses over an actuator communication channel. To account for packet loss, we consider the two protocols proposed in Schenato et al. (2007). The first one in which the packet loss is not monitored, termed UDP-like for its similarity to the communication protocol. The second protocol is termed as TCP-like and monitors the packet loss realisation by sending a packet receipt acknowledgement message back to the receiver. Both protocols and the systems they constitute are depicted in Fig. 1 and Fig. 2. In the literature these are termed *like* due to previous transmissions being monitored and not retransmitted (Schenato et al. (2007)), (Sinopoli et al. (2005)), (Sinopoli et al. (2008)), (Mo et al. (2013)). In Zhang et al. (2016) the authors consider a deterministic DoS attack strategy on a system with power constraints whereas we consider a random DoS attack construction that operates within a no-detection region using an MPC formulation. We first study attack sequences that are constructed as an independent and identically distributed (IID) process. The rationale for this attack construction stems from the simplicity of the attack implementation and the robust attack performance for a wide range of system parameters. We then propose a non-stationary random attack construction and show that for some systems their dynamics can be exploited by the attacker to improve upon the IID construction.

⋆ This work is supported by Rolls-Royce, ESPRC, and The Control, Monitoring and Systems Engineering UTC at The University of Sheffield.

2. PLANT MODEL

We consider the plant model given by

$$X_{k+1} = \mathbf{A}X_k + \mathbf{B}\mathbf{V}_k U_k + W_k, \tag{1}$$

where $\mathbf{A} \in \mathbb{R}^{n\times n}$ is the dynamics matrix, $X_k \in \mathbb{R}^n$ describes the state of the plant at time step $k \in \mathbb{N}$, $\mathbf{B} \in \mathbb{R}^{n\times m}$ is the control matrix, $U_k \in \mathbb{R}^m$ is the vector of control inputs at the k-th time step, $W_k \in \mathbb{R}^n$ is the process noise modelled by a Gaussian distributed vector of random variables with mean $\mathbf{0} \in \mathbb{R}^n$ and covariance matrix $\Sigma_W \in S^n_{++}$ where, S^n_{++} is the set of $n \times n$ symmetric positive definite matrices, and $\mathbf{V}_k \in S^m_+$ is the packet loss variable where the i-th diagonal entry is an IID Bernoulli random variable with mean μ_i. We assume that the current state of the plant is determined by the vector of Gaussian distributed random variables X_k with mean $\overline{X} \in \mathbb{R}^n$ and covariance matrix $\Sigma_X \in S^n_{++}$.

In this paper, we adopt the MPC formulation used in Casbolt et al. (2019) to describe the plant model in (1) over the prediction horizon $N \in \mathbb{N}_+$, resulting in the prediction model given by

$$\chi_k = \Phi X_k + \Gamma \nu_k \Upsilon_k + \Lambda \Xi_k, \tag{2}$$

where $\Phi \in \mathbb{R}^{Nn\times n}$ is the dynamics matrix over the prediction horizon, $\chi_k \in \mathbb{R}^{Nn}$ is the state prediction vector, $\Gamma \in \mathbb{R}^{Nn\times Nm}$ is

Copyright lies with the authors

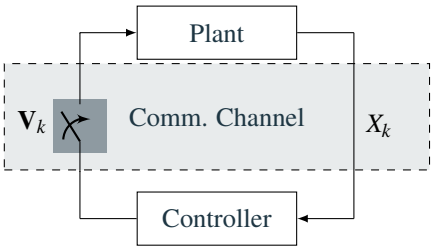

Fig. 1. The UDP-like protocol where the realisation of the packet loss $\mathbf{V}_k$ is not monitored

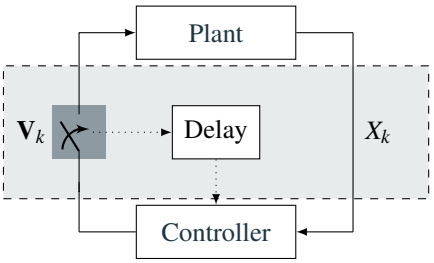

Fig. 2. The TCP-like protocol where the realisation of the packet loss $\mathbf{V}_k$ is transmitted back to the controller.

the propagation matrix for the control law over the prediction horizon, $\Upsilon_k \in \mathbb{R}^{mN}$ is the realisation at the k-th time step of the control law, $\Lambda \in \mathbb{R}^{Nn\times Nn}$ is the propagation matrix for the process noise, $\Xi_k \in \mathbb{R}^{Nn}$ is the process noise over the prediction horizon, and ν_k is a diagonal matrix with the Bernoulli random variables describing the packet losses over the prediction horizon along the diagonal. All the terms in (2) are presented in (6). Due to the lossy communication between the controller and the plant, the operator implements a communication protocol to monitor the state of the packets transmitted to the plant. We adopt the two protocol paradigms proposed by Schenato et al. (2007), namely a UDP-like protocol that does not monitor the channel and a TCP-like protocol that acknowledges receipt of the packet from the controller by sending an *acknowledgement* message to the controller over an auxiliary channel. The difference between both protocol paradigms is depicted in Fig. 1 and Fig. 2, which show the UDP-like and the TCP-like protocols, respectively. The information set available to the controller is determined by the choice of the protocol. We define the information sets as

$$\mathscr{I}_k = \begin{cases} \mathscr{F}_k = \left\{X^k, \mathbf{V}^{k-1}\right\}, & \text{TCP-like,} \\ \mathscr{G}_k = \left\{X^k\right\}, & \text{UDP-like,} \end{cases} \tag{3}$$

where $\mathbf{V}^{k-1} = \{\mathbf{V}_0, \mathbf{V}_1, \ldots, \mathbf{V}_{k-1}\}$, $X^k = \{X_0, X_1, \ldots, X_k\}$ and all sets are monotonically increasing, i.e. there is a filtration such that $\mathscr{I}_k \subseteq \mathscr{I}_{k+1}$. It is shown in Casbolt et al. (2019) that for both protocols the optimal control law is determined by the mean of the packet loss variable $\bar{\nu} \triangleq \mathbb{E}[\nu_k]$. Following in the steps of Schenato et al. (2007), the performance of the controller is characterised by a linear quadratic Gaussian (LQG) cost function. The description of the cost function in the MPC framework proposed in Casbolt et al. (2019) is the cost function

$$J^*(\mathscr{I}_k) \triangleq \min_{\Upsilon_k} \left\{\mathbb{E}\left[X_k^\mathsf{T}\mathbf{Q}X_k + \chi_k^\mathsf{T}\Omega\chi_k + \Upsilon_k{}^\mathsf{T}\nu_k^\mathsf{T}\Psi\nu_k\Upsilon_k \middle| \mathscr{I}_k\right]\right\}, \tag{4}$$

where $\Omega \in S_{++}^{Nn}$ is the state penalty diagonal matrix, $\Psi \in S_{++}^{Nm}$ is the input penalty diagonal matrix, and the diagonal matrix $\mathbf{Q} \in S_{++}^{n}$. The optimal control law for (4) is obtained in Casbolt et al. (2019) for each protocol and shown to be

$$\Upsilon^*_{k|\mathscr{I}_k} = \begin{cases} \Upsilon^*_{k|\mathscr{F}_k} = -\left(\Psi + \Delta^\Gamma\bar{\nu}\right)^{-1}\mathbf{F}X_k, & \text{TCP-like,} \\ \Upsilon^*_{k|\mathscr{G}_k} = -\left(\Psi + \Delta^\Gamma\bar{\nu} + \left(\mathbf{I}\odot\Delta^\Gamma\right)(1-\bar{\nu})\right)^{-1}\mathbf{F}X_k, & \text{UDP-like,} \end{cases} \tag{5}$$

where $\odot$ is the element-wise Hadamard product.

2.1 Attack Model

The performance of the controller is determined by the mean of the packet losses in the actuation channel. In view of this, we study the security risk posed by an attacker that governs the statistics of the packet losses on the actuation channel. In practice, this can be achieved by the attacker via DoS attacks over the communication channel. We are not concerned with the particular implementation of the DoS attacks, instead we study the packet loss attack strategy that aims to disrupt the operation of the controller. In particular, we consider the case in which the attacker constructs the attack sequence by designing a random distribution. The rationale for this stems from the fact that the operator expects the packet losses to be IID, and therefore, the attacker mimics the *nominal* operation of the channel. That being the case, the optimal attack construction is characterised by the probability of packet loss in the actuation channel, described by the diagonal matrix $\mathbf{V}_k^\alpha \in S_{++}^m$ where the i-th diagonal entry is an IID Bernoulli random variable with mean μ_i^α.

To achieve this, the attacker has knowledge of the information set given by

$$\mathscr{A}_k = \{\mathbf{A}, \mathbf{B}, \Sigma_W, \bar{\nu}, \Omega, \Psi, \mathscr{I}_k\}. \tag{7}$$

It is shown later that knowledge of the state of the plant is not necessary to construct the optimal attack and is only required to compute the cost induced by the attack for a particular realisation of the state variables. The controller operates under the assumption that the packet losses over the actuation channel are IID with a mean defined by $\mathbf{M} \triangleq \mathbb{E}[\mathbf{V}_k]$ for $k \in \mathbb{N}$, with $\mathbf{M} \in S_{++}^m$. By changing the statistics of the actuation channel, the attacker induces a different distribution over the sequence of packet losses. To distinguish the induced case from the nominal case the induced sequence of packet losses is defined as the diagonal matrix $\mathbf{V}_k^\alpha$ as above. Similarly, the channel is characterised by $\mathbf{M}^\alpha \triangleq \mathbb{E}[\mathbf{V}_k^\alpha]$ for $k \in \mathbb{N}$, with $\mathbf{M}^\alpha \in S_{++}^m$. Note that the mean does not depend on the time step k, and therefore, the sequence of random variables describing the packet loss in the i-th position is IID. The sequence of packet losses over the prediction horizon is described by the diagonal matrix ν_k^α with the Bernoulli sequences along the diagonal and $\bar{\nu}^\alpha \triangleq \mathbb{E}[\nu_k^\alpha]$ for $k = 1, 2, \ldots, N$.

The objective of the attacker, in contrast to the objective of the controller, is to maximise the cost function (4). The cost function of the attacker is

$$J_A(\mathscr{A}_k) \triangleq \min_{\Upsilon^*_{k|\mathscr{I}_k}} \left\{\mathbb{E}\left[X_k^\mathsf{T}\mathbf{Q}X_k + \chi_k^\mathsf{T}\Omega\chi_k + \Upsilon_k^\mathsf{T}\nu_k^{\alpha\mathsf{T}}\Psi\nu_k^\alpha\Upsilon_k \middle| \mathscr{A}_k\right]\right\},$$

and the optimal attack construction is defined as

$$J_A^*(\mathscr{A}_k) \triangleq \max_{\bar{\nu}^\alpha} \left\{J_A(\mathscr{A}_k)\right\}. \tag{8}$$

Note that, the cost function of the operator in (4) is nested inside (8), i.e. the attacker chooses the worst case packet loss mean

$$\underbrace{\begin{pmatrix} X_{k+1} \\ X_{k+2} \\ \vdots \\ X_{k+N} \end{pmatrix}}_{\chi_k} = \underbrace{\begin{pmatrix} \mathbf{A} \\ \mathbf{A}^2 \\ \vdots \\ \mathbf{A}^N \end{pmatrix}}_{\Phi} X_k + \underbrace{\begin{pmatrix} \mathbf{B} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{AB} & \mathbf{B} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{A}^{N-1}\mathbf{B} & \dots & \mathbf{AB} & \mathbf{B} \end{pmatrix}}_{\Gamma} \underbrace{\begin{pmatrix} \mathbf{V}_k & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_{k+1} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{V}_{k+N-1} \end{pmatrix}}_{\nu_k} \underbrace{\begin{pmatrix} U_k^n \\ U_{k+1}^n \\ \vdots \\ U_{k+N-1}^n \end{pmatrix}}_{\Upsilon_k} + \underbrace{\begin{pmatrix} \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{A} & \mathbf{I} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{A}^{N-1} & \dots & \mathbf{A} & \mathbf{I} \end{pmatrix}}_{\Lambda} \underbrace{\begin{pmatrix} W_k \\ W_{k+1} \\ \vdots \\ W_{k+N-1} \end{pmatrix}}_{\Xi_k} \tag{6}$$

under the assumption that the controller operates optimally. The state estimation performed by the attacker accounts for the true statistics of the actuation channel to produce the state prediction

$$\hat{\chi}_k^\alpha \triangleq \mathbb{E}\left[\chi_k \middle| \mathscr{A}_k\right] = \Phi X_k + \Gamma \bar{\nu}^\alpha \Upsilon_k, \tag{9}$$

and the state error prediction of the attack construction for the two protocols is given by

$$\mathbf{E}_{k|\mathscr{F}_k}^\alpha \triangleq \chi_k - \mathbb{E}\left[\chi_k \middle| \mathscr{F}_k, \nu_k^\alpha\right] = \Lambda \Xi_k, \tag{10a}$$

$$\mathbf{E}_{k|\mathscr{G}_k}^\alpha \triangleq \chi_k - \mathbb{E}\left[\chi_k \middle| \mathscr{G}_k\right] = \Gamma(\nu_k^\alpha - \bar{\nu}^\alpha)\Upsilon_k + \Lambda \Xi_k, \tag{10b}$$

for the TCP-like protocol and the UDP-like protocol, respectively. Note that the TCP-like prediction includes the knowledge of ν_k^α, this is the same knowledge the operator has when calculating this expectation. Specifically, the acknowledgement link removes the effect of the actuation from the estimation. Even though at time step k the realisation of $\mathbf{V}_{k+1}$ is not known, it is known at time step $k+1$, and therefore, the operator knows that any chosen actuation law does not affect the estimation at future time steps as this contribution can be removed from the estimation. This is not possible if the acknowledgement channel is not a perfect channel. The proof of the error trajectories is analogous to the proof in Casbolt et al. (2019) and is omitted. We describe the attack induced cost by rewriting (8) in terms of the state prediction and the state prediction error as in Casbolt et al. (2019) which yields

$$\begin{aligned} J_A^*(\mathscr{A}_k) = \max_{\bar{\nu}^\alpha} \Bigg\{ & \min_{\Upsilon_{k|\mathscr{A}_k}^*} \Big\{ \mathbb{E}\Big[X_k^\mathsf{T} \mathbf{Q} X_k + \Upsilon_k^\mathsf{T} \nu_k^{\alpha\mathsf{T}} \Psi \nu_k^\alpha \Upsilon_k \\ & + (\hat{\chi}_k^\alpha + \mathbf{E}_k^\alpha)^\mathsf{T} \Omega (\hat{\chi}_k^\alpha + \mathbf{E}_k^\alpha) \Big| \mathscr{A}_k \Big] \Big\} \Bigg\}, \\ = X_k^\mathsf{T} \left(\mathbf{Q} + \Delta^\Phi\right) X_k + \max_{\bar{\nu}^\alpha} \Bigg\{ & \min_{\Upsilon_{k|\mathscr{A}_k}^*} \left\{ \mathbb{E}\left[\mathbf{E}_k^{\alpha\mathsf{T}} \Omega \mathbf{E}_k^\alpha \middle| \mathscr{A}_k \right] \right\} \\ & + \Upsilon_{k|\mathscr{A}_k}^{*\mathsf{T}} \bar{\nu}^\alpha \left(2\mathbf{F} X_k + \left(\Delta^\Gamma \bar{\nu}^\alpha + \Psi\right) \Upsilon_{k|\mathscr{A}_k}^* \right) \Big\}, \end{aligned} \tag{11}$$

where $\Delta^\Phi = \Phi^\mathsf{T} \Omega \Phi$, $\Delta^\Gamma = \Gamma^\mathsf{T} \Omega \Gamma$, and $\mathbf{F} = \Gamma^\mathsf{T} \Omega \Phi$.

2.2 Monitoring of Packet Losses and Attack Detection

The optimal control law for both protocols is determined by the mean of packet losses as shown in (5). In view of this, the operator monitors the average number of losses on the actuation channel to check that it agrees with the postulated statistic used to construct the control law. Given that the packet losses form a Bernoulli IID sequence, the distribution is fully characterised by the mean of packet losses $\mathbf{M}$. To that end, the system computes the average number of packet losses over each dimension on the channel up to time step $k \in \mathbb{N}$ thus producing the estimate for dimension i given by

$$\hat{\mu}_{i,k} = \frac{1}{k} \sum_{j=1}^{k} (\mathbf{V}_j)_{i,i}, \tag{12}$$

where $(\mathbf{V}_j)_{i,i}$ describes the i,i-th element of $\mathbf{V}_j$. The resulting estimate of the mean probability of packet loss at time step k is given by

$$\widehat{\mathbf{M}}_k = \operatorname{diag}(\hat{\mu}_{1,k}, \hat{\mu}_{2,k}, \dots, \hat{\mu}_{m,k}). \tag{13}$$

The system uses the estimate to check whether the actuation channel is nominal. In this setting, nominal operation entails that the estimated mean does not deviate significantly from the postulated mean used by the controller to implement the control law. Specifically, the operator determines a safe operation region shaped as a hypercube centered around $\mathbf{M}$ and with edge lengths determined by $\mathbf{L} \in S_+^m$, where the structure of the lengths is such that $\mathbf{L} = \operatorname{diag}(\varepsilon_1, \varepsilon_2, \dots, \varepsilon_m)$ to account for the different detection thresholds $\{\varepsilon_i\}_{i=1}^m$ for each dimension of the actuation channel. The resulting safe operation region is given by

$$\mathscr{C}(\mathbf{M}, \mathbf{L}) = \left\{ \widehat{\mathbf{M}}_k \in [0,1]^{m \times m} : -\mathbf{L} \preceq \widehat{\mathbf{M}}_k - \mathbf{M} \preceq \mathbf{L} \right\}. \tag{14}$$

In this setting, an attack is declared at time step $k \in \mathbb{N}$ if $\widehat{\mathbf{M}}_k \notin \mathscr{C}(\mathbf{M}, \mathbf{L})$. Otherwise, normal operation of the system continues and the operator monitors the packet losses by updating its estimate $\widehat{\mathbf{M}}_k$ at time step $k \in \mathbb{N}$. Note that the operator does not incorporate a monitoring performance metric in the cost function; instead, the packet loss monitoring procedure operates concurrent to the system operation but independently of the controller. The detection criteria chosen is optimal, the false alarm rate is determined by $\mathbf{L}$, if the operator decides upon a time varying $\mathbf{L}_k$ instead of a fixed $\mathbf{L}$ then instead define the $\mathscr{C}(\mathbf{M}, \mathbf{L})$ above as the intersection of all the time varying regions. In view of this, the attack construction is concerned with two performance metrics: the cost increase induced by the attack on the performance of the controller and satisfying that the calculated average of the packet losses induced by the attack conforms to the safe region defined by (14). In the following, we discuss optimal attack construction strategies. It should be noted that the attacker has access to the variables $\mathbf{L}$ and $\mathbf{M}_k$.

3. IID ATTACK CONSTRUCTION

As a result of having access to different information sets, the UDP-like protocol error trajectory given in (10b) depends on the mean of the control variable for the attacker, while the TCP-like protocol does not depend on the mean of the control variable, as shown in (10a). For that reason, the derivation is presented separately for the each protocol.

3.1 UDP-like Protocol

The optimal attack strategy for the UDP-like protocol is the solution to the optimisation problem

$$\max_{\bar{\nu}^\alpha} \; J_A(\mathscr{A}_k) \tag{15a}$$

$$\text{s.t.} \quad \mathbf{M}^\alpha \in \mathscr{C}(\mathbf{M}, \mathbf{L}). \tag{15b}$$

Note that the maximisation aims to increase the cost incurred by the controller as a result of the packet losses induced by

the attack while the constraint aims to keep the attack within the safe operation region. Additionally, the maximisation in (15) and the minimisation of the control law in Casbolt et al. (2019) differ in that $\mathbf{M}^{\alpha} \neq \mathbf{M}$, and therefore, the terms within $\Upsilon^{*}_{k|\mathscr{G}_k}$ described in (5) do not cancel. In the UDP-like setting the information set, $\mathscr{A}_k$, does not have access to the previous realisations of packet losses for estimation, that is, $\mathscr{I}_k = \mathscr{G}_k$ in (7). Using Lemma 1 from Casbolt et al. (2019) and (11) yields the equivalent cost function given by

$$J^{*}_{A}(\mathscr{A}_k) = X_k^{\mathsf{T}}\left(\mathbf{Q}+\Delta^{\Phi}\right)X_k + \operatorname{tr}\left(\Delta^{\Lambda}\Sigma_{\Xi}\right) + \max_{\bar{\nu}^{\alpha}}\Big\{\Upsilon^{*\mathsf{T}}_{k|\mathscr{G}_k}\bar{\nu}^{\alpha}(2\mathbf{F}X_k + \left(\Delta^{\Gamma}\bar{\nu}^{\alpha}+\Psi+\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\left(\mathbf{I}-\bar{\nu}^{\alpha}\right)\right)\Upsilon^{*}_{k|\mathscr{G}_k}\Big)\Big\}, \quad (16)$$

where $\Delta^{\Lambda} = \Lambda^{\mathsf{T}}\Omega\Lambda$ and the maximisation is subject to $\mathbf{M}^{\alpha} \in \mathscr{C}(\mathbf{M},\mathbf{L})$. In the following, without loss of generality and for the sake of presentation clarity, it is assumed that all actuators for the system share a single communication channel as in Schenato et al. (2007). This simplifies the attack construction while displaying the same properties of the general attack construction. That being the case, $\bar{\nu}^{\alpha}$ is a diagonal matrix with equal entries, and therefore, $\bar{\nu}^{\alpha} = \alpha\mathbf{I}$ where $\alpha \in \mathbb{R}$ is the control variable of the attacker and $\alpha\mathbf{I} \in \mathscr{C}(\mathbf{M},\mathbf{L})$. Similarly, the detection region $\mathscr{C}(\mathbf{M},\mathbf{L})$ is described, in this case, by the interval $\mathscr{C}(\mu,\varepsilon)$ where $\mu \in [0,1)$ is the mean of the Bernoulli random variable describing the packet losses in the scalar case and $\varepsilon \in [0,1]$ denotes the detection threshold set by the operator. Within this setting, the attack strategy is characterised by the attack design parameter α. In view of this, substituting α as the control variable in (16) reduces the optimisation problem to

$$J^{*}_{A}(\mathscr{A}_k) = X_k^{\mathsf{T}}\left(\mathbf{Q}+\Delta^{\Phi}\right)X_k + \operatorname{tr}\left(\Sigma_{\Xi}\Delta^{\Lambda}\right) + \max_{\alpha\in\mathscr{C}(\mu,\varepsilon)}\left\{\Upsilon^{\mathsf{T}}_{k|\mathscr{G}_k}\alpha\left(\alpha\Delta^{\Gamma}+(1-\alpha)\left(\mathbf{I}\odot\Delta^{\Gamma}\right)+\Psi-2\mathbf{G}_{\mathscr{G}_k}\right)\Upsilon_{k|\mathscr{G}_k}\right\}.$$

Note that the first two terms on the right hand side of the equation above are constants that do not depend on α. Therefore, it is sufficient to maximise the last term. Substituting $-\mathbf{G}^{-1}_{\mathscr{G}_k}\mathbf{F}X_k$ for $\Upsilon_{k|\mathscr{G}_k}$ we write the term inside the maximisation as

$$f(\alpha) \triangleq X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{G}_k}\alpha\left(\alpha\Delta^{\Gamma}+(1-\alpha)\left(\mathbf{I}\odot\Delta^{\Gamma}\right)+\Psi - 2\mathbf{G}_{\mathscr{G}_k}\right)\mathbf{G}^{-1}_{\mathscr{G}_k}\mathbf{F}X_k. \quad (17)$$

The function (17) is concave, convex, or linear in α depending on the system parameters, and therefore, the attacker has no control over the convexity of the cost function used for the attack construction. However, the information set available to the attacker determines the convexity of the cost function, and therefore, the attacker is able to construct the optimal attack by solving (17) for any system parameters. In the following lemma we show that for the convex and linear systems the optimal attack construction is equivalent.

Lemma 1. Let (17) be convex or linear in α over $\mathscr{C}(\mu,\varepsilon)$. Then it's maximum is given by

$$\max\{f(\alpha)\} = \max\left\{f(\min\{\mathscr{C}(\mu,\varepsilon)\}), f(\max\{\mathscr{C}(\mu,\varepsilon)\})\right\}.$$

Proof. Assume there is a maximum of (17), $f(a)$, such that $a \in \operatorname{Int}\{\mathscr{C}(\mu,\varepsilon)\}$, we prove by contradiction that this is false, and therefore, the maximum is on the boundary of $\mathscr{C}(\mu,\varepsilon)$. By the definition of convexity, for $\delta > 0$ it holds that

$$f(a) \geq \max\{f(a+\delta), f(a-\delta)\}, \quad (18)$$

$$f(a) \geq tf(a+\delta)+(1-t)f(a-\delta). \quad (19)$$

It follows that $f(a)$ is greater than any point of the line connecting $f(a+\delta)$ and $f(a-\delta)$) however, this breaks the convexity assumption of (17), and therefore, the maximum is on the boundary. This concludes the proof. ■

When the function is concave there is a third maximising possibility, the case for which the global maximum of the function exists within the interval $\mathscr{C}(\mu,\varepsilon)$. The following lemma describes this case.

Lemma 2. Let (17) be concave in α over $\mathscr{C}(\mu,\varepsilon)$. Then the maximum of the function is given by

$$\max\{f(\alpha)\} = \max\Big\{f(\min\{\mathscr{C}(\mu,\varepsilon)\}), f(\max\{\mathscr{C}(\mu,\varepsilon)\}), f\left(\mathbb{1}_{\mathscr{C}(\mu,\varepsilon)}(\alpha_{\max})\,\alpha_{\max}\right)\Big\}, \quad (20)$$

where $\mathbb{1}_{\mathscr{B}}(\alpha_{\max})$ is the indicator function as a function of $\alpha_{\max}$ over the set $\mathscr{B}$ and $\alpha_{\max} \in \mathbb{R}$ is the global maximum of $f(\alpha)$.

Proof. In the concave case a global maximum exists, but is not necessarily within the interval $\mathscr{C}(\mu,\varepsilon)$, and therefore, we restrict the domain to the safe operation region with the indicator function $\mathbb{1}_{\mathscr{C}(\mu,\varepsilon)}(\alpha)$. The concavity of the function implies

$$f'(\alpha) = \Upsilon^{*\mathsf{T}}_{k|\mathscr{G}_k}\left(2\alpha\Delta^{\Gamma}+(1-2\alpha)\left(\mathbf{I}\odot\Delta^{\Gamma}\right)+\Psi-2\mathbf{G}_{\mathscr{G}_k}\right)\Upsilon^{*}_{k|\mathscr{G}_k}, \quad (21)$$

$$f''(\alpha) = 2\Upsilon^{*\mathsf{T}}_{k|\mathscr{G}_k}\left(\Delta^{\Gamma}-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\Upsilon^{*}_{k|\mathscr{G}_k} < 0, \quad (22)$$

where (22) follows from the strict concavity of (17). Setting (21) equal to zero gives

$$\Upsilon^{*\mathsf{T}}_{k|\mathscr{G}_k}\left(2\alpha\Delta^{\Gamma}+(1-2\alpha)\left(\mathbf{I}\odot\Delta^{\Gamma}\right)+\Psi-2\mathbf{G}_{\mathscr{G}_k}\right)\Upsilon^{*}_{k|\mathscr{G}_k} = 0,$$

which results in

$$2\alpha X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{G}_k}\left(\Delta^{\Gamma}-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}^{-1}_{\mathscr{G}_k}\mathbf{F}X_k = X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{G}_k}\left(2\mathbf{G}_{\mathscr{G}_k}-\Psi-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}^{-1}_{\mathscr{G}_k}\mathbf{F}X_k. \quad (23)$$

It follows from the strict concavity of (17), as in (22), that

$$X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{G}_k}\left(\Delta^{\Gamma}-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}^{-1}_{\mathscr{G}_k}\mathbf{F}X_k \neq 0. \quad (24)$$

In view of this (23) can be solved for α yielding

$$\alpha_{\max} = \frac{1}{2}h^{-1}_{\mathrm{UDP}}\Big(X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{G}_k}\left(2\mathbf{G}_{\mathscr{G}_k}-\Psi-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}^{-1}_{\mathscr{G}_k}\mathbf{F}X_k\Big), \quad (25)$$

where $h_{\mathrm{UDP}} \triangleq \Upsilon^{*\mathsf{T}}_{k|\mathscr{G}_k}\left(\Delta^{\Gamma}-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\Upsilon^{*}_{k|\mathscr{G}_k}$. The global maximum is the solution when $\alpha_{\max} \in \mathscr{C}(\mu,\varepsilon)$, i.e. the term $\alpha_{\max}\mathbb{1}_{\mathscr{C}(\mu,\varepsilon)}(\alpha_{\max})$ in (20). When $\alpha_{\max} \notin \mathscr{C}(\mu,\varepsilon)$ the solution follows as in the convex scenario by noticing that the inequality is strict and in the opposite direction. Therefore, if $\alpha_{\max} \notin \mathscr{C}(\mu,\varepsilon)$ the attack construction reverts to selecting the value of α on the maximising boundary. For a concave function this is equivalent to finding the boundary that is closest to $\alpha_{\max}$. Let $a,b \in \mathscr{C}(\mu,\varepsilon)$ and assume $f(a) > f(b)$ and that $|a-\alpha_{\max}| < |b-\alpha_{\max}|$, then

$$f(b) < tf(a)+(1-t)f(\alpha_{\max}). \quad (26)$$

However, this line segment lies above the function which contradicts the fact that this function is concave, and therefore, the maximising α is on the boundary that is closest to $\alpha_{\max}$. This concludes the proof. ■

Note that the α_{max} attack construction provides a globally optimal performance for the attacker from within the safe

operation region. In fact, it also provides a lower probability of attack detection as it allows the attacker to operate away from the boundary.

The following lemma highlights that an attack that minimises the cost of the operator is not achieved by setting $\alpha = 1$. In the following we show that the optimal attack construction does not necessarily imply increasing the number of packet losses incurred by the operator. Indeed, there exist system parameters for which the optimal attack entails increasing the number of actuations. Whilst this might not be implementable in all attack scenarios, it is feasible to envision settings in which the attacker has full control of the actuation channel and can set the packet loss statistics at will. The following lemma captures this notion, namely, that the performance of the operator does not necessarily improve with the average number of received packets. Reiterating that the operator assumes a mean packet loss, and in doing so, creates an opportunity for the attacker to exploit the channel.

Lemma 3. For any choice of system parameters it holds that

$$\min_{a\in[0,1)} f(a) \leq \min\{f(1), f(\mu)\}, \tag{27}$$

where f is defined in (17).

Proof. Setting (21) equal to zero, and substituting in $\alpha = 1$ yields

$$f'(1) = \Upsilon^{*\mathsf{T}}_{k|\mathscr{G}_k}\left(2\left(\mathbf{I}-\bar{\nu}\right)\Delta^{\Gamma} - \Psi - \left(3\mathbf{I}-2\bar{\nu}\right)\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\Upsilon^{*}_{k|\mathscr{G}_k}.$$

For this to be a minimising solution it needs to hold that, Ψ is equal to $2\left(\mathbf{I}-\bar{\nu}\right)\Delta^{\Gamma} - \left(3\mathbf{I}-2\bar{\nu}\right)\left(\mathbf{I}\odot\Delta^{\Gamma}\right)$. Due to Ψ being a diagonal matrix and the structure of Δ^{Γ} it is only possible for this equality to hold in a system with $\mathbf{A} = \mathbf{0}$ and a diagonal $\mathbf{B}$. In this scenario, $\Delta^{\Gamma} = \left(\mathbf{I}\odot\Delta^{\Gamma}\right)$, this results in, $\Psi = -\left(\mathbf{I}\odot\Delta^{\Gamma}\right)$. By assumption $\Psi \succ 0$, however, it is shown in Lemma 5 that $\left(\mathbf{I}\odot\Delta^{\Gamma}\right) \succ 0$ which is a contradiction, and therefore, $f'(\alpha = 1) \neq 0$ and thus $\alpha = 1$ is not a minimising solution. Substituting $\alpha\mathbf{I} = \mu\mathbf{I} = \bar{\nu}$ in (21) results in

$$\begin{aligned} f'(\alpha) &= \Upsilon^{*\mathsf{T}}_{k|\mathscr{G}_k}\left(2\bar{\nu}\Delta^{\Gamma} + \left(\mathbf{I}-2\bar{\nu}\right)\left(\mathbf{I}\odot\Delta^{\Gamma}\right) + \Psi - 2\mathbf{G}_{\mathscr{G}_k}\right)\Upsilon^{*}_{k|\mathscr{G}_k}, \\ &= -X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{G}_k}\left(\Psi + \left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}^{-1}_{\mathscr{G}_k}\mathbf{F}X_k. \end{aligned} \tag{28}$$

This is not equal to 0 due to $\Psi + \left(\mathbf{I}\odot\Delta^{\Gamma}\right) \succ 0$. Therefore, (28) is strictly negative and not a minimising solution. This concludes the proof. ■

Theorem 4. Let $\mathscr{A}_k = \{\mathbf{A}, \mathbf{B}, \Sigma_W, \bar{\nu}, \Omega, \Psi, \mathscr{G}_k\}$ be the information set available to construct the attack, then the optimal mean packet loss probability for an IID attack is given by

$$\alpha^*_{\text{UDP}} = \max\Big\{f\left(\min\{\mathscr{C}(\mu,\varepsilon)\}\right), f\left(\max\{\mathscr{C}(\mu,\varepsilon)\}\right), f\left(\mathbb{1}_{\mathscr{C}(\mu,\varepsilon)}(\alpha_{\max})\,\alpha_{\max}\right)\Big\},$$

where

$$f(a) \triangleq X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{G}_k} a\left(a\Delta^{\Gamma} + (1-a)\left(\mathbf{I}\odot\Delta^{\Gamma}\right) + \Psi - 2\mathbf{G}_{\mathscr{G}_k}\right)\mathbf{G}^{-1}_{\mathscr{G}_k}\mathbf{F}X_k.$$

Proof. The result follows from the application of Lemma 1 for the convex and linear cases, Lemma 2 for the concave case, and by noticing that the set of solutions for the convex and linear cases is a subset of of the set of solutions of the concave case. This concludes the proof. ■

3.2 TCP-like Protocol

The optimal attack strategy for the TCP-like protocol is the solution to the optimisation problem

$$\max_{\bar{\nu}^{\alpha}} \; J_A(\mathscr{A}_k), \tag{29a}$$

$$\text{s.t.} \quad \mathbf{M}^{\alpha} \in \mathscr{C}(\mathbf{M}, \mathbf{L}). \tag{29b}$$

Note that in this case the information set $\mathscr{A}_k$ contains the realisations of the packet losses as given in $\mathscr{F}_k$. For that reason, the optimisation problem differs from that in (15) in that the cost function exhibits a different structure induced by the conditioning of the previous packet loss realisations. From (11) and Lemma 1 in Casbolt et al. (2019) with substitution of the optimal control law under the TCP-like protocol in (5) yields

$$\begin{aligned} J^*(\mathscr{A}_k) &= X_k^{\mathsf{T}}\left(\mathbf{Q} + \Delta^{\Phi}\right)X_k + \operatorname{tr}\left(\Delta^{\Lambda}\Sigma_{\Xi}\right) \\ &+ \max_{\bar{\nu}^{\alpha}}\left\{X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{F}_k}\bar{\nu}^{\alpha}\left(\Delta^{\Gamma}\bar{\nu}^{\alpha} + \Psi - 2\mathbf{G}_{\mathscr{F}_k}\right)\mathbf{G}^{-1}_{\mathscr{F}_k}\mathbf{F}X_k\right\}, \end{aligned} \tag{30}$$

where the maximisation is subject to $\mathbf{M}^{\alpha} \in \mathscr{C}(\mathbf{M}, \mathbf{L})$. As with the UDP-like protocol attack construction, it is assumed without loss of generality, that all actuators share a single communication channel (Schenato et al. (2007)). Therefore, $\bar{\nu}^{\alpha} = \alpha\mathbf{I}$. Noting that the first two terms in (30) do not depend on $\bar{\nu}^{\alpha}$ and that $\mathbf{G}_{\mathscr{F}_k} = \left(\Delta^{\Gamma}\bar{\nu} + \Psi\right)$ as shown in (Casbolt et al. (2019)), then the term inside the maximisation can be rewritten as

$$g(\alpha) \triangleq -X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{F}_k}\alpha\left(\Delta^{\Gamma}\left(2\bar{\nu} - \alpha\mathbf{I}\right) + \Psi\right)\mathbf{G}^{-1}_{\mathscr{F}_k}\mathbf{F}X_k. \tag{31}$$

Differentiating (31) results in

$$g'(\alpha) = -X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{F}_k}\left(\Delta^{\Gamma}\left(2\bar{\nu} - 2\alpha\mathbf{I}\right) + \Psi\right)\mathbf{G}^{-1}_{\mathscr{F}_k}\mathbf{F}X_k, \tag{32}$$

$$g''(\alpha) = 2X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{F}_k}\Delta^{\Gamma}\mathbf{G}^{-1}_{\mathscr{F}_k}\mathbf{F}X_k. \tag{33}$$

Lemma 5. Let the pair $(\mathbf{A}, \mathbf{B})$ be reachable and the state penalty matrix Ω be positive definite. Then the function defined by (31) is convex in α over $\mathscr{C}(\mu,\varepsilon)$ almost surely.

Proof. It follows from (33) that if $\Delta^{\Gamma} \succ 0$ and $X_k \neq \mathbf{0}$ then (33) is strictly greater than zero. Therefore, (31) is convex in α over $\mathscr{C}(\mu,\varepsilon)$. It is shown in (Seber, 2007, p.225, 10.31(c)) that when $\operatorname{rank}(\Gamma) = \max\{Nn, Nm\}$ and Ω is positive definite then $\Delta^{\Gamma} \succ 0$. Since $(\mathbf{A}, \mathbf{B})$ is a reachable pair then $\operatorname{rank}\left[\mathbf{B}, \mathbf{AB}, \ldots, \mathbf{A}^{N-1}\mathbf{B}\right] = \max\{n, m\}$. Therefore, due to the triangular structure of Γ we have that $\operatorname{rank}(\Gamma) = \max\{Nn, Nm\}$. Under these assumptions (33) is convex in α over $\mathscr{C}(\mu,\varepsilon)$ when $X_k \neq 0$, which holds with probability 1. This concludes the proof. ■

Theorem 6. Consider a system operating with a TCP-like protocol.

$$g(\alpha) \triangleq -X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}^{-1}_{\mathscr{F}_k}\alpha\left(\Delta^{\Gamma}\left(2\bar{\nu} - \alpha\mathbf{I}\right) + \Psi\right)\mathbf{G}^{-1}_{\mathscr{F}_k}\mathbf{F}X_k,$$

Therefore, the optimal choice of α is

$$\alpha^*_{\text{TCP}} = \max\Big\{g\left(\min\{\mathscr{C}(\mu,\varepsilon)\}\right), g\left(\max\{\mathscr{C}(\mu,\varepsilon)\}\right)\Big\}.$$

Proof. Note from Lemma 5 that $g(\alpha)$ is convex therefore, as shown in Lemma 3, α^*_{TCP} is known to be on the boundary. This concludes the proof.

Note that due to the convexity of (31) the solution of (32) results in the minimising value of α, which interestingly is not $\alpha\mathbf{I}=\bar{\nu}$, or $\alpha=1$ but instead is given by

$$\alpha_{\min}=\frac{1}{2}h_{\mathrm{TCP}}^{-1}\left(X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\left(2\Delta^{\Gamma}\bar{\nu}+\Psi\right)\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k\right), \quad (34)$$

where $h_{TCP}=X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{F}_k}^{-1}\Delta^{\Gamma}\mathbf{G}_{\mathscr{F}_k}^{-1}\mathbf{F}X_k>0$. When considering the TCP-like protocol without detection constraints, additional insight can be obtained by analysing the attack construction

$$g(1)=X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{F}_k}^{-1}\left(\Delta^{\Gamma}\left(\mathbf{I}-2\bar{\nu}\right)-\Psi\right)\mathbf{G}_{\mathscr{F}_k}^{-1}\mathbf{F}X_k,\ g(0)=0. \quad (35)$$

From (35), if $\Delta^{\Gamma}(\mathbf{I}-2\bar{\nu})\succ\Psi$ the maximising value of α is 1 or 0. The $\Delta^{\Gamma}(\mathbf{I}-2\bar{\nu})$ term is the state penalty matrix Ω weighted by the reachability of the system and the packet loss probability. The terms in (35) capture the average impact of actuation in the cost reduction with respect to the input penalty matrix Ψ. Therefore, the optimal attack is 1 when the average cost increase per actuation is greater than the average penalty induced by the actuation. As a result, for a system with a high probability of packet loss that penalises state error more than actuation, the optimal attack strategy is to allow perfect communication, i.e. all packets are received by the plant. Additionally, for $\bar{\nu}\succ\frac{1}{2}\mathbf{I}$ the optimal attack strategy is $\bar{\nu}=\mathbf{0}$. That being the case, for a system with a low probability of packet losses the operator could simplify their detection criteria to a one-sided test.

4. COST INCREASE ANALYSIS

In this section we evaluate the cost increase induced by the optimal IID attack by comparing the expected cost when an attack is present to the expected cost when no attack is present, i.e. $\mathbb{E}[J_A^*(\mathscr{A}_k)]-\mathbb{E}[J^*(\mathscr{I}_k)]$. The expected cost increase of the three attack strategies are studied separately. The analysis is carried out for the case $\mathscr{C}(\mu,\varepsilon)=[0,1)$, i.e. the extreme cases of the average attack packet drop. Note that there is no loss of generality as the case with detection constraints can be analysed following the same approach with the appropriate scaling. Additionally, when considering the different cases of α^* it should be noted that for a given set of detection parameters α^* is unique. Since, the detection parameters, μ,ε, are not fixed. The region $[0,1)$ is continuous with respect to α^*.

4.1 UDP-like Cost Analysis

Attack performance when $\alpha^*\to 0$. We first analyse the case when the attacker losses all the packets and induces the cost

$$\mathbb{E}\left[J_A^0(\mathscr{A}_k)\right]\triangleq\lim_{\alpha^*\to 0}\mathbb{E}\left[J_A(\mathscr{A}_k)\right]. \quad (36)$$

The expected cost when there is an attack is given by

$$\mathbb{E}\left[J_A^0(\mathscr{A}_k)\right]=\mathrm{tr}\left(\Sigma_X\left(\mathbf{Q}+\Delta^{\Phi}\right)+\Sigma_{\Xi}\Delta^{\Lambda}\right)+\max\{f(\alpha)\}. \quad (37)$$

Since (17) is continuous in α we have that $\alpha^*\to 0$ implies $f(\alpha^*)\to 0$, and therefore, the cost increase is

$$\mathbb{E}\left[J_A^0(\mathscr{A}_k)\right]-\mathbb{E}\left[J^*(\mathscr{G}_k)\right]=X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\bar{\nu}\mathbf{F}X_k>0. \quad (38)$$

Note that the $\alpha^*\to 0$ attack strategy forces the system into open loop, and therefore, the expected cost increase coincides with the expected cost reduction introduced by the controller when there is no attack present in the communication channel.

Attack performance when $\alpha^*=1$. In this case, the attacker allows successful reception of all packets, i.e. the actuation communication channel is perfect. Surprisingly, there exist systems for which the cost increase, given by

$$\mathbb{E}\left[J_A^1(\mathscr{A}_k)\right]\triangleq\mathbb{E}\left[J_A(\mathscr{A}_k)\right]\Big|_{\alpha^*=1}, \quad (39)$$

is positive despite the fact that the communication channel of the operator improves. Evaluation of (17) with perfect communication results in

$$\begin{aligned}\mathbb{E}\left[J^1(\mathscr{A}_k)\right]=&\,\mathrm{tr}\left(\Sigma_X\left(\mathbf{Q}+\Delta^{\Phi}\right)+\Sigma_{\Xi}\Delta^{\Lambda}\right)\\&+X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\left(\mathbf{I}-2\bar{\nu}\right)\left(\Delta^{\Gamma}-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k\\&-X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\left(\left(\mathbf{I}\odot\Delta^{\Gamma}\right)+\Psi\right)\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k. \quad (40)\end{aligned}$$

Unlike the $\alpha^*\to 0$ case, the $\alpha^*=1$ construction does not guarantee an increase in cost for every system. In fact, the cost only increases when

$$\begin{aligned}X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\left(\mathbf{I}-2\bar{\nu}\right)\left(\Delta^{\Gamma}-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k\ \geq\\ X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\left(\left(\mathbf{I}\odot\Delta^{\Gamma}\right)+\Psi\right)\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k>0. \quad (41)\end{aligned}$$

However, all variables that determine (41) are system parameters known by the attacker, and therefore, the attacker decides the optimal attack strategy accordingly. The expected cost increase is

$$\begin{aligned}&\mathbb{E}\left[J_A^1(\mathscr{A}_k)\right]-\mathbb{E}\left[J^*(\mathscr{G}_k)\right]\\&=X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\left(\Delta^{\Gamma}+\Psi\right)\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k+\left(\bar{\nu}-2\mathbf{I}\right)X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k.\end{aligned}$$

Attacker performance when $\alpha^*=\mathbb{1}_{\mathscr{C}(\mu,\varepsilon)}(\alpha_{\max})\alpha_{\max}$. We tackle next the introduction of a general detection constraint. In this case, the expected cost for the attacker is

$$\begin{aligned}\mathbb{E}\left[J^{\alpha_{\max}}(\mathscr{A}_k)\right]&\triangleq\mathbb{E}\left[J_A(\mathscr{A}_k)\right]\Big|_{\alpha^*=\alpha_{\max}}\\&=\mathrm{tr}\left(\Sigma_X\left(\mathbf{Q}+\Delta^{\Phi}\right)+\Sigma_{\Xi}\Delta^{\Lambda}\right)+f(\alpha_{\max}).\end{aligned}$$

Algebraic maniupulation of $f(\alpha_{\max})$ and substituting (25) yields:

$$f(\alpha_{\max})=\frac{h_{\mathrm{UDP}}^{-1}}{4}\left(X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\left(2\mathbf{G}_{\mathscr{G}_k}-\Psi-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k\right)^2,$$

where the inequality comes from the fact that f is concave when $\alpha_{\max}$ is a feasible optimal attack strategy. The resulting cost increase is

$$\begin{aligned}&\mathbb{E}\left[J^{\alpha_{\max}}(\mathscr{A}_k)\right]-J^*(\mathscr{G}_k)=X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\bar{\nu}\mathbf{F}X_k\\&+\frac{1}{4}h_{\mathrm{UDP}}^{-1}\left(X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{G}_k}^{-1}\left(2\mathbf{G}_{\mathscr{G}_k}-\Psi-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\mathbf{G}_{\mathscr{G}_k}^{-1}\mathbf{F}X_k\right)^2>0.\end{aligned}$$

Note that the inequality is strict, i.e. the attack guarantees a performance loss of the operator. As mentioned previously this attack strategy is only feasible when $\alpha_{\max}\in\mathscr{C}(\mu,\varepsilon)$ and (17) is concave. Additionally, $\mathbb{E}\left[J^0(\mathscr{A}_k)\right]-\mathbb{E}\left[J^*(\mathscr{G}_k)\right]$ is a upper bounded by the cost increase induced by the $\alpha^*=\mathbb{1}_{\mathscr{C}(\mu,\varepsilon)}(\alpha_{\max})\alpha_{\max}$ strategy.

4.2 TCP-like Cost Analysis

The cost increase analysis for TCP-like protocols contains only two attack strategies. The analysis is again performed on the $\mathscr{C}(\mu,\varepsilon)=[0,1)$ interval.

Attacker performance when $\alpha^* \to 0$. For the attack construction that forces the system into open loop the expected cost is given by

$$\mathbb{E}\left[J^0\left(\mathscr{A}_k\right)\right] = \operatorname{tr}\left(\Sigma_X\left(\mathbf{Q}+\Delta^{\Phi}\right)+\Sigma_{\Xi}\Delta^{\Lambda}\right) + \lim_{\alpha^*\to 0} g(\alpha^*).$$

Since (35) is continuous in α we have that $\alpha^* \to 0$ implies $g(\alpha^*) \to 0$. Therefore, the expected cost increase is:

$$\mathbb{E}\left[J^0\left(\mathscr{A}_k\right)\right] - \mathbb{E}\left[J^*\left(\mathscr{F}_k\right)\right] = X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{F}_k}^{-1}\bar{\nu}\mathbf{F}X_k > 0.$$

As with the UDP-like protocol, by implementing the $\alpha^* \to 0$ attack strategy the attacker forces the system into open loop. Note that the cost increase for the TCP-like protocol and the UDP-like protocol under the $\alpha^* \to 0$ strategy differ only in the $\mathbf{G}_{\mathscr{I}_k}$ designed by the controller, i.e. on the available information.

Attacker performance when $\alpha^* = 1$. In the TCP case, the attack that provides a perfect communication channel induces an expected cost given by

$$\mathbb{E}\left[J^1\left(\mathscr{A}_k\right)\right] = \operatorname{tr}\left(\Sigma_X\left(\mathbf{Q}+\Delta^{\Phi}\right)+\Sigma_{\Xi}\Delta^{\Lambda}\right) + g(1).$$

Therefore, it follows form (35) that the expected cost increase induced by the $\alpha^* = 1$ strategy is

$$\begin{aligned}&\mathbb{E}\left[J^1\left(\mathscr{A}_k\right)\right] - \mathbb{E}\left[J^*\left(\mathscr{F}_k\right)\right]\\ &= \underbrace{X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{F}_k}^{-1}\left(\Delta^{\Gamma}\left(\mathbf{I}-2\bar{\nu}\right)-\Psi\right)\mathbf{G}_{\mathscr{F}_k}^{-1}\mathbf{F}X_k}_{>0} + \underbrace{X_k^{\mathsf{T}}\mathbf{F}^{\mathsf{T}}\mathbf{G}_{\mathscr{F}_k}^{-1}\bar{\nu}\mathbf{F}X_k}_{\mathbb{E}\left[J^0\left(\mathscr{A}_k\right)\right]-\mathbb{E}\left[J^*\left(\mathscr{F}_k\right)\right]},\end{aligned}$$

where the first term is strictly positive following the assumption that the $\alpha^* = 1$ attack construction is optimal. Therefore, $\alpha* = 1$ is strictly greater than the $\alpha^* \to 0$ attack strategy only when $\Delta^{\Gamma}\left(\mathbf{I}-2\bar{\nu}\right) \succ \Psi$, which is the condition needed for the $\alpha^* = 1$ construction to be optimal.

5. NON-STATIONARY RANDOM ATTACKS

Since the plant given in (1) is Markovian, it seems reasonable to assume that the attacker should be able to exploit the memory of the system in the construction of the attack. In that sense, the IID attack construction does not provide sufficient flexibility to incorporate the time dependency between consecutive packet losses. Motivated by this insight, in the following we investigate the extension of random attacks to non-IID settings. Specifically, we consider the case in which the statistics of the attack are non-stationary. The resulting non-stationary attack construction extends the IID attack construction to an attack that corrupts a system with independent actuator channels. As in the IID case, the aim of the attacker is to increase the cost function while remaining in the safe operation region by adjusting the value of $\mathbf{M}^{\alpha}$. Additionally, the attack construction is no longer restricted to a constant $\mathbf{M}^{\alpha}$, i.e. $\mathbf{M}_k^{\alpha} \triangleq \mathbb{E}[\mathbf{V}_k^{\alpha}]$ for $k \in \mathbb{N}$. The derivation of the non-stationary attack construction is equivalent to the IID attack construction up to (16). The reason for the necessity of a different derivation stems from the fact that $\bar{\nu}^{\alpha} \neq \alpha\mathbf{I}$ since $\mathbf{M}_k^{\alpha} \neq \mathbf{M}_{k+1}^{\alpha}$ for $k \in \mathbb{N}$.

We first consider the non-stationary attack construction for the UDP-like protocol. Notice that for the non-stationary construction maximising (16) is equivalent to is equivalent to maximising the function

$$\begin{aligned}f\left(\bar{\nu}^{\alpha}\right) &= \Upsilon_{\mathscr{G}_k}^{*\mathsf{T}}\bar{\nu}^{\alpha}\left(\Delta^{\Gamma}\bar{\nu}^{\alpha}+\Psi+\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\left(\mathbf{I}-\bar{\nu}^{\alpha}\right)-2\mathbf{G}_{\mathscr{G}_k}\right)\Upsilon_{\mathscr{G}_k}^{*},\\ &= \operatorname{tr}\left(\bar{\nu}^{\alpha}\left(\left(\Delta^{\Gamma}-\left(\mathbf{I}\odot\Delta^{\Gamma}\right)\right)\bar{\nu}^{\alpha}+\left(\mathbf{I}\odot\Delta^{\Gamma}\right)+\Psi-2\mathbf{G}_{\mathscr{G}_k}\right)\Upsilon_{\mathscr{G}_k}^{*}\Upsilon_{\mathscr{G}_k}^{*\mathsf{T}}\right),\end{aligned}$$

where the maximisation is subject to $\mathbf{M}_k^{\alpha} \in \mathscr{C}\left(\mathbf{M},\mathbf{L}\right)$ for $k \in \mathbb{N}$. Letting $\Delta^H = \Delta^{\Gamma} - \left(\mathbf{I}\odot\Delta^{\Gamma}\right)$, and substituting $\mathbf{G}_{\mathscr{G}_k}$ allows the optimal attack strategy for the UDP-like protocol to be posed as a quadratic optimisation problem (QP) given by

$$\begin{aligned}\max_{\bar{\nu}^{\alpha}} \quad & \operatorname{tr}\left(\left[\bar{\nu}^{\alpha}\Delta^{H}\bar{\nu}^{\alpha}-\bar{\nu}^{\alpha}\left(\left(\mathbf{I}\odot\Delta^{\Gamma}\right)+\Psi+2\bar{\nu}\Delta^{H}\right)\right]\Upsilon_{\mathscr{G}_k}^{*}\Upsilon_{\mathscr{G}_k}^{*\mathsf{T}}\right),\\ \text{s.t.} \quad & \mathbf{M}_k^{\alpha} \in \mathscr{C}\left(\mathbf{M},\mathbf{L}\right) \qquad \text{for } k \in \mathbb{N}. \end{aligned} \tag{42}$$

Note that the set of IID attack strategies is a subset of the strategies generated with this formulation. Therefore, if IID is indeed the optimal attack then the proposed non-stationary attack construction coincides with the strategy presented in the previous section. If however, the cost induced by the non-stationary attack is greater than that induced by $\alpha\mathbf{I}$, then it follows that memory in the attack yields larger cost increases while satisfying the same detection constraints. Performing the same analysis for the TCP-like system results in an analogous QP formulation given by

$$\begin{aligned}\max_{\bar{\nu}^{\alpha}} \quad & \operatorname{tr}\left(\left[\bar{\nu}^{\alpha}\Delta^{\Gamma}\bar{\nu}^{\alpha}-\bar{\nu}^{\alpha}\left(2\bar{\nu}\Delta^{\Gamma}+\Psi\right)\right]\Upsilon_{\mathscr{F}_k}^{*}\Upsilon_{\mathscr{F}_k}^{*\mathsf{T}}\right),\\ \text{s.t.} \quad & \mathbf{M}_k^{\alpha} \in \mathscr{C}\left(\mathbf{M},\mathbf{L}\right) \qquad \text{for } k \in \mathbb{N}. \end{aligned} \tag{43}$$

In the TCP-like scenario the terms Δ^{Γ}, Ψ, and $\Upsilon_{\mathscr{F}_k}^{*}\Upsilon_{\mathscr{F}_k}^{*\mathsf{T}}$ are positive semidefinite. Both UDP-like and TCP-like QP formulations can be modified to include IID attacks on non-scalar systems provided that the additional constraint $\mathbf{M}_k^{\alpha} = \mathbf{M}_{k+1}^{\alpha}$ for $k \in \mathbb{N}$ is included.

6. NUMERICAL RESULTS

The comparison between the IID attack and the non-stationary attack is conducted over two communication channels for the same system. To that end, we use the same test system as in Casbolt et al. (2019). The first simulation is performed over a scalar communication channel with $\mathbf{M} = 0.7$ by averaging 1000 realisations of the state trajectories and is shown in Fig. 3 and 4. Additionally, when under attack the UDP-like system trajectory depicted in Fig. 4 displays a larger change from the nominal state trajectory when compared with TCP-like trajectory depicted in Fig. 3. For the second channel model, i.e. $\mathbf{M} = \left(\begin{smallmatrix}0.7 & 0\\ 0 & 0.01\end{smallmatrix}\right)$, the non-stationary attack results in a larger increase from the nominal state trajectory when compared to the IID attack as shown in Fig. 5 and Fig. 6, which are obtained by averaging 1000 realisations of the state trajectories. As in the previous case, the results suggest that the system operating with a UDP-like protocol is more vulnerable to attacks than a system operating with a TCP-like protocol. The nominal terminal cost of the system with no attack present is 7.671 for the UDP-like protocol shown in Fig. 6. Interestingly, the terminal cost induced by the non-stationary attack is 18.217 while the terminal cost induced by the IID attack is 13.26. This suggests that UDP-like protocols are more vulnerable to non-stationary attacks than to IID attacks. However, the difference in the induced cost for the scalar channel shown in Fig. 4 is not as significant as that for the multiple input channel. Surprisingly, for the TCP-like case shown in Fig. 5 the performance of the IID attack outperforms the non-stationary attack. Specifically, the terminal cost induced by the

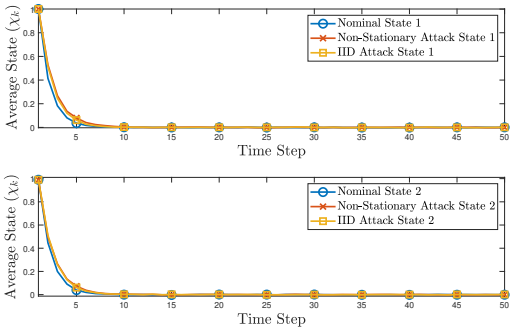

Fig. 3. System with TCP-like protocol with $\mathbf{A} = \left(\begin{smallmatrix} 1.03 & 0.005 \\ 0.35 & 0.5 \end{smallmatrix}\right)$, $\mathbf{M} = 0.7$, and $\varepsilon = 0.1$.

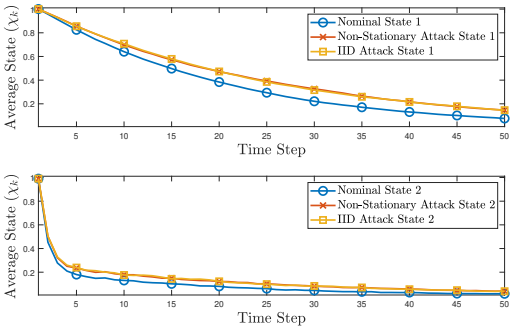

Fig. 4. System with UDP-like protocol with $\mathbf{A} = \left(\begin{smallmatrix} 1.03 & 0.005 \\ 0.35 & 0.5 \end{smallmatrix}\right)$, $\mathbf{M} = 0.7$, and $\varepsilon = 0.1$.

IID attack is 8.689 whereas the terminal cost induced by the non-stationary attack is 8.525. This small difference is due to numerical error in the simulations, indeed when the accuracy is increased the non-stationary cost converges to the IID attack cost. These results seem to suggest that there is no significant advantage in implementing non-stationary attacks in TCP-like systems.

7. CONCLUSION AND COMPARISON

We have characterised the optimal IID attack construction for UDP-like and TCP-like systems with lossy actuation channels. The attacks are envisioned as DoS attacks over the actuation communication channel which results in packet losses being induced by the attacker. Under the assumption that the operator monitors the state of the channel with the average packet loss as the decision statistic, we have shown that the optimal attack strategy does not always increase the number of packet losses. In fact, we have characterised the effect of the system parameters over the solution structure and shown that three different scenarios emerge for which the attack strategy is different. Interestingly, under both protocols the attacker only needs to know Δ^{Γ}, Ψ, and $\mathbf{M}$ to decide the optimal strategy, unless the system operates with a UDP-like protocol and the function is concave, in which case all system parameters must be known. For all cases, the cost increase of the optimal IID construction has been characterised and analysed. We have also shown that the IID attack construction is not optimal by proposing an achievability scheme that constructs attacks with non-stationary statistics. It is shown numerically that the proposed non-stationary attack outperforms the IID attack in most

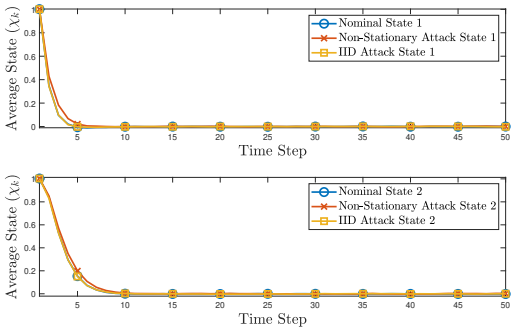

Fig. 5. System with TCP-like protocol with $\mathbf{A} = \left(\begin{smallmatrix} 1.03 & 0.005 \\ 0.35 & 0.5 \end{smallmatrix}\right)$, $\mathbf{M} = \left(\begin{smallmatrix} 0.7 & 0 \\ 0 & 0.01 \end{smallmatrix}\right)$, and $\mathbf{L} = 0.1\mathbf{I}$.

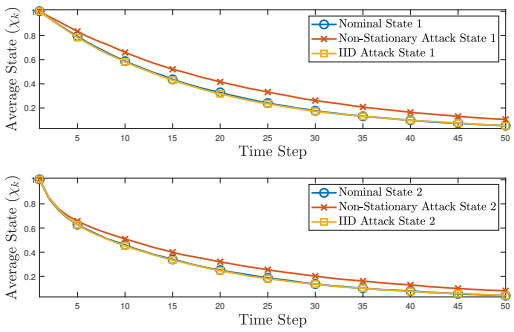

Fig. 6. System with UDP-like protocol with $\mathbf{A} = \left(\begin{smallmatrix} 1.03 & 0.005 \\ 0.35 & 0.5 \end{smallmatrix}\right)$, $\mathbf{M} = \left(\begin{smallmatrix} 0.7 & 0 \\ 0 & 0.01 \end{smallmatrix}\right)$, and $\mathbf{L} = 0.1\mathbf{I}$.

settings although at the expense of increased computational complexity.

REFERENCES

Casbolt, W., Jones, B., and Esnaola, I. (2019). Optimal control over multiple input lossy channels. *arXiv preprint arXiv:1911.07548*.

Colbert, E.J.M. and Kott, A. (2016). *Cyber-security of SCADA and Other Industrial Control Systems*. Springer International Publishing.

Mo, Y., Garone, E., and Sinopoli, B. (2013). LQG control with Markovian packet loss. *In Proc. European Control Conference*, 2380–2385.

Schenato, L., Sinopoli, B., Franceschetti, M., Poolla, K., and Sastry, S.S. (2007). Foundations of control and estimation over lossy networks. *In Proc. IEEE*, 95(1), 163–187.

Seber, G.A.F. (2007). *A Matrix Handbook for Statisticians*. Wiley-Interscience, New York, NY, USA.

Sinopoli, B., Schenato, L., Franceschetti, M., Poolla, K., and Sastry, S.S. (2005). Optimal control with unreliable communication: the tcp case. *In Proc. American Control Conference*, 5, 3354–3359.

Sinopoli, B., Schenato, L., Franceschetti, M., Poolla, K., and Sastry, S.S. (2008). Optimal linear LQG control over lossy networks without packet acknowledgment. *Asian Journal Control*, 10(1), 3–13.

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.