Parameter Privacy versus Control Performance: Fisher Information Regularized Control

Ingvar Ziemann, Henrik Sandberg

► To cite this version:

Ingvar Ziemann, Henrik Sandberg. Parameter Privacy versus Control Performance: Fisher Information Regularized Control. 2020. hal-02318237v2

HAL Id: hal-02318237

<https://hal.science/hal-02318237v2>

Preprint submitted on 16 Mar 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Parameter Privacy versus Control Performance: Fisher Information Regularized Control

Ingvar Ziemann, Henrik Sandberg

Abstract—This article introduces and solves a new privacy-related optimization problem for cyber-physical systems where an adversary tries to learn the system dynamics. In the context of linear quadratic systems, we consider the problem of achieving a small cost while balancing the need for keeping knowledge about the model’s parameters private. To this end, we formulate a Fisher information regularized version of the linear quadratic regulator with cheap cost. Here the control operator is allowed to not only control the plant but also mask its state by injecting further noise. Within the class of linear policies with additive noise, we solve this problem and show that the optimal noise distribution is Gaussian with state dependent covariance. Next, we prove that the optimal linear feedback law is the same as without regularization. Finally, to motivate our proposed scheme, we formulate for scalar systems an equivalent maximin problem for the worst-case scenario in which the adversary has full knowledge of all other inputs and outputs. Here, our policies are maximin optimal with respect to maximizing the variance over all asymptotically unbiased estimators.

I. INTRODUCTION

The advent of cyper-physical systems (CPS) poses many new problems to the secure and effective operation of modern control systems [1]. Not only do imminent threats arise due to the potential of adversaries to alter control trajectories but the potential for information about the system to fall into the wrong hands also poses significant risks for both privacy and security. To this end, [2] studies learning-based attacks – attacks where the adversary initially has little or no knowledge of the system dynamics but uses system identification to launch an effective attack. Given such exposure it becomes interesting to ask whether the adversary’s learning attempts are preventable while – importantly – maintaining reasonable control performance.

To make matters precise, this article treats linear stochastic systems of the form

$$x_{t+1} = Ax_t + Bu_t + w_{t+1} + v_{t+1}, \quad (1)$$

where the state $x = (x_t)$ is driven by a sequence of random disturbances $w = (w_t)$. Beside the ordinary control input $u = (u_t)$, which we assume to be linear in the current state x_t , we assume that the control operator also has the

possibility to inject further noise $v = (v_t)$ by at each stage choosing its density, p_t . This ability to inject further noise is not all that different from the watermarking schemes proposed in for instance [3]. In contrast to watermarking however, instead of creating a physical difference in state, it is the increased variance of the adversary’s attack, once based on their parameter estimate, which serves to raise detectability.

In the context of these systems this article studies the trade-off between control performance and the adversary’s ability to estimate the A -matrix given knowledge of the state sequence (x_t) and the control input (u_t) . We assume that our adversary wants to find an estimator $\hat{A}$, that minimizes

$$\text{tr } \mathbf{E}(A - \hat{A})(A - \hat{A})^\top \quad (2)$$

using the input and output sequences, $(x_t), (u_t)$, which they could then later use for malicious purposes. The control operator is thus faced with a certain trade-off – a decision needs to be made whether to optimize control performance or to ensure that as little information as possible about A is available to the adversary. Clearly, also other parameters could be of interest to the adversary, but due to the importance of the A -matrix in control we consider it an interesting case.

We would also like to point out that under the given information structure and stability of the closed loop system, succesful estimation of A will always be possible asymptotically. In this article we discuss how to prevent the adversary from learning for as long as possible, at the smallest possible extra cost. This is very much in line with the work [2] in which a two-phase attack is considered. They consider a model where the adversary first passively obtains a trajectory of length T of the system, and then using these is constrained to execute an undetectable attack in the second phase. Our scheme thus makes it harder for this attacker to remain undetectable for a given sampling period of length T .

Since (2) is hard to maximize directly, we instead optimize a quantity which lower bounds (2), known as the trace of Fisher information of A , $I^T(p)$. To support this relaxation, we prove in Theorem 5.3 that for scalar systems in the worst-case scenario where the adversary knows the disturbance w – for instance it might be them who inject it – then these problems are equivalent.

Considering the trade-off between parameter privacy and control performance, via Fisher information and through a

Ingvar Ziemann (ziemann@kth.se) and Henrik Sandberg (hsan@kth.se) are with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden.

This work was supported in part by the Swedish Research Council (grant 2016-00861), and the Swedish Foundation for Strategic Research (Project CLAS). The authors would also like to express their gratitude to David Umsonst for valuable discussions and feedback on an earlier draft. Moreover, we would like to thank three anonymous referees for their helpful comments.

trade-off parameter λ , yields the relaxed problem

$$\min_{u,p} \limsup_{T \rightarrow \infty} \frac{1}{T} \left(I^T(p_1, \dots, p_T) + \lambda \sum_{t=1}^T \mathbf{E} x_t^\top x_t \right),$$

subject to the dynamics (1). In the class of linear control laws we are able to show that the optimal density for v_{t+1} is Gaussian with covariance proportional to the closed loop reachability Gramian from w and inversely proportional to the current state energy. Next, in Theorem 4.3, we prove that the optimal linear control law is the same as in the unregularized case, $\lambda \rightarrow \infty$. This is then finally related to a maximin problem for a simplified version of the objective in (2) in Theorem 5.3.

The idea to minimize Fisher information dates back to the robust statistics community, see [4] for an overview. More recently, [5], [6] introduce it as a measure of privacy and it is from them we draw much of our inspiration. The authors thereof use Fisher information as a guarantee of state privacy as through the classical Cramér-Rao lower bound. This is also the main motivation for using it in our work, as the Cramér-Rao lower bound states that for any adversary only having access to state and controller information, the trace variance of their estimators at time T is bounded below by $1/I^T$. In contrast to the present work however, they focus on state estimation and do not consider the situation with control.

Other approaches to statistical privacy include differential privacy [7] which has been considered in the context of filtering in [8] and control in [9]–[11]. Similar in spirit to our work [12] also considers a privacy-regularized version of the linear quadratic regulator but instead uses directed information (a variant of Shannon information) as a measure of privacy.

A. Organization

The rest of the paper is organized as follows: In §2, we give a precise formulation of the problem and give some preliminaries. Next, §3 discusses the optimal concealment problem regarding how to find the density of the noise v , which optimally trades between cost and Fisher information. This is then related to the control problem in §4, where we derive the optimal linear control law. In §5, we show that for scalar A , the proposed scheme also has a nice maximin property with respect to the adversary's estimation variance. We work through and plot the 1-dimensional case in more detail in §6. Finally, §7 concludes.

II. PROBLEM FORMULATION

As mentioned in the introduction, our aim is to control the system (1) while keeping information about A private. Our notion of privacy about A is here restricted to keeping the variance of the trace of all unbiased estimators based on knowledge of x and u large. Ideally then, we would want to find v which gives the largest value of

$$\min_A \max_v \text{tr} \mathbf{E}(A - \hat{A})(A - \hat{A})^\top \quad (3)$$

subject to keeping control cost below some previously specified level. Treating this problem directly is nontrivial but since by the Cramér-Rao lower bound [13],

$$\text{tr} \mathbf{E}(A - \hat{A})(A - \hat{A})^\top \geq (I^T)^{-1}$$

a lower bound on the objective (3) may be obtained by instead minimizing the trace of Fisher information, I^T .

Definition 2.1: Trace Fisher information with respect to the A -matrix is given by

$$I^T(p) = \text{tr} \int \nabla_A^2 \ln p(v) dp(v).$$

Using this insight, we first formulate and treat the relaxed version of (3) below, which we dub Fisher Information Regularized Control.

Problem 1: Find the optimal policies of

$$\min_{u \in \mathcal{L}, p \in \mathcal{M}} \limsup_{T \rightarrow \infty} \frac{1}{T} \left(I^T(p) + \lambda \sum_{t=1}^T \mathbf{E} x_t^\top x_t \right),$$

s. t. $x_{t+1} = Ax_t + Bu_t + w_{t+1} + v_{t+1}(p).$

We assume that the state evolves in $\mathbb{R}^n$ so that $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$ and that the control input $u = (u_t)$ is chosen from the set $\mathcal{L}$ of linear time-invariant feedback controllers of the form $u_t = Kx_t$ with $K \in \mathbb{R}^{m \times n}$. The external noise sequence (w_t) is assumed independent and identically distributed Gaussian with mean zero and identity covariance. Moreover, the control operator is allowed to inject further mean zero noise (v_t) by choice of a Markov kernel $p = (p_t)$ from a subset of smooth Markov kernels, $\mathcal{M}$. Note further that we assume for simplicity that the controller is chosen at no external cost – cheap control.

Definition 2.2: The set $\mathcal{M}$ consists of Markov kernels with densities of the form

$$p(v_1, v_2, \dots, v_T | x) = p_1(v_1)p_2(v_2|x_1) \dots p_T(v_T|x_{T-1})$$

where each p_t is a smooth (C^∞) function of its first argument and is a distribution with conditional mean zero.

Note that under Markovianity, trace Fisher information satisfies a chain rule and splits into a sum over conditional marginals. This merits the following definition.

Definition 2.3: Conditional trace Fisher information of p_t with respect to A is given by

$$I_t(p_t) = \text{tr} \int \nabla_A^2 \ln p_t(v_{t+1}) dp_t(v_{t+1}).$$

Problem 1 is essentially a version of the linear quadratic regulator where an added value is placed on the system's A -matrix being hard to estimate given state information. This hardness of estimation is captured by the functional $I^T(p)$ which measures the extent to which knowledge of x and u allows determination of A .

It should be noted that we do not actually know that the optimal causal controller should be linear. Nonetheless, due to prevalence of linear controllers and for reasons of analytical tractability we restrict attention to this class. Moreover, the smoothness assumption on $\mathcal{M}$ is with little loss of generality

since Gaussian mixtures, which are smooth, are dense in the sense of weak convergence in the set of probability measures, [14].

A. Preliminaries from Linear Algebra

For a given matrix M we denote by M^+ any pseudoinverse. That is, M^+ satisfies $M^+M = I - \pi_{\ker M}$ where $\pi_{\ker M}$ is the projection onto the kernel of M . If M is positive semi-definite, the eigenvalues coincide with the singular values and moreover one may define the square root $\sqrt{M}$ which is any matrix $\sqrt{M}$ satisfying $\sqrt{M}^\top \sqrt{M} = M$. For a fixed square matrix L , we denote by Γ_L^t the associated time- t reachability gramian,

$$\Gamma_L^t := \sum_{j=0}^t L^j (L^j)^\top$$

and for $t = \infty$ simply as Γ_L . We would also like to recall that a pair (A, B) is called stabilizable if there exists a matrix K such that $\sigma(A + BK) < 1$. Since we here only consider the case with state cost, the associated infinite horizon variance optimal control law is given directly as $u_t = K^* x_t$, where

$$K^* \in \arg \min_K \text{tr} \Gamma_{A+BK}.$$

III. OPTIMAL CONCEALMENT

Here we derive the optimal noise and control policies for Problem 1. We first take the optimal linear feedback controller as a given and derive the optimal concealing noise as a function thereof. This will later be used in Lemma 4.1 to characterize the optimal linear feedback controller by a finite-dimensional optimization problem.

We shall consider the finite-time version of Problem 1 for fixed T .

Problem 2: Find the optimal policies and associated value of

$$\min_{u \in \mathcal{L}, p \in \mathcal{M}} \frac{1}{T} \left(I^T(p) + \lambda \sum_{t=1}^T \mathbf{E} x_t^\top x_t \right),$$

s. t. $x_{t+1} = Ax_t + Bu_t + w_{t+1} + v_{t+1}(p)$.

Next, using Markovianity and linearity we show that the problem of finding the optimal Markov kernel can be broken down into a family of sub-problems.

Lemma 3.1: For the finite time version, Problem 2, the optimal density¹ at time t , p_t , minimizes

$$\min_{p_t} I_t(p_t) + \lambda \mathbf{E} v_t \Gamma_L^{T-t} v_t$$

where Γ_L is any closed loop reachability gramian induced by L .

¹We are suppressing the dependence on x_t in p_t to avoid cumbersome notation.

Proof: We may write,

$$\begin{aligned} I^T(p) &= \text{tr} \int \nabla_A^2 \ln p(v) dp(v) \\ &= \text{tr} \int \nabla_A^2 \left(\sum_{t=1}^T \ln p_t \right) dp \\ &= \sum_{t=1}^T \mathbf{E} \text{tr} \int \nabla_A^2 \ln p_t dp_t \\ &= \sum_{t=1}^T \mathbf{E} I_t(p_t) \end{aligned}$$

This is essentially the chain rule for Fisher information, and we see that each p_t is only active in one part of the sum.

Next, we isolate the impact of v_t in the quadratic component of the cost. Observe that for fixed $L = A + BK$, the closed loop dynamics are

$$x_{t+1} = \sum_{k=0}^{t+1} L^{t+1-k} (w_k + v_k).$$

Recalling the finite time gramian

$$\Gamma_L^t = \sum_{j=0}^t L^j (L^j)^\top$$

the second part of the cost becomes

$$\sum_{t=1}^T \mathbf{E} x_t' x_t = \mathbf{E} \left(\sum_{k=0}^T w_k' \Gamma_L^{T-k} w_k \right) + \mathbf{E} \left(\sum_{k=1}^T v_k' \Gamma_L^{T-k} v_k \right).$$

Note in particular that p_{t+1} and v_{t+1} only show up in one term of the sum and so varying any other term does not affect optimality. ■

In particular, for the infinite-time problem, as T goes to infinity, the energy from v_t is of the simple form

$$\mathbf{E} v_t \Gamma_L v_t + o(1).$$

Before attempting to solve the partial optimization problem, we note its convexity properties.

Lemma 3.2: The functional

$$I_t(p_t) + \lambda \mathbf{E} v_t \Gamma_L^{T-t} v_t$$

is convex and lower semicontinuous over the set of smooth densities.

Proof: That Fisher information possesses these properties is well-known [4]. The second term obviously satisfies this property and thus so does the sum of both terms. ■

Given convexity, to solve the problem in Lemma 3.1 it suffices to solve the PDE giving the first order condition for optimality, found below. For a formal justification of sufficiency see e.g. [15] or [16].

Lemma 3.3: Given any closed loop L , the optimal noise density for the finite time problem 2, given all information at time t , $p_{t+1} = f^2$, satisfies

$$n \text{tr} (x_t x_t^\top \nabla_v^2 f(v)) + \left(\mu(v) - \frac{1}{4} \lambda v^\top \Gamma_L^{T-t} v \right) f(v) = 0 \quad (4)$$

where the dual variable $\mu(v)$ is chosen such that $\int f^2 dv = 1$ and $\int v f^2 dv = 0$.

Proof: The derivation is similar to that of Theorem 2 in [6] and is omitted. ■

We are now prepared to characterize the optimal noise density in terms of the closed loop reachability gramian.

Theorem 3.4: Given any closed loop L and associated trajectory, the optimal noise density p_t at time t , conditioned on x_t , is Gaussian with mean zero and covariance

$$\sqrt{nx_t(\Gamma_L^{T-t}\lambda)^+x_t^\top}.$$

Before commencing with the proof, let us comment on the interpretation of the result. The fact the optimal (conditional) density is Gaussian can essentially be guessed from maximum entropy principles. Indeed, it is a well known fact that for a given variance (cost), the distribution with unbounded support which maximizes entropy is mean zero normal at that variance [17] - the Gaussian family gives us the most added uncertainty at the best possible price.

Moreover, we see that the noise covariance is inversely proportional to $\sqrt{\Gamma_L}$, which is natural as Γ_L directly measures the input-to-cost of noise. The covariance is also proportional to the current state x_t which is natural since a larger current state leads to a better signal to noise ratio for the upcoming measurement. Note also that $p_t^*(v)$ must result in the same information for any choice of pseudo-inverse. This is a consequence of the fact that the trace Fisher information is really a function of the singular values of the covariance of $v_{t+1} \sim p_t^*(\cdot)$ above.

Proof: To find the solution of the PDE, (4), we make an ansatz of the form

$$f(v) = c_1 e^{-c_2 v^\top Q v}.$$

Showing that this indeed verifies the PDE with the variance stated in the theorem is rather straightforward and the details are thus omitted for brevity. ■

Interestingly, the conditional trace Fisher information functional does not depend on the state.

Lemma 3.5: Conditional trace Fisher information evaluated at p^* is

$$I_t(p^*) = \lambda \text{tr} \Gamma_L^{T-t}.$$

Proof: Write

$$\begin{aligned} I_t &= n \int \text{tr} \left[\frac{1}{p_{t+1}^*(v)} x_t x_t^\top \nabla_v p_{t+1}^* (\nabla_v p_{t+1}^*)^\top \right] dv \\ &= n \int \text{tr} (x_t x_t^\top \lambda \Gamma_L^{T-t} (n x_t x_t^\top)^+) p_{t+1}^* dv \\ &= \lambda \text{tr} \Gamma_L^{T-t} \end{aligned}$$

as required. ■

Together with the Cramér-Rao lower bound, this implies the following.

Corollary 3.6: For any closed loop L and added noise $v \sim p^*$, the adversary's estimation error is lower bounded as

$$\min_{\hat{A}} \text{tr} \mathbf{E}(A - \hat{A})(A - \hat{A}) \geq \frac{1}{\lambda \sum_{t=1}^T \text{tr} \Gamma_L^{T-t}}.$$

This lower bound becomes larger when the adversary does not have knowledge of w , in which case an additional term of order $1/T$ proportional to the covariance of Σ can be added by the data processing inequality for Fisher Information.

In either case, to ensure a constant lower bound, λ must be of order $1/T$. That is, our preference for low state variance is inversely related through the tuning parameter λ to our preference for the adversary to have a bad estimate of A . This may seem somewhat discouraging at first, but should be compared to detector performance. In certain cases, even a small increase in variance can be enough to infer the presence of an adversary which bases its attacks on its estimate of A .

IV. OPTIMAL CONTROL

At this point, we are equipped with the structure of the optimal noise kernel p^* through Theorem 3.4 but do not yet know its covariance as the optimal closed loop L remains unknown. Our task is now to simplify Problem 1 further using Theorem 3.4 to move toward a solution of the entire problem.

We may now insert the solution in Theorem 3.4 using the Lemma in Problem 1 to obtain a finite-dimensional optimization problem.

Lemma 4.1: Given the optimal set of densities of Theorem 3.4, Problem 1 may be written as

$$\begin{aligned} \min_{K \in \mathcal{L}} \quad & \text{tr} \Gamma_L + \limsup_{t \rightarrow \infty} \frac{1}{T} \mathbf{E} \left(\sum_{t=1}^T x_t^\top x_t \right), \\ \text{s. t.} \quad & x_{t+1} = (A + BK)x_t + w_{t+1} + v_{t+1}, \\ & \Gamma_L = \sum_{t=1}^{\infty} ((A + BK)^t)^\top (A + BK)^t. \end{aligned}$$

Proof: We evaluate $I^T(p)$ at the optimal Markov kernel given by Theorem 3.4.

$$I^T(p^*) = \sum_{t=1}^T \mathbf{E} \text{tr} \int \nabla_A^2 \ln p_t^* dp_t^* = \sum_{t=1}^T \mathbf{E} I_t(p_t^*)$$

Now note that each of the $I_t(p_t^*)$ are just (the trace of) Fisher Information as described by Lemma 3.5. Therefore

$$\frac{1}{T} I^T(p^*) = \frac{1}{T} \sum_{t=1}^T \lambda \text{tr}(\Gamma_L^{T-t}) = \lambda \text{tr} \Gamma_L.$$

Adding the quadratic part of the cost, dividing by λ and taking limits, the result follows. ■

Lemma 4.2: Assume that (A, B) is stabilizable and that L is a stable closed loop. Then under the optimal density p^* , the state cost limit as $T \rightarrow \infty$ exists and is equal to

$$\frac{1}{T} \sum_{t=1}^T \mathbf{E} x_t^\top x_t \rightarrow \text{tr}(\Gamma_L) \left(1 + \frac{1 + \sqrt{2\pi\lambda + 1}}{\pi\lambda} \right).$$

In the lemma above the first term is due cost induced by the process noise w and the second term is due to added noise v .

Proof: To see that the sequence above converges, it suffices to note that $\mathbf{E} x_t^\top x_t$ is either constant² or monotone

²which occurs if and only if $A - BK = 0$.

increasing and bounded. Details of this are omitted for brevity.

To compute the limit, write using the law of total probability,

$$\begin{aligned}
& \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \sum_{t=1}^T x_t^\top x_t \\
&= \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbf{E} (w_t^\top \Gamma_L w_t + v_t^\top \Gamma_L v_t) \\
&= \text{tr} \Gamma_L + \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbf{E} \mathbf{E} v_{t-1}^\top \Gamma_L v_t \\
&= \text{tr} \Gamma_L + \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbf{E} \text{tr} \sqrt{\lambda^{-1} x_{t-1}^\top \Gamma_L x_{t-1}} \quad (5)
\end{aligned}$$

Here, we replaced Γ_L^{T-t} by its limit Γ_L , which is justified since the rate of convergence is exponential and the averaging occurs at a linear rate.

To evaluate (5), we note that

$$\mathbf{E} \text{tr} \sqrt{\lambda^{-1} x_{t-1}^\top \Gamma_L x_{t-1}} = \sqrt{\mathbf{E} \lambda^{-1} x_{t-1}^\top \Gamma_L x_{t-1}} \sqrt{\frac{2}{\pi}} \quad (6)$$

where this last equality uses that $\sqrt{\lambda^{-1} x_{t-1}^\top \Gamma_L x_{t-1}}$ is equal in distribution to the absolute value of a Gaussian random variable with variance $\mathbf{E} \lambda^{-1} x_{t-1}^\top \Gamma_L x_{t-1}$. Using (6) we find that (5) is equal to

$$\text{tr} \Gamma_L + \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \sqrt{\mathbf{E} \lambda^{-1} x_{t-1}^\top \Gamma_L x_{t-1}} \sqrt{\frac{2}{\pi}}$$

Now the right and left hand sides are Cesàro limits of functions of $\mathbf{E} x_t^\top x_t$, which exist as true limits. Thus the left hand side is equal to X , satisfying the following system of equations

$$\begin{aligned}
X &= \text{tr} \Gamma_L + \sqrt{\text{tr} \lambda^{-1} \Gamma_L Y} \sqrt{\frac{2}{\pi}}, \\
X &= \text{tr} Y.
\end{aligned}$$

Solving this finishes the proof. $\blacksquare$

Combining lemmas 4.1 and 4.2 shows that the optimal cost is linear in $\text{tr} \Gamma_L$. This cost has exactly the same minimizers as the unregularized cheap control problem. These observations result in the following Theorem.

Theorem 4.3: For any stabilizable (A, B) , a solution of Problem 1 exists and is given by

$$\begin{aligned}
u_t^* &= BK^* x_t \\
p_t^*(v_{t+1}) &\propto \exp \left(-\frac{1}{2} v^\top \sqrt{\lambda \Gamma_L (n x_t x_t^\top) + v} \right),
\end{aligned}$$

where the closed loop is given by $L = A + BK^*$, with

$$K^* \in \arg \min_K \text{tr} \Gamma_L.$$

Remark 4.4: The sole difference in policy between the standard cheap cost problem is thus to add mean zero Gaussian noise with the covariance specified above. Since the state is available, this amounts to sampling white Gaussian

noise and pre-multiplying with a function of the state, which is computationally tractable.

The optimal noise distribution aside, there is an interesting system identification based intuition behind why u_t^* is independent of λ . In system identification, a possibility to efficiently estimate a model parameter is often characterized by a notion known as persistence of excitation [18]. Roughly speaking, the more noise passes through the closed loop system, the easier it is to identify the model's parameters. Since the goal of bringing the plant to a low state variance is exactly the opposite of it being persistently excited, the result is in some sense expected. The only question that remains is in what direction noise should be added, and here as shown in Theorem 3.4, the answer is to add noise in a direction inversely proportional to reachability gramian of the closed loop system weighted by the last state.

V. A PRIVACY-PERFORMANCE GAME

The information bound in Corollary 3.6 tells us that we are stopping the adversary from learning A at a rate of $1/\text{tr} \Gamma_L$. If A is a scalar, $n = 1$, then the trace becomes irrelevant and the bound is attainable (in an asymptotic sense at least).

In this section, we show that for scalar A and in the worst-case scenario, where the adversary has knowledge of the disturbance sequence w , this rate is indeed optimal with respect to the adversary's estimator variance at a given control performance; that is, it Fisher Information Regularized Control is maximin optimal for the following game.

Problem 3: The parameter privacy-performance game is

$$\begin{aligned}
& \sup_{u, v} \inf_{\hat{A}} \limsup_{T \rightarrow \infty} T \mathbf{E} (A - \hat{A})^2, \\
& \text{s. t. } x_{t+1} = Ax_t + Bu_t + w_{t+1} + v_{t+1},
\end{aligned}$$

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \sum_{t=1}^T x_t^\top x_t \leq C.$$

where it is now assumed that $A \in \mathbb{R}$ and under the regularity conditions on $\hat{A}, u$ and v of Assumption 1 below.

Assumption 1: The set of admissible policies for Problem 3 are:

- A1. The control sequence, u , is chosen over all linear time-invariant controllers $\mathcal{L}$.
- A2. The noise sequence, v , is decided by a smooth Markov kernel from $\mathcal{M}$.
- A3. The estimator $\hat{A}$ is asymptotically unbiased and is based on knowledge of x, u, w .
- A4. The controlled Markov chain, $x(u, v)$, has an invariant measure π , which has finite Fisher information with respect to $A \in \mathbb{R}$.
- A5. There exists a π -integrable function M such that

$$\left\| \nabla^3 \log p_t(v|\tilde{A}) \right\| \leq M(v)$$

for all $\tilde{A}$ in a neighborhood of A and some norm $\|\cdot\|$ on $\mathbb{R}$.

The fact that the regularized controller is maximin optimal is based on the following lemma which shows that the opponent can attain the lower bound in Corollary 3.6 by

maximum likelihood estimation, uniformly for all policies admissible to Problem 3.

Lemma 5.1: Under Assumption 1, the adversary's maximum likelihood estimator asymptotically attains the Cramér-Rao lower bound.

Proof: Denote by $\hat{A}_T$ the maximum likelihood estimator and by l^T the associated log-likelihood, given the first T samples. A Taylor expansion of ∇l^T around A , shows that

$$\begin{aligned} \nabla l^t(\hat{A}_T) &= \nabla l^T(A) + \nabla^2 l^T(A)(\hat{A}_T - A) \\ &\quad + (\hat{A}_T - A)^\top \nabla^3 l^T(\tilde{A}_T)(\hat{A}_T - A) \end{aligned}$$

where we have suppressed the dependence on the state sequence x in l^t . Here $\tilde{A}_T$ is enclosed in a spherical shell generated by the intersection of balls of radius $\|A\|$ and $\|\hat{A}_T\|$. Moreover, by optimality of $\hat{A}_t$, the left hand side above is zero, so that

$$\begin{aligned} \sqrt{T}(A - \hat{A}_T) &= \left(\frac{1}{T} \nabla^2 l^T(A) + (\hat{A}_T - A)^\top \frac{1}{T} \nabla^3 l^T(\tilde{A}_T) \right)^{-1} \\ &\quad \times \frac{\nabla l^T(A)}{\sqrt{T}}. \end{aligned} \quad (7)$$

Since $\hat{A}_T \rightarrow A$ a.s. by stability of the closed loop system, we are done if we prove that

$$\frac{1}{T} \mathbf{E}[\nabla l^T(A)(\nabla l^T(A))^\top] \rightarrow I(p^*), \quad (8)$$

$$\frac{1}{T} \nabla^2 l^T(A) \rightarrow I(p^*) \text{ and that,} \quad (9)$$

$$\frac{1}{T} \nabla^3 l^T(\tilde{A}_T) \text{ is bounded a.s.} \quad (10)$$

By Markovianity of (v_t) , the log-likelihood splits into a sum of increments of the form

$$l_t(v_{t+1}|A) = l_t(x_{t+1} - (A + BK)x_t - w_{t+1}).$$

Next, observe that since w is assumed to be Gaussian noise with non-degenerate covariance, the Markov (Harris) chain x is positive recurrent. Using this, we see that each $\nabla l_t(A)$, $\nabla^2 l_t(A)$ and $\nabla^3 l_t(\hat{A}_T)$ are π -integrable functions of the positive Harris chain with invariant measure π and by the (Birkhoff/von Neumann) Ergodic Theorem for positive Harris chains, [19], we see that both (8), (9), and (10) converge to their expectations. The result follows by applying the variance to the limit of (7), which we have just shown to exist. ■

Next, we prove that our proposed strategy through Theorem 4.3 satisfies Assumption 1.

Lemma 5.2: The control and noise policy of Theorem 4.3 satisfy Assumption 1.

Proof: Except for x_t , p_t only depends on t through Γ_L^{T-t} which converges to Γ_L by stability. Using this we see that p_t converges to a stationary conditional distribution. From this it follows that x has an invariant measure. The integrability assumptions are also satisfied since the functions $\nabla^k l_t$ are just weighted fourth moments of x (which is a sum of Gaussian and roots of Gaussian random variables). ■

Using these lemmas, we are able to show that control with Fisher information as a regularizer is actually maximin optimal with respect to Problem 3.

Theorem 5.3: If a solution to

$$C = \Gamma_L + \Gamma_L \frac{1 + \sqrt{2\pi\lambda + 1}}{\pi\lambda}$$

exists for some $\lambda \geq 0$ using the optimal policy given in Theorem 4.3 with the largest such λ is also maximin optimal for Problem 3 under Assumption 1.

Note that the necessity of the solution existing merely means that the control problem of attaining cost C needs to be feasible, $C \geq \Gamma_L$. In this $\lambda \rightarrow \infty$ reduces to the ordinary optimal control problem without regularization.

Proof: According to Lemma 4.2 the optimal cost at a given λ is

$$\Gamma_L + \Gamma_L \frac{1 + \sqrt{2\pi\lambda + 1}}{\pi\lambda}. \quad (11)$$

Since we are constrained to costs less than or equal to C , choosing the largest λ solving (11) equal to C maximizes the inverse of the Fisher information subject to our constraints. By Lemma 5.1, the optimal adversarial estimator has variance precisely equal to this for all admissible strategies of the control operator. Since Lemma 5.2 states the policies of Theorem 4.3 are admissible under Assumption 1, it is maximin optimal for the control operator to maximize this lower bound (as it is the value of the best response). ■

VI. THE 1-D CASE WORKED OUT

To get a feeling for how a system evolves under the regularized problem, let us plot the dynamics for the 1-D case below. Here, we specialize further and assume that the system is governed by

$$x_{t+1} = ax_t + u_t + w_{t+1} + v_{t+1}$$

where w_t has variance 1 and $\lambda = 1$. A Monte Carlo average of the trajectory of this system's cost under the optimal policy in Theorem 4.3, the theoretical average cost and adversarial estimator variance are given below, in Figure 1.

Fig. 1. 1-D Cost and Estimator Variance. The blue line gives the empirical average cost and the black line its eventual limit. The red line is the Cramér-Rao lower bound on the adversary's optimal estimation performance.

We see that the control performance quickly reaches the theoretical long-run average cost predicted by Theorem 5.3.

Meanwhile, the adversary's variance is kept above 0 at a rate of $1/T$.

Next, in Figure 2, we plot the trade-off curve between control performance and asymptotic variance of the adversary's estimator.

Fig. 2. Parameter Privacy versus Control Performance: Here, the colored region represents the achievable region for asymptotic estimator variance and cheap control performance.

This final figure thus characterizes the trade-off between privacy about the scalar a and cheap control performance. Our proposed scheme is able to achieve all rates along the boundary between the two regions.

VII. DISCUSSION AND CONCLUSION

In this article we have characterized the trade-off between cheap control and keeping knowledge about the A -matrix private in a variance optimal manner. We also introduced a computationally tractable procedure to solve the trade-off problem, dubbed Fisher Information Regularized Control. This gives an approach for how to mitigate the initial phase of learning-based attacks, as considered in [2]. In particular, this would lessen the requirements on the detector when the attack is ongoing. It would also be satisfying to show that a maximin-type characterization of the proposed scheme is valid for larger classes of systems.

As for further work, noting the close theoretical analogue to the work of [6], it would be interesting to attempt to generalize the results here to the case where state privacy is the central matter. One might expect similar results although our work would have to be extended to cover partial observability for this to make sense. This however would be interesting in its own right. In a similar vein, it would also be interesting to cover the case where the adversary tries to learn other system parameters or different linear functions thereof. It is also clear to us that we have made a number of simplifying assumptions pertaining to the system's parameters. In particular, we would very much like to extend

our results to the case where the cost is also allowed to depend on the control input, u .

REFERENCES

- [1] M. S. Chong, H. Sandberg, and A. M. Teixeira, "A tutorial introduction to security and privacy for cyber-physical systems," in *2019 18th European Control Conference (ECC)*, IEEE, 2019, pp. 968–978.
- [2] M. J. Khojasteh, A. Khina, M. Franceschetti, and T. Javidi, "Learning-based attacks in cyber-physical systems," *CoRR*, vol. abs/1809.06023, 2018.
- [3] Y. Mo, S. Weerakkody, and B. Sinopoli, "Physical authentication of control systems: Designing watermarked control inputs to detect counterfeit sensor outputs," *IEEE Control Systems Magazine*, vol. 35, no. 1, pp. 93–109, 2015.
- [4] P. J. Huber, *Robust Statistics*. John Wiley & Sons, 2004, vol. 523.
- [5] F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4726–4734, 2017.
- [6] —, "Ensuring privacy with constrained additive noise by minimizing fisher information," *Automatica*, vol. 99, pp. 275–288, 2019.
- [7] C. Dwork, "Differential privacy," *Automata, languages and programming*, pp. 1–12, 2006.
- [8] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2013.
- [9] Y. Wang, Z. Huang, S. Mitra, and G. E. Dullerud, "Entropy-minimizing mechanism for differential privacy of discrete-time linear feedback systems," in *53rd IEEE Conference on Decision and Control*, IEEE, 2014, pp. 2130–2135.
- [10] J. Cortés, G. E. Dullerud, S. Han, J. Le Ny, S. Mitra, and G. J. Pappas, "Differential privacy in control and network systems," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, IEEE, 2016, pp. 4252–4272.
- [11] S. Han and G. J. Pappas, "Privacy in control and dynamical systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 309–332, 2018.
- [12] T. Tanaka, M. Skoglund, H. Sandberg, and K. H. Johansson, "Directed information and privacy loss in cloud-based control," in *2017 American Control Conference (ACC)*, IEEE, 2017, pp. 1666–1672.
- [13] L. Wasserman, *All of statistics: a concise course in statistical inference*. Springer Science & Business Media, 2013.
- [14] P. Billingsley, *Convergence of probability measures*. John Wiley & Sons, 2013.
- [15] D. G. Luenberger, *Optimization by vector space methods*. John Wiley & Sons, 1997.
- [16] E. Zeidler, *Nonlinear functional analysis and its applications: III: variational methods and optimization*. Springer Science & Business Media, 2013.
- [17] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [18] L. Ljung, *System Identification: Theory for the User*. Pearson Education, 1998.
- [19] O. Kallenberg, *Foundations of modern probability*. Springer Science & Business Media, 2006.