Preprints of the 20th World Congress
The International Federation of Automatic Control
Toulouse, France, July 9-14, 2017

State Estimation with Secrecy against Eavesdroppers

Anastasios Tsiamis, Konstantinos Gatsis, George J. Pappas

Department of Electrical and Systems Engineering, University of Pennsylvania, 200 South 33rd Street, Philadelphia, PA 19104, United States (e-mails: [atsiamis, kgatsis, pappasg]@seas.upenn.edu).

Abstract: We study the problem of remote state estimation, in the presence of an eavesdropper. An authorized user estimates the state of a linear plant, based on the data received from a sensor, while the data may also be intercepted by the eavesdropper. To maintain confidentiality with respect to state, we introduce a novel control-theoretic definition of perfect secrecy requiring that the user's expected error remains bounded while the eavesdropper's expected error grows unbounded. We propose a secrecy mechanism which guarantees perfect secrecy by randomly withholding sensor information, under the condition that the user's packet reception rate is larger than the eavesdropper's interception rate. Given this mechanism, we also explore the tradeoff between user's utility and confidentiality with respect to the eavesdropper, via an optimization problem. Finally, some examples are studied to provide insights about this tradeoff.

Keywords: Secrecy, Privacy, Security, Eavesdropping Attacks, Remote Estimation.

1. INTRODUCTION

The recent emergence of the Internet of Things as a collection of wirelessly connected sensor and actuator devices has given rise to significant cyber-security concerns (Cárdenas et al., 2008; Sandberg et al., 2015). Research efforts have targeted, for example, denial-of-service attacks (Amin et al., 2009; Gupta et al., 2010), privacy issues (Le Ny and Pappas, 2014), as well as data integrity of compromised sensors (Fawzi et al., 2014; Mo et al., 2014; Pajic et al., 2014). However, the broadcast nature of the wireless medium opens up further vulnerabilities in such connected systems (Zou et al., 2016). A fundamental vulnerability is confidentiality against *eavesdroppers* who may intercept the transmitted information. This becomes particularly crucial for sensor or actuator data, which convey critical information about the physical system state.

Encryption and cryptography-based tools are commonly employed for confidential communication (Katz and Lindell, 2014), dating back to the work of (Shannon, 1949). These approaches rely on encrypting communication messages with public or private keys, in order to achieve confidentiality against computationally limited eavesdroppers. These are generic tools, typically employed at intermediate layers of the communication protocol stack, and they do not take into account the characteristics of the application for which confidentiality is required, or the characteristics of the physical layer used for message communication, e.g., the wireless medium.

Additionally, it is possible to exploit the physical layer in order to achieve confidentiality, usually termed secrecy in this context (Wyner, 1975; Liang et al., 2008; Oggier and Hassibi, 2011; Regalia et al., 2015). This approach models the eavesdropper as overhearing the communication over a channel that is degraded compared to the legitimate channel. This channel disparity may be exploited using information-theoretic tools to achieve a positive communication rate with secrecy. This information-theoretic approach may be applied to problems in remote estimation and control, as discussed in preliminary works (Li et al., 2011; Wiese et al., 2016). However, the construction of practical secrecy codes is still an active area of investigation (Regalia et al., 2015).

In this paper, we take an alternative approach by introducing a novel control-theoretic definition of secrecy and by designing simple mechanisms that meet this definition. More specifically, we consider a sensor transmitting the outputs of an unstable linear dynamical system to a legitimate user, while an eavesdropper tries to intercept the sent messages compromising confidentiality. Communication follows the packet-based paradigm commonly used in networked control systems (Sinopoli et al., 2004; Hespanha et al., 2007), where the user and the eavesdropper respectively receive and intercept packets with different success rates. Our definition of perfect secrecy requires that any state estimate devised by the eavesdropper, suffers from an unbounded expected error in the limit, while the legitimate user still tracks the state with a bounded error (Section 2).

In Section 3, we show that perfect secrecy is possible as long as the packet success rate of the user is larger than the packet interception rate of the eavesdropper (Theorem 1). This is achieved by a simple mechanism, which randomly withholds information with the appropriate rate at the sensor. Then, estimation at both the user and the eavesdropper, follows two respective intermittent Kalman filter problems (Sinopoli et al., 2004). To achieve perfect secrecy, we exploit both the unstable dynamics and

⋆ This work was supported by NSF CNS-1505799 grant and the Intel-NSF Partnership for Cyber-Physical Systems Security and Privacy.

the inferiority of the eavesdropper's rate. The latter is similar in spirit to the degraded channel assumption in information-theoretic approaches (Regalia et al., 2015). Our approach differs by explicitly considering a control-theoretic definition of secrecy and by subsequently employing control-theoretic tools. Furthermore, our packet-based communication model yields a simple secrecy mechanism, in contrast to the problem of developing appropriate coding in, e.g., Li et al. (2011); Wiese et al. (2016).

In Section 4, we relax the perfect secrecy requirement, by seeking to achieve the minimum estimation error at the user, as long as the eavesdropper's error is larger than a desired lower bound. Due to the lack of analytical expressions for the intermittent Kalman filter error, we relax the problem by replacing the errors with known upper and lower bounds. The resulting optimization problem can be solved efficiently using bisection, and can be used as a tool for approximate, yet quantitative, analysis of the tradeoff between secrecy and utility. In Section 5, we illustrate this analysis in a scalar system example, which reveals how the parameters of the channel influence this tradeoff. We also present numerical results, which demonstrate the effect of employing the proposed secrecy mechanism to a second order system. We conclude with remarks in Section 6.

After this paper was submitted, a paper was written on a similar setting (Leong et al., 2017). The approach is different than ours since it utilizes acknowledgment signals from the user back to the sensor. In our paper, we assume that acknowledgments are not available.

2. PROBLEM FORMULATION

The considered remote estimation architecture is shown in Figure 1 and consists of a sensor observing a dynamical system, a legitimate user, and an eavesdropper. We consider the following linear dynamical system:

$$\begin{aligned} x(k+1) &= Ax(k) + w(k) \\ y(k) &= Cx(k) + v(k) \end{aligned} \tag{1}$$

where $x(k) \in \mathbb{R}^n$ is the state, $y(k) \in \mathbb{R}^m$ is the output and $k \in \mathbb{N}$ is the (discrete) time. To have a well posed estimation problem, we assume that (A, C) is detectable. Signals $w(k) \in \mathbb{R}^n$ and $v(k) \in \mathbb{R}^m$ are the process and measurement noise respectively. They are modeled as independent Gaussian random variables with zero mean and covariance matrices Q and R respectively. The initial state x_0 is also a Gaussian random variable with zero mean and covariance Σ_0. Covariance matrices R, Q, Σ_0 are assumed to be positive definite. In more compact notation $R, Q, \Sigma_0 \succ 0$, where $\succ$ ($\succeq$) denotes comparison in the positive definite (semidefinite) cone.

We consider an unstable system, i.e., its spectral radius is $\rho(A) = \max_i |\lambda_i(A)| > 1$. From a security point of view, the problem is more interesting when the system is unstable, as otherwise the eavesdropper can always predict that a stable system is close to equilibrium without even eavesdropping.

The sensor communicates over a channel with two outputs/receivers as shown in Figure 1. The input to the channel is denoted by $\hat{y}(k)$. The first output, denoted by $\hat{y}_1(k)$, is the authorized one to the user, while the second, denoted by $\hat{y}_2(k)$, is the unauthorized one to the eavesdropper. The communication is organized in packets, which consist of sufficiently large number of bits to neglect quantization errors (Sinopoli et al., 2004; Hespanha et al., 2007; Gatsis et al., 2014).

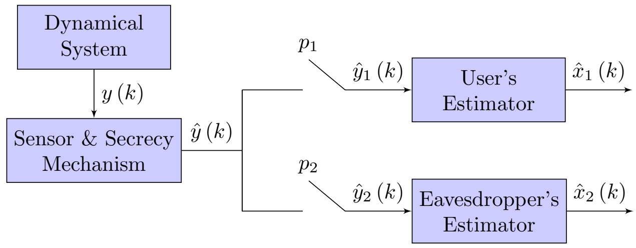

Fig. 1. A sensor collects the output $y(k)$ of the dynamical system (1). Then, it produces a distorted version $\hat{y}(k)$ of the output, and sends it to the channel. The user and the eavesdropper receive the distorted output with probability p_1 and p_2 respectively, and compute the minimum square error estimates $\hat{x}_1(k)$ and $\hat{x}_2(k)$.

Communication with the user is unreliable, i.e., may undergo packet drops. Additionally, communication is not secure against the eavesdropper, i.e., the latter may intercept transmitted packets. Packet drops and packet interceptions are modeled as independent and identically distributed (i.i.d.) over time and across outputs – see Remark 1 for further discussion on this model. In particular, we denote by p_1 the probability that a packet of the authorized output is received by the user. Similarly, p_2 is the probability that a packet of the unauthorized output is intercepted. Thus, the channel model is the following:

$$\hat{y}_i(k) = \begin{cases} \hat{y}(k), & \text{with prob. } p_i \\ \varepsilon, & \text{with prob. } 1 - p_i \end{cases} \tag{2}$$

for $i = 1, 2$, where symbol ε, is used to represent the "no information" outcome.

Our goal is to design a secrecy mechanism at the sensor, so that communication over the channel conveys sufficient plant state information to the user, but limited state information to the eavesdropper. In particular, the sensor transmits a distorted version $\hat{y}(k) \in \mathbb{R}^m \cup \{\varepsilon\}$ of the output $y(k)$ of the system (1) at each time step k over the channel. The symbol ε indicates that no information is sent. The secrecy mechanism dictates how $\hat{y}(k)$ is selected, possibly randomly, given all the available information at the sensor at time k, i.e. past measurements $y(t)$, $t \leq k$ and past messages $\hat{y}(t)$, $t < k$. The mechanism does not know the success of previous packets (values of $\hat{y}_i(k)$); no acknowledgment from the user is assumed.

In this architecture, all system and noise parameters A, C, Q, R, Σ_0 as well as the probabilities p_1, p_2 are assumed to be public knowledge, available to all involved entities, i.e., the sensor, the user, and the eavesdropper (see also Remark 2). Moreover, both the user and the eavesdropper know how the mechanism works. Under those assumptions, both the user and the eavesdropper use their received information to obtain a mean square error estimate of the system state. Let us denote by $\hat{x}_i(k)$ the mean square error estimate at the user ($i = 1$) and the eavesdropper ($i = 2$), defined as:

$$\hat{x}_i(k) = \mathbb{E}\left\{x(k) \,|\, \boldsymbol{\hat{y}}_{i,k}\right\} \tag{3}$$

where $\hat{\boldsymbol{y}}_{i,k} = (\hat{y}_i(0), \ldots, \hat{y}_i(k))$. The corresponding estimation error covariance is given by

$$P_i(k|k) = \mathbb{E}\left\{(x(k) - \hat{x}_i(k))(x(k) - \hat{x}_i(k))^\intercal | \hat{\boldsymbol{y}}_{i,k}\right\} \quad (4)$$

We also define the covariance of the prediction error as

$$P_i(k+1) = AP_i(k|k)A^\intercal + Q \quad (5)$$

with $P_i(0) = \Sigma_0$ at time $k = 0$. Throughout this work, we quantify the uncertainty about the state using the expected value of the prediction error.

We are now ready to introduce our notion of secrecy, requiring that the eavesdropper's uncertainty grows unbounded, whereas the user's uncertainty remains bounded. We term this *perfect expected secrecy.*

Definition 1. (Perfect Expected Secrecy). Given the system (1) and the channel model (2), we say that a secrecy mechanism achieves perfect expected secrecy if and only if, for any initial condition $\Sigma_0 \succ 0$, both of the following conditions hold:

$$\lim_{k\to\infty} \mathbf{tr}\, \mathbb{E}\{P_2(k)\} = \infty \quad (6)$$

$$\limsup_{k\to\infty} \mathbf{tr}\, \mathbb{E}\{P_1(k)\} < \infty \quad (7)$$

where $\mathbf{tr}$ is the trace operator. ◇

This notion of secrecy is asymptotic; the eavesdropper can maintain a trivial open loop estimate $\hat{x}_2(k) = 0$ that has unbounded but finite expected prediction error at any time step k. We also note that the secrecy constraints are required in expectation, not almost surely. By our model, there is always a non-zero probability event that the eavesdropper successfully intercepts a long sequence of messages and intuitively maintains a good state estimate.

To have a well-posed problem, we assume that the user has bounded uncertainty under nominal system operation, i.e. when there are no secrecy concerns and the sensor sends all output measurements.

Assumption 1. If the mechanism $\hat{y}(k) = y(k)$ is employed for all $k \geq 0$, then the user's expected error is bounded

$$\limsup_{k\to\infty} \mathbf{tr}\, \mathbb{E}\{P_1(k)\} < \infty \quad (8)$$

for any initial condition $\Sigma_0 \succ 0$.

The goal of the following section is to propose a simple secrecy mechanism that achieves perfect expected secrecy exploiting the channel model and the system dynamics. Later on, in Section 4 we explore the tradeoff between secrecy against the eavesdropper and utility to the user.

Remark 1. Modeling the user reception as an i.i.d. sequence implies a lossy memoryless channel, commonly assumed in networked control systems (Hespanha et al., 2007). The assumption that packet interception at the eavesdropper is also i.i.d. is novel, and can be similarly thought to model a lossy memoryless channel. In practical scenarios the eavesdropper cannot perfectly intercept all messages, e.g., it overhears the communication from a distance. Randomness in the interception may be attributed to random varying channel conditions of the wireless medium. ◇

Remark 2. The assumption that the system designer knows exactly the eavesdropper's channel model is common in formulations of physical layer security problems (Regalia et al., 2015). In our case, knowing the packet interception rate p_2 (or some upper bound), is less restrictive than knowing the exact eavesdropper's channel model. Alternatively, the value p_2 can be thought of as a level of confidence the system designer has on the ability of an eavesdropper to intercept the messages or not. ◇

3. PERFECT EXPECTED SECRECY

In this section, we explore sufficient conditions, under which, we can achieve perfect expected secrecy. In particular, we propose a simple mechanism, which consists of flipping a coin with success probability p at each time k to decide whether to transmit the sensor's output measurement over the communication channel. If sent, and if the respective packet is not dropped, it also reaches the user and/or the eavesdropper. With probability $1 - p$, on the other hand, no message is sent, hence, neither the user nor the eavesdropper receive any information. Intuitively, the proposed mechanism tries to achieve secrecy by randomly withholding sensor information. By selecting this probability p, we can control, to some extent, the amount of the information availability in both channel outputs, though not independently. Formally, the secrecy mechanism has the following form:

$$\hat{y}(k) = \begin{cases} y(k) \text{ with prob. } p \\ \varepsilon \text{ with prob. } 1-p \end{cases}, \ \forall k \geq 0. \quad (9)$$

With no acknowledgments available, the mechanism's decisions (9) are independent of the packet drops in the channel.

The next theorem states that a sufficient condition for perfect expected secrecy, is that the authorized output of the channel is more reliable than the unauthorized one, i.e. $p_1 > p_2$. If this condition holds, then we can use the secrecy mechanism (9) in order to satisfy (6), (7), by carefully selecting values for the probability p.

Theorem 1. (Conditions for Perfect Secrecy). Consider system (1) and the channel model (2). Under Assumption 1, perfect expected secrecy is achieved within the family of mechanisms (9) if and only if

$$p_1 > p_2. \quad (10)$$

In particular, there exists a probability $p_c \in [0, 1)$ such that all probabilities p satisfying

$$\frac{p_c}{p_1} < p \leq \min\left\{\frac{p_c}{p_2}, 1\right\} \quad (11)$$

are exactly those, which guarantee perfect expected secrecy. ◇

The condition $p_1 > p_2$ is a reasonable requirement for secrecy in many cases of practical interest. For example, as mentioned in Remark 1, when the eavesdropper intercepts the communication from some distance while the user is physically closer to the sensor, the user experiences better reception.

In the rest of this section, we present the proof of Theorem 1 and we characterize the probability p_c. Due to the channel model (2) and the mechanism (9), at each time k, the user receives $y(k)$ with probability $\hat{p}_1 = pp_1$, and gets no information with probability $1 - \hat{p}_1$. Similarly, the eavesdropper receives $y(k)$ with probability $\hat{p}_2 = pp_2$, and no information with probability $1 - \hat{p}_2$. Hence, the estimation problems of the user and the eavesdropper, are actually two separate estimation problems with intermittent observations (Sinopoli et al., 2004). Let $\gamma_i(k) = 1$ when $\hat{y}_i(k) \neq \varepsilon$ and $\gamma_i(k) = 0$, when $\hat{y}_i(k) = \varepsilon$. Then, the optimal estimates (3) follow the expressions for the intermittent Kalman filter given by:

$$\begin{aligned} \hat{x}_i(k+1) &= A\hat{x}_i(k) + \gamma_i(k+1) K_i(k+1) \\ &\times (y(k+1) - CA\hat{x}_i(k)) \end{aligned} \tag{12}$$

for $i = 1, 2$, where $K_i(k) = P_i(k) C^\intercal (CP_i(k) C^\intercal + R)^{-1}$ is the standard Kalman filter gain matrix. The prediction error (5) evolves recursively as:

$$P_i(k+1) = g_{\gamma_i(k)}(P_i(k)) \tag{13}$$

where the function g is defined as

$$\begin{aligned} g_\lambda(X) &= AXA^\intercal + Q \\ &- \lambda AXC^\intercal (CXC^\intercal + R)^{-1} CXA^\intercal \end{aligned} \tag{14}$$

for any $\lambda \in [0, 1]$ and any $X \succeq 0$ in $\mathbb{R}^{n \times n}$. Notice that in contrast to the classical Kalman filter, here $P_i(k)$ is stochastic and depends on the random sequence $\gamma_i(k)$ of successful receptions.

The estimation performance with intermittent observations varies with the probability of correct packet reception, in our case either $\hat{p}_1$ or $\hat{p}_2$. Specifically, there exists a critical probability p_c, which determines a kind of phase transition for the evolution of the expected prediction error covariance matrix $\mathbb{E}\{P_i(k)\}$. If $\hat{p}_i > p_c$, then the expected error covariance matrix is bounded. On the other hand, if $\hat{p}_i \leq p_c$ then the expected error is unbounded. The following lemma extends results of (Sinopoli et al., 2004) in our setup – see Remark 3 .

Lemma 2. Given the system (1) and the secrecy mechanism (9), there exists a critical probability $p_c \in [0, 1)$ such that:

$$\lim_{k\to\infty} \mathbf{tr}\, \mathbb{E}\{P_i(k)\} = \infty \text{ if } \hat{p}_i \leq p_c,\ \forall \Sigma_0 \succ 0 \tag{15}$$

$$\sup_{k\geq 0} \mathbf{tr}\, \mathbb{E}\{P_i(k)\} \leq M_{\Sigma_0}, \text{ if } \hat{p}_i > p_c,\ \forall \Sigma_0 \succ 0 \tag{16}$$

where $i = 1, 2$ and M_{Σ_0} is a positive constant, depending on the initial covariance Σ_0. ◇

Proof. The proof is included in the Appendix.

Now we can take advantage of the phase transition according to Lemma 2 to chose mechanism (9) for perfect secrecy, and thus prove Theorem 1. First, we select p to be small enough so that the eavesdropper's error $\mathbf{tr}\, \mathbb{E}\{P_2(k)\}$ grows unbounded, so that condition (6) of perfect expected secrecy is satisfied. According to (15) in Lemma 2, this is guaranteed by selecting $\hat{p}_2 = pp_2 \leq p_c$. On the other hand, p should not be too small, so that the user's error $\mathbf{tr}\, \mathbb{E}\{P_1(k)\}$ stays bounded and condition (7) of perfect expected secrecy is satisfied. To achieve this, from condition (16) of Lemma 2, we could select $\hat{p}_1 = pp_1 > p_c$. Combining both conditions, and due to the fact that $p \in [0, 1]$, it is sufficient select p within the interval $p_c/p_1 < p \leq \min\{p_c/p_2, 1\}$. What remains to show is that this interval is nonempty. By the condition $p_2 < p_1$, we obtain that $p_c/p_1 < p_c/p_2$. It remains to argue that also $p_c/p_1 < 1$. By Assumption 1, the user's error is finite under no secrecy mechanism – in our case when $p = 1$. By (16) in Lemma 2, this can occur only if $p_1 > p_c$. This completes the sufficiency part of Theorem 1.

Now let us argue about the necessity part of Theorem 1. Both conditions $pp_2 \leq p_c < pp_1$ are necessary within the family of mechanisms (9). If $pp_2 > p_c$ or $pp_1 \leq p_c$, then one of the conditions of perfect expected secrecy is violated, for according to Lemma 2 either the eavesdropper's error is bounded or the user's error grows unbounded. Hence, condition $p_2 < p_1$ is necessary within the family of mechanisms (9), which completes the proof of Theorem 1.

Theorem 1 describes that it is possible to achieve perfect secrecy by selecting the probability p of our mechanism to lie within a specific interval, in particular $(p_c/p_1, \min\{p_c/p_2, 1\}]$. However, this approach might still not be constructive. Computing the critical probability value p_c is hard in general, apart from some special cases (Plarre and Bullo, 2009; Mo and Sinopoli, 2008). In the following section, we explore the tradeoff between secrecy and utility via an optimization framework, and by this we also obtain computationally efficient methods to tune the probability p of our secrecy mechanism.

Remark 3. The presence of the eavesdropper makes the problem different than the one presented in Sinopoli et al. (2004). The main goal in the analysis of the intermittent Kalman filter, has been to guarantee bounded error for a user. In contrast, in our work we also require unbounded error for the eavesdropper. Lemma 2 is an extension, as it sheds more light to the unboundedness case. In Sinopoli et al. (2004), it was only shown that unboundedness occurs for some $\Sigma_0 \succeq 0$, while, here, we prove that it occurs for all $\Sigma_0 \succ 0$.

4. TRADEOFF BETWEEN SECRECY AND UTILITY

The notion of perfect secrecy, according to Definition 1, requires infinite estimation error at the eavesdropper. This might be too conservative in practice. At the same time, the definition requires that the legitimate user has a bounded estimation error, but this might be impractically large sometimes. An apparent tradeoff arises between the requirement for secrecy and the utility to the user. In this section, we seek to explore this tradeoff via an optimization framework.

More specifically, consider our secrecy mechanism (9). If the sensor sends measurements more frequently, by increasing the probability p, the user maintains a better state estimate. At the same time, however, the eavesdropper is able to intercept more messages, hence, the level of secrecy declines. Conversely, secrecy is reinforced and the user's error deteriorates if the sensor withholds information at a higher rate. We are interested, then, in designing our mechanism (9) to minimize the error at the user, as long as the eavesdropper error remains larger than some desired bound. This is an optimization problem of the form

$$\begin{aligned} \underset{p\in[0,1]}{\text{minimize}} \quad & \limsup_{k\to\infty} \mathbf{tr}\left(\mathbb{E}\{P_1(k)\}\right) \\ \text{subject to} \quad & \liminf_{k\to\infty} \mathbf{tr}\left(\mathbb{E}\{P_2(k)\}\right) \geq M \end{aligned} \tag{17}$$

where the design variable is the probability p of mechanism (9). Here, $M > 0$ is a positive constant describing some desired level of eavesdropper's error. This is a relaxed notion of secrecy as compared to the perfect secrecy in Definition 1. The latter can be recovered as $M \to \infty$.

Unfortunately, we cannot express the objective function or the constraint of problem (17) as a function of our secrecy mechanism (9) in closed form. The reason is that, to the best of our knowledge, there are no closed form expressions for the expected error of the intermittent Kalman filter. Nonetheless, we can exploit well-known upper and lower bounds on the error introduced in Sinopoli et al. (2004). In particular, we have the following two results.

Proposition 3. Consider system (1) and secrecy mechanism (9). Let p_ℓ be a probability defined as follows

$$p_\ell = 1 - 1/\rho^2(A) \tag{18}$$

Then, the eavesdropper's error is asymptotically lower bounded by

$$\liminf_{k\to\infty} \mathbf{tr}\,\mathbb{E}\left\{P_2\left(k\right)\right\} \geq \mathbf{tr}\,S\left(p\right). \tag{19}$$

If $pp_2 > p_\ell$, $S\left(p\right)$ is defined as the positive definite solution of

$$S\left(p\right) = \left(1 - pp_2\right) AS\left(p\right) A^\mathsf{T} + Q \tag{20}$$

while if $pp_2 \leq p_\ell$, $\mathbf{tr}\,S\left(p\right)$ is defined to be ∞. ⋄

Proposition 4. Consider system (1) and mechanism (9). Let p_u be a probability defined as

$$p_u = \inf\left\{\lambda \in [0,1]:\ \exists X \text{ with } X \succeq g_\lambda\left(X\right)\right\} \tag{21}$$

where $g_\lambda\left(X\right)$ is defined in (14). Then, the user's estimation error is asymptotically upper bounded by:

$$\limsup_{k\to\infty} \mathbf{tr}\,\mathbb{E}\left\{P_1\left(k\right)\right\} \leq \mathbf{tr}\,V\left(p\right). \tag{22}$$

If $pp_1 > p_u$, $V\left(p\right)$ is defined as the positive definite solution of:

$$V\left(p\right) = g_{pp_1}\left(V\left(p\right)\right), \tag{23}$$

while if $pp_1 \leq p_u$, $\mathbf{tr}\,V\left(p\right)$ is defined to be ∞. ⋄

Propositions 3, 4 follow from the proofs of Theorems 3, 4 in (Sinopoli et al., 2004).

We can utilize the bounds of the preceding propositions to design a desired secrecy mechanism, by relaxing the optimization problem (17). In particular, we relax the constraint of (17) by requiring that the lower bound $S\left(p\right)$ on the eavesdropper's estimation error is larger than the desired value M. Moreover we relax the objective with the upper bound $V\left(p\right)$ on the user's estimation error. Thus, the relaxation of problem (17) has the form:

$$\begin{aligned} &\underset{p\in[0,1]}{\text{minimize}} && \mathbf{tr}\,V\left(p\right) \\ &\text{subject to} && \mathbf{tr}\,S\left(p\right) \geq M \end{aligned} \tag{24}$$

Even though Problem (24) is not convex, it has a specific structure that allows us to solve it efficiently. In particular, both the objective value $\mathbf{tr}\,S\left(p\right)$ and the constraint $\mathbf{tr}\,V\left(p\right)$ are monotonically decreasing functions with respect to p (see also Lemma 6 in the Appendix). The next theorem hence explicitly describes the optimal solution.

Theorem 5. Consider system (1) and the mechanism (9). Let $\mathbf{tr}\,S\left(p\right)$, $\mathbf{tr}\,V\left(p\right)$ be the lower and upper bounds, as defined in Proposition 3 and Proposition 4 respectively. Then, the optimal solution of problem (24) is given by

$$p^\star = \max\{p \in [0,1]:\ \mathbf{tr}\,S\left(p\right) \geq M\}. \tag{25}$$

Proof. In Lemma 6, included in the Appendix, we proved that $\mathbf{tr}\,S\left(p\right)$, $\mathbf{tr}\,V\left(p\right)$ are non-increasing functions of p. As a result, problem (24) is equivalent to the optimization problem $\max\{p \in [0,1]:\ \mathbf{tr}\,S_2\left(p\right) \geq M\}$. □

Capitalizing on the above result and monotonicity, we may further devise an algorithm to find the optimal solution p^* of problem (24), based on a bisection search. This process is presented in Algorithm 1. It takes as inputs the system and noise parameters A, C, Q, R, probability p_2, the desired bound $M > 0$ on the eavesdropper's error and a positive constant ε. This constant ε represents the absolute tolerance within which we want to compute $p^\star$.

The steps of the algorithm are the following. First, the probability p_ℓ is computed, which according to (18), depends only on the spectral radius of matrix A, and, thus, can be readily computed. Then, a bisection on the interval $p \in [0,1]$ is performed in order to solve (25). At each bisection step, the algorithm evaluates the function $\mathbf{tr}\,S(p)$ at the midpoint p of the current interval. If $pp_2 \leq p_l$, then the algorithm sets $\mathbf{tr}\,S(p) = \infty$. Otherwise, the algorithm solves the linear matrix equality (20) with respect to $S(p)$ for the given p, e.g., solved as a Lyapunov equation. Then, the computed value $\mathbf{tr}\,S(p)$ is compared with the desired value M, and the half-interval for the next bisection step is selected, based on the fact that $\mathbf{tr}\,S(p)$ is decreasing in p. Algorithm 1 terminates when $p^\star$ is known to lie within an interval of ε tolerance. The algorithm terminates after at most $-\log_2\left(\varepsilon\right)$ iterations, since at each iteration, the bisection search ends up with half of the interval of the previous step.

Algorithm 1 Optimal probability $p^\star$ of problem (24)

Input: A, C, Q, R, p_2, M, ϵ
Output: optimal solution $p^\star$

Compute p_ℓ.
Do bisection on p. Values u and l are the upper and lower bounds in every step of the bisection.
Set $\ell = 0$, $u = 1$ for the initial bounds.
while $|u - \ell| < \epsilon$ **do**
    Set $p = \left(\ell + u\right)/2$
    Set $\mathbf{tr}\,S\left(p\right) = \infty$ if $pp_2 \leq p_l$. Else, compute $\mathbf{tr}\,S\left(p\right)$.
    if $\mathbf{tr}\,S\left(p\right) < M$ **then**
        Set $u = p$
    else
        Set $l = p$
    end if
end while
return p

Finally, we point out that after Algorithm 1 returns the optimal probability $p^\star$, we may also evaluate the optimal objective value $\mathbf{tr}\,V\left(p^\star\right)$ of (24). As long as $p^\star p_1 > p_u$ this optimal value is finite and can be computed by solving (23) with respect to $V\left(p^\star\right)$. The value of probability p_u can be computed by a quasi-convex optimization problem, while equation (23) can be solved by a semidefinite program (see Sinopoli et al. (2004) for both methods).

5. EXAMPLES

In this section, we present analytical expressions for the problem of secrecy in the special case of scalar systems, as well as numerical results for a second order system. In the former case, we are interested in the tradeoff between secrecy and utility to the user, as captured by the mechanisms p^* that solve problem (24). In particular, we examine how the solution depends on the channel's parameters and how the solution varies, as the required bound M on the eavesdropper error varies.

Here all parameters A, C, R, Q, $\Sigma_0 \in \mathbb{R}$ as well as the upper bound $V(p) \in \mathbb{R}$ on the user's error and the lower bound $S(p) \in \mathbb{R}$ on the eavesdropper's error, as defined in (20) and (23) respectively, are scalars. Moreover, the scalar system is one of the special cases where we know exactly the critical probability and $p_\ell = p_u = p_c = 1 - 1/A^2$ (Sinopoli et al., 2004). The interesting case to study is when secrecy is needed, i.e. when the eavesdropper has bounded expected error without any secrecy mechanism ($p = 1$). According to Lemma 2, this happens when $p_2 > p_c = 1 - 1/A^2$. Also, since $S(p) \geq S(1)$, the secrecy constraint $S(p) > M$ with $M < S(1)$ is trivially satisfied. Thus, we also assume that $M \geq S(1)$ in this section.

From (19), we can find that $S(p) = Q/\big(1 - (1 - pp_2)A^2\big)$. Thus, by Theorem 5, we deduce that the optimal secrecy mechanism $p^\star$ that solves problem (24) is the probability value that satisfies $S(p^\star) = M$. So, for the scalar case, we can obtain the following closed form expression for p^*:

$$p^\star = \frac{p_c}{p_2} + \frac{Q}{Mp_2A^2}. \tag{26}$$

Observe that the first term $\frac{p_c}{p_2}$, on the right hand side, is a probability that guarantees perfect expected secrecy if $p_1 > p_2$ (see the discussion of Section 3). Hence as $M \to \infty$, we may recover perfect expected secrecy.

Next, we evaluate the upper bound on the user's error when using the above mechanism p^*, i.e., the optimal objective value of problem (24). Recall that if $p^\star p_1 \leq p_c$, then $V(p^\star) = \infty$. Otherwise, if $p^\star p_1 > p_c$, then we can solve for the positive solution $V(p^\star)$ of the (quadratic) equation (23). So, in the latter case, we obtain:

$$V(p^\star) = \frac{\beta + \sqrt{\beta^2 + 4QRC^2\left[(A^2-1)\left(\frac{p_1}{p_2} - 1\right) + \frac{p_1}{Mp_2}\right]}}{2C^2\left[(A^2-1)\left(\frac{p_1}{p_2} - 1\right) + \frac{p_1}{Mp_2}\right]}$$

where $\beta = \left(A^2 - 1\right)R + Q^2C$.

The expression for $V(p^\star)$ as a function of M captures the tradeoff between the guaranteed utility to the user and the secrecy level M at the eavesdropper. Figure 2 plots this expression as a function of M for different values of the channel probabilities p_1, p_2. The system parameters were $A = 1.2$, $C = 1$, $Q = 1$ and $R = 1$. Interestingly, as the ratio p_1/p_2 increases, the tradeoff between secrecy and efficiency improves, meaning that a better estimation error can be guaranteed at the user at a given secrecy level.

In the second case, we study the estimation performance over time. We consider a second order system with parameters $A = \begin{bmatrix} 1.2 & 1 \\ 0 & 1.1 \end{bmatrix}$, $C = [1\ 0]$, $R = 1$, $Q = \Sigma_0 = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 2 \end{bmatrix}$ and probabilities $p_1 = 0.9$, $p_2 = 0.6$. In Figure 3, we compare the user's and eavesdropper's expected prediction error $\mathbf{tr}\,\mathbb{E}\{P_i(k)\}$, for $i = 1, 2$, under the mechanism $p = p_\ell/p_2 = 0.51$. The expected errors are approximately computed via Monte Carlo simulation with

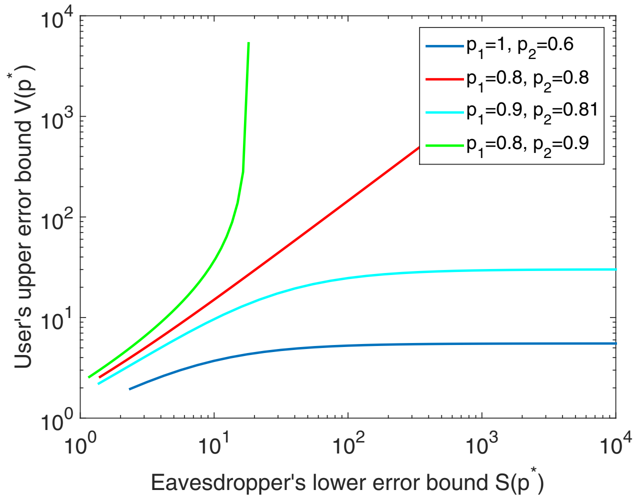

Fig. 2. This figure shows different tradeoffs between optimum user's upper bound $V(p^\star)$ and eavesdropper's lower bound $S(p^\star)$ in problem (24), as we vary p_1, p_2. As the ratio p_1/p_2 increases, tradeoff between secrecy and estimation efficiency improves. As we observe, tradeoff is in favour of the user when $p_1 > p_2$, but in favor of the eavesdropper when $p_2 > p_1$.

10000 iterations. We observe that perfect expected secrecy is achieved; the user's expected error is bounded while the eavesdropper's grows unbounded with an exponential rate. Until now we have only explored how the user and eavesdropper errors behave in expectation. Let us now present one random time sample of the behavior of the actual estimation errors $\|\hat{x}_i(k) - x(k)\|_2$, for user and eavesdropper $i = 1, 2$ respectively for the same second-order system. We compare the estimation errors between the user and the eavesdropper for two cases; i) when no secrecy mechanism is employed ($p = 1$), ii) using the secrecy mechanism with probability $p = 0.51$. We use the same randomly generated sample for the noise and packet drop sequences, in both cases.

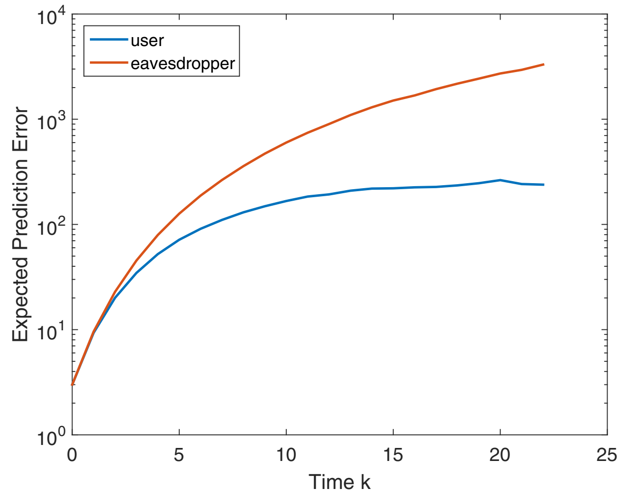

Fig. 3. This figure compares the user's and eavesdropper's expected prediction error. The expected values are approximated via Monte Carlo simulation with 10000 iterations. Notice that perfect expected secrecy is achieved.

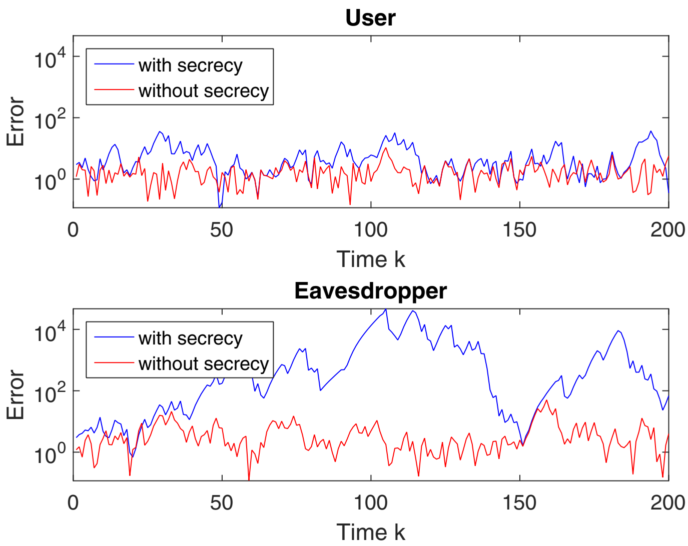

Fig. 4. This figure shows how the estimation errors $\|\hat{x}_1(k) - x(k)\|_2$, $\|\hat{x}_2(k) - x(k)\|_2$ evolve over time for two cases. In the first case, no secrecy mechanism is employed ($p = 1$). In the second case, a secrecy mechanism with $p = 0.51$ is used. The same randomly generated sequence of noises and packet drops was used in both cases. The packet drops of the user are independent of the packet drops of the eavesdropper.

In Figure 4 we see that there is benefit in employing the secrecy mechanism. When no secrecy is used, the user's and eavesdropper's errors are comparable, and the confidentiality of the system's state is compromised. The time average $\frac{1}{N}\sum_{k=1}^{N}\|\hat{x}_i(k) - x(k)\|_2$ of the errors, for $N = 200$, are 1.9 for the user and 4.6 for the eavesdropper respectively. However, in the later case, where a secrecy mechanism with $p = 0.51$ is used, the relative gain we have improves. The time average of the absolute errors are 6.2 for the user and 2922.7 for the eavesdropper respectively.

However, the eavesdropper's error drops to small values infinitely often, as one can see in Figure 4. This occurs because whenever the eavesdropper successfully intercepts a measurement, its prediction covariance matrix decreases dramatically – see also Wu et al. (2014) for a related discussion. In other words, the fact that the eavesdropper's error $\mathbf{tr}\,\mathbb{E}\{P_2(k)\}$ grows unbounded in expectation, by our secrecy definition, does not imply that the actual sample $\mathbf{tr}\,P_2(k)$ is always large.

6. CONCLUSION

In this paper, we considered the novel problem of remote estimation in the presence of an eavesdropper. Requiring confidentiality makes the problem challenging, and forces us to modify the communication scheme, by adding a secrecy mechanism to the sensor. A simple secrecy mechanism, which randomly withholds measurements, can guarantee perfect expected secrecy, if the user's reception ability is better than the eavesdropper's. This mechanism creates a tradeoff between confidentiality and utility, which can be approximated by solving an optimization problem. The numerical results of the second order system show that our proposed mechanism can be useful in practice. Still, it guarantees secrecy in expectation – with nonzero probability the eavesdropper's uncertainty can be small. In future work, we will address more general channel models and alternative mechanisms for higher performance. We will also study whether the condition, that the user's rate is higher than the eavesdropper's, is necessary for general mechanisms. Other extensions include the case where the sensor receives packet acknowledgements from the user.

REFERENCES

Amin, S., Cárdenas, A.A., and Sastry, S.S. (2009). Safe and secure networked control systems under denial-of-service attacks. In *International Workshop on Hybrid Systems: Computation and Control*, 31–45. Springer Berlin Heidelberg.

Cárdenas, A.A., Amin, S., and Sastry, S. (2008). Research challenges for the security of control systems. In *HotSec*.

Fawzi, H., Tabuada, P., and Diggavi, S. (2014). Secure estimation and control for cyber-physical systems under adversarial attacks. *IEEE Transactions on Automatic Control*, 59(6), 1454–1467.

Gatsis, K., Ribeiro, A., and Pappas, G.J. (2014). Optimal power management in wireless control systems. *IEEE Transactions on Automatic Control*, 59(6), 1495–1510.

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

Hespanha, J.P., Naghshtabrizi, P., and Xu, Y. (2007). A survey of recent results in networked control systems. *Proceedings of the IEEE*, 95(1), 138.

Katz, J. and Lindell, Y. (2014). *Introduction to modern cryptography*. CRC press.

Le Ny, J. and Pappas, G.J. (2014). Differentially private filtering. *IEEE Transactions on Automatic Control*, 59(2), 341–354.

Leong, A.S., Quevedo, D.E., Dolz, D., and Dey, S. (2017). Remote state estimation over packet dropping links in the presence of an eavesdropper. *arXiv preprint arXiv:1702.02785*.

Li, H., Lai, L., and Zhang, W. (2011). Communication requirement for reliable and secure state estimation and control in smart grid. *IEEE Transactions on Smart Grid*, 2(3), 476–486.

Liang, Y., Poor, H.V., and Shamai, S. (2008). Secure communication over fading channels. *IEEE Transactions on Information Theory*, 54(6), 2470–2492.

Mo, Y., Hespanha, J.P., and Sinopoli, B. (2014). Resilient detection in the presence of integrity attacks. *IEEE Transactions on Signal Processing*, 62(1), 31–43.

Mo, Y. and Sinopoli, B. (2008). A characterization of the critical value for Kalman filtering with intermittent observations. In *47th IEEE Conference on Decision and Control , 2008*, 2692–2697.

Oggier, F. and Hassibi, B. (2011). The secrecy capacity of the mimo wiretap channel. *IEEE Transactions on Information Theory*, 57(8), 4961–4972.

Pajic, M., Weimer, J., Bezzo, N., Tabuada, P., Sokolsky, O., Lee, I., and Pappas, G.J. (2014). Robustness of attack-resilient state estimators. In *ICCPS'14: ACM/IEEE 5th International Conference on Cyber-Physical Systems (with CPS Week 2014)*, 163–174. IEEE Computer Society.

Plarre, K. and Bullo, F. (2009). On Kalman filtering for detectable systems with intermittent observations. *IEEE Transactions on Automatic Control*, 54(2), 386–390.

Regalia, P.A., Khisti, A., Liang, Y., and Tomasin, S. (Eds.) (2015). Secure Communications via Physical-Layer and Information-Theoretic Techniques [Special Issue]. *Proceedings of the IEEE*, 103(10).

Sandberg, H., Amin, S., and Johansson, K.H. (Organizers) (2015). Cyberphysical Security in Networked Control Systems [Special Issue]. *IEEE Control Systems*, 35(1).

Shannon, C.E. (1949). Communication theory of secrecy systems. *Bell system technical journal*, 28(4), 656–715.

Sinopoli, B., Schenato, L., Franceschetti, M., Poolla, K., Jordan, M.I., and Sastry, S.S. (2004). Kalman filtering with intermittent observations. *IEEE transactions on Automatic Control*, 49(9), 1453–1464.

Wiese, M., Johansson, K.H., Oechtering, T.J., Papadimitratos, P., Sandberg, H., and Skoglund, M. (2016). Uncertain wiretap channels and secure estimation. In *2016 IEEE International Symposium on Information Theory (ISIT)*, 2004–2008.

Wu, J., Shi, G., and Johansson, K.H. (2014). Probabilistic convergence of Kalman filtering with nonstationary intermittent observations. In *53rd IEEE Conference on Decision and Control*, 3783–3788. IEEE.

Wyner, A.D. (1975). The wire-tap channel. *The Bell System Technical Journal*, 54(8), 1355–1387.

Zou, Y., Zhu, J., Wang, X., and Hanzo, L. (2016). A survey on wireless security: Technical challenges, recent advances, and future trends. *Proceedings of the IEEE*, 104(9), 1727–1765.

Appendix A. PROOF OF RESULTS

Proof of Lemma 2

Since (A, C) is detectable, $\left(A, Q^{\frac{1}{2}}\right)$ is controllable and A is unstable, the result of Theorem 2 in (Sinopoli et al., 2004) readily applies. By this result, there exists a $p_c \in [0, 1)$, such that

$$\exists \Sigma_0 \succeq 0 : \lim_{k\to\infty} \mathbf{tr}\, \mathbb{E}\{P_i(k)\} = \infty, \text{ if } \hat{p}_i \leq p_c \tag{A.1}$$

$$\sup_{k\geq 0} \mathbf{tr}\, \mathbb{E}\{P_i(k)\} \leq M_{\Sigma_0}, \text{ if } \hat{p}_i > p_c, \forall \Sigma_0 \succeq 0 \tag{A.2}$$

where $i = 1, 2$ and M_{Σ_0} is a positive constant, depending on the initial covariance Σ_0. This result directly implies the statement (16). Hence it remains to show that (15) holds for any positive definite initial condition $\Sigma_0 \succ 0$.

Let $\Sigma'_0 \succeq 0$ be one initial condition with $P'_i(k)$ the respective error sequence, for which $\lim_{k\to\infty} \mathbf{tr}\, \mathbb{E}\{P'_i(k)\} = \infty$, according to (A.1). Also, let $P_i(k)$ be the error sequence with some arbitrary initial condition $\Sigma_0 \succ 0$. Both covariance errors $P'_i(k)$ and $P_i(k), k \geq 0$ are random variables that depend on the sequence of packet successes $\gamma_i(k), k \geq 0$ according to (13). To complete the proof, it is sufficient to find a positive constant $\beta > 1$ such that $\mathbf{tr}\, \mathbb{E}\{P'_i(k)\} \leq \beta\, \mathbf{tr}\, \mathbb{E}\{P_i(k)\}$ for all $k \geq 0$. Then, in the limit $\lim_{k\to\infty} \mathbf{tr}\, \mathbb{E}\{P_i(k)\} \geq (1/\beta) \lim_{k\to\infty} \mathbf{tr}\, \mathbb{E}\{P'_i(k)\} = \infty$ and the result (15) follows. Note that as $\Sigma_0 \succ 0$, we can find a large enough positive constant $\beta > 1$ such that $\beta\Sigma_0 \succeq \Sigma'_0$. Then, we claim that for any fixed packet success sequence $\gamma_i(k)$, $k \geq 0$, we have $P'_i(k) \preceq \beta P_i(k)$ for all $k \geq 0$. Hence, averaging over all possible success sequences, we obtain $\mathbf{tr}\, \mathbb{E}\{P'_i(k)\} \leq \beta\, \mathbf{tr}\, \mathbb{E}\{P_i(k)\}$.

Finally, to prove that $P'_i(k) \preceq \beta P_i(k)$ for all $k \geq 0$, we use induction. At $k = 0$ we have $\Sigma'_0 \preceq \beta\Sigma_0$, which is the same as $P'_i(0) \preceq \beta P_i(0)$. Suppose that $P'_i(k) \preceq \beta P_i(k)$. Then, at $k + 1$ we employ the relation (13)- (14). Notice that in (14), $g_\lambda(X)$ is an increasing function of X with respect to the positive semidefinite cone for any $\lambda \in [0, 1]$ (see Lemma 1 in Sinopoli et al. (2004)). We can also see that $g_\lambda(X)$ is an increasing function of Q and R as well. Hence, for $\beta > 1$, since $\beta Q \succ Q$, $\beta R \succ R$, we obtain, that $\beta g_\lambda(X) \succeq g_\lambda(\beta X)$. Then, at time k, from (13), we have

$$\begin{aligned}\beta P_i(k+1) =& \beta g_{\gamma_i(k)}(P_i(k)) \\ \succeq& g_{\gamma_i(k)}(\beta P_i(k)) \\ \succeq& g_{\gamma_i(k)}(P'_i(k)) = P'_i(k+1).\end{aligned}$$

The first inequality comes from the fact that $\beta > 1$, while the second comes from monotonicity of $g_\lambda(X)$ with respect to $X$and the induction hypothesis at time step $k-1$. □

Lemma 6. The bounds $\mathbf{tr}\, S(p)$, $\mathbf{tr}\, V(p)$, defined in (20) and (23), defined to be ∞ when $pp_2 \leq p_l$ and $pp_1 \leq p_u$ respectively, are non-increasing functions of p.

Proof. First, we prove that $\mathbf{tr}\, S(p)$ is non-increasing with p. From the proof of Theorem 3 in (Sinopoli et al., 2004), we know that when $pp_2 > p_l$, then $S(p) = \lim_{k\to\infty} S_{k+1}(p)$, is the limit of a sequence of matrices $S_{k+1}(p) = m_p(S_k(p))$, where $m_p(X) = (1 - pp_2) AXA^\intercal + Q$ and $S_0(p) = 0$. Now observe that $m_p(X)$ is decreasing with p and increasing with X (with respect to the positive semidefinite cone). Given two different probabilities $1 \geq \lambda_1 \geq \lambda_2 > p_l/p_2$, we will use the monotonicity of $m_p(X)$ to show by induction that $S_k(\lambda_1) \preceq S_k(\lambda_2)$, for all $k \geq 0$. For $k = 0$, it is true, since $S_0(\lambda_i) = 0$, $i = 1, 2$. Now, assume that $S_{k-1}(\lambda_1) \preceq S_{k-1}(\lambda_2)$ holds. Then,

$$\begin{aligned}S_k(\lambda_1) &= m_{\lambda_1}(S_{k-1}(\lambda_1)) \\ &\preceq m_{\lambda_2}(S_{k-1}(\lambda_1)) \\ &\preceq m_{\lambda_2}(S_{k-1}(\lambda_2)) = S_k(\lambda_2)\end{aligned}$$

where the first inequality comes from monotonicity of $m_p(X)$ with respect to p and $\lambda_1 \geq \lambda_2$; the second inequality comes from monotonicity with respect to X and the induction hypothesis. Taking the trace in both sides we have that $\mathbf{tr}\, S_k(\lambda_1) \leq \mathbf{tr}\, S_k(\lambda_2)$ and as $k \to \infty$, we obtain $\mathbf{tr}\, S(\lambda_1) \leq \mathbf{tr}\, S(\lambda_2)$. Since $\mathbf{tr}\, S(p)$ is extended to ∞ for $pp_2 \leq p_l$, $\mathbf{tr}\, S(p)$ is non-increasing in all of $[0, 1]$.

The proof that $\mathbf{tr}\, V(p)$ is non-increasing with p is similar. From the proof of Theorem 4 in Sinopoli et al. (2004), we have that when $pp_1 > p_u$, then $V(p)$ is the limit of a sequence of matrices $V_{k+1}(p) = g_{pp_1}(V_k(p))$, with $V_0(p) = \Sigma_0$. Function $g_{pp_1}(X)$ is defined in Proposition 4 and it is also decreasing with respect to p and increasing with respect to X (see Lemma 1 in (Sinopoli et al., 2004)). Repeating exactly the same argument as before, we can show that if $1 \geq \lambda_1 \geq \lambda_2 > p_u/p_1$, then $V_k(\lambda_1) \preceq V_k(\lambda_2)$, for any $k \geq 0$. Therefore, as $k \to \infty$, we obtain $\mathbf{tr}\, V(\lambda_1) \leq \mathbf{tr}\, V(\lambda_2)$. Since $V(p)$ is extended to ∞ for $pp_1 \leq p_u$, then $V(p)$ is non-increasing in all of $[0, 1]$. □