Output-Feedback linear quadratic robust control under actuation and deception attacks

João P. Hespanha Shaunak D. Bopardikar

***Abstract*— We consider output-feedback robust control of a linear system subject to disturbances and noise and in presence of an attacker who 1) can corrupt the measured output (deception attack) and, 2) can introduce perturbations to the control signal (actuation attack). We consider an open-loop control problem over a finite horizon which models the scenario where feedback control could be stopped if one is certain that an attack is ongoing. We formulate this problem as a zero-sum game between a defender that selects the control signal based on a measured output and an attacker that selects the attack signals. The game has asymmetric information in that the defender only knows the measured output, whereas the attacker knows additional information, which includes the value of initial conditions and disturbances/measurement noise. The main contributions are (i) sufficient conditions for the existence of a Nash equilibrium corresponding to a saddle-point for the defender and (ii) a computationally efficient procedure to compute a pair of policies that form a Nash equilibrium for the game. We apply the procedure to a finite horizon linear quadratic control problem.**

I. Introduction

Networked control systems (NCS) have become prevalent in several domains over the past few decades. Although the systems enable advancements in performance and features, significant vulnerabilities have been reported in domains such as industrial power plants [1], automotives [2], and water networks [3]. Attackers may leverage flaws in the design of the NCS to launch coordinated deception attacks by covertly modifying the controller signals and the measurement values to ensure that while the system performs abnormally, the measurements appear to be perfectly normal. This paper addresses one such problem of controlling a system under coordinated actuation and deception attacks.

In the CPSs security research literature, authors have been considering various models of attacks. In [4] and [5], the authors have considered *false data injections* on static estimators. This attack is modeled as corruption of measurements that are used for state estimation. Conditions on the system properties that prevent these attacks are derived in [6], where the authors show that an attack exists if and only if the system dynamics have an unstable mode and the associated eigenvector satisfies a technical assumption. More recent work, [7],[8] and [9], provide methods to change the system model so that a class of *stealthy attacks* on the actuators or sensors can be detected. In [10] the authors provide a more general framework to analyze several types of attacks on power systems and networks. In particular, general conditions for attack detection and identifiability for descriptor linear time-invariant systems are defined. Fundamental limits on the reconstruction of the initial state from attacked measurements have been analyzed in [11]. This result has been extended to the case of noise in the sensors along with a coding theoretic interpretation for the noiseless case in [12].

Game theory provides a rigorous framework to model adversarial reasoning problems. Early work in control theory was based on mini-max design, to account for worst-case realizations for the uncertainty [13]. The problem addressed in this paper is also related to a type of game termed as *signaling game*. Recent works in this area have addressed the problem of estimation of a variable using one or multiple sensors which try to deceive a remote receiver, which has access to side information, in believing a different value for the quantity to be measured [14]. [15] analyzed the set-up in multiple dimensions and characterized both, simultaneous and sequential formulations. Game theory has been applied in the past to cases in which the exact statistics of the input and measurement noise are unknown [16], and extended to the case of state constraints [17].

This paper considers the output-feedback robust control of a linear system subject to disturbances and noise that cannot be measured and an attacker that can both, corrupt the measured output (deception attack) and introduce perturbations to the control signal (actuation attack). We consider an open-loop control problem over a finite horizon which, in principle, could be solved repeatedly in a receding horizon fashion.

João P. Hespanha is with the Department of Electrical and Computer Engineering, University of California Santa Barbara, CA USA. Email: hespanha@ece.ucsb.edu. Shaunak D. Bopardikar is with the Department of Electrical and Computer Engineering, Michigan State University. Email: shaunak@egr.msu.edu. This research was partially supported by the ONR grant N00014-16-1-2710 and the NSF grants CNS-1329650 and ECCS-1608880.

We formulate this problem as a game between player P_{U} who selects the control signal U based on a measured output Y and player P_{A} who selects the attack A. We consider a zero-sum game where P_{U} wants to minimize a controlled output Z, whereas P_{A} want to maximize it. However, the game has asymmetric information in that P_{U} only knows Y, whereas P_{A} knows additional information, which includes the value of initial conditions and disturbances/measurement noise. We place no physical constraints on the attack A, other than one that player P_{U} wins if P_{A} selects an attack that would lead to a measured output Y that is incompatible with the absence of an attack. This models a realistic scenario where feedback control would be stopped if one is certain that an attack is taking place.

The main results of this paper are:

1) Sufficient conditions for the existence of a Nash equilibrium corresponding to a saddle-point for P_{U} (Section III), and
2) A computationally efficient procedure to compute a pair of policies (U, A) that form a Nash equilibrium (Section IV).

This paper is organized as follows. Section II presents a formal problem statement. The main results are presented in Sections III and IV. Numerical results on finite horizon control of an LTI system are presented in Section V. Conclusions and future directions are presented in Section VI. Due to space constraints, complete proofs are made available in an online technical report [18].

II. Problem Formulation

We first describe an attack-free version of our robust control problem and then introduce the attack model.

A. The attack-free case

Consider a control problem where we want to select a vector of control *inputs* $U \in \mathbb{R}^{n_u}$ to minimize the norm of a *controlled output* vector

$$Z = Z_0 + F_x x_0 + F_u U + F_d D \in \mathbb{R}^{n_z}, \tag{1}$$

where $D \in \mathbb{R}^{n_d}$ denotes an unknown *disturbance*, $x_0 \in \mathbb{R}^{n_x}$ an unknown *initial state*, and $Z_0 \in \mathbb{R}^{n_z}$ a known constant vector. We shall see in Section II-C that this model is applicable to the control of a linear system, in which case U, D, and Z include the values of the control input, a disturbance/noise, and the controlled output, respectively, over a finite time horizon. In this case, we would be solving an open-loop control problem over a finite horizon. While this problem could be solved repeatedly in a receding horizon fashion to generate a feedback control, here we focus on a single horizon.

To minimize $\|Z\|$, the control designer knows a bound Δ for the norm of the disturbance and has access to a *measured output vector*

$$Y = G_x x_0 + G_d D \in \mathbb{R}^{n_y} \tag{2}$$

that provides information about x_0 and D. For simplicity, we assume that the effect of U has been subtracted from Y. A robust controller would then select U to minimize the following worst-case criteria:

$$\min_{U \in \mathbb{R}^{n_u}} \max_{\substack{\hat{x}_0 \in \mathbb{R}^{n_x},\, \hat{D} \in \mathbb{R}^{n_d}, \\ \|\hat{D}\| \leqslant \Delta,\, Y = G_x \hat{x}_0 + G_d \hat{D}}} \|Z_0 + F_x \hat{x}_0 + F_u U + F_d \hat{D}\|^2,$$

where the inner maximization considers worst-case values for the initial condition/disturbance compatible with the measured output Y. This minimization is finite if

$$\ker G_x \subset \ker F_x,$$

which means that

$$G_x x_0 = 0 \quad \Rightarrow \quad F_x x_0 = 0. \tag{3}$$

This is essentially an observability condition that requires an "unobservable" initial condition x_0 to have no effect on Z.

B. The attack case

In this paper, we are interested in a two-player version of the problem described above: as before one player P_{U} selects the control $U \in \mathbb{R}^{n_u}$ to minimize the norm of the controlled output Z based on the measured output Y, but we now have a second player P_{A} who tries to maximize the norm of Z by injecting an attack $A \in \mathbb{R}^{n_a}$ both in Z and Y to maximize the norm of Z:

$$Z = Z_0 + F_x x_0 + F_u U + F_d D + F_a A \in \mathbb{R}^{n_z}, \tag{4}$$
$$Y = G_x x_0 + G_d D + G_a A \in \mathbb{R}^{n_y}. \tag{5}$$

The injection of A in Z corresponds to an *actuation attack*, where P_{A} tries to directly steer Z, while the injection of A in Y is a *deception attack* where P_{A} tries to compromise the information available to P_{U}.

The game has *asymmetric information* in that P_{A} knows the exact values of the variables x_0 and D, whereas P_{U} only has indirect information about these variables. Specifically, P_{U} knows a constant Δ that upper bounds $\|D\|$ and observes the measured output Y. In a practical problem, the control designer may not know if the system is under attack so we include in our problem formulation a basic attack detection mechanism: if the measured output Y cannot have been produced by the attack-free model (2), then P_{U} declares that it is under attack and wins the game, regardless of the value of Z.

Definition 1 (Nash equilibrium): We say that a pair (A, U) is a *Nash equilibrium* for the game if none of the players regrets its choice once it learns the choice of the other player. Specifically, this means the following:

1) Given that P_{U} selects the action U, P_{A}'s best response is A in the sense that this action maximizes $\|Z\|$, subject to the constraint that its action must result in a measured output Y that is compatible with the attack-free model (2). Specifically, A achieves the maximum of

$$J_a(U, x_0, D) \coloneqq \max_{\substack{A \in \mathbb{R}^{n_a}, \tilde{x}_0 \in \mathbb{R}^{n_x}, \\ \tilde{D} \in \mathbb{R}^{n_d}}} \|Z_0 + F_x x_0 + F_u U + F_d D + F_a A\|^2$$
$$\text{subject to } \|\tilde{D}\| \leqslant \Delta,$$
$$G_x x_0 + G_d D + G_a A = G_x \tilde{x}_0 + G_d \tilde{D}. \qquad (6)$$

Therefore the action A must result in a measured output that is compatible with the attack-free model (2).

2) Given that P_{A} selects the action A (which must be compatible with the attack-free model (2)), $\mathsf{P}_{\mathrm{U}}'s$ best response to the measurement $Y \coloneqq G_x x0 + G_d D + G_a A$ is U in the sense that this action minimizes $\|Z\|$ under worst-case assumptions on x_0, D, and A. Specifically, U achieves the minimum of

$$J_u(Y) \coloneqq \min_{U \in \mathbb{R}^{n_u}} \max_{\substack{\hat{x}_0 \in \mathbb{R}^{n_x}, \hat{D} \in \mathbb{R}^{n_d}, \\ \hat{A} \in \mathbb{R}^{n_a}}} \|Z_0 + F_x \hat{x}_0 + F_u U + F_d \hat{D} + F_a \hat{A}\|^2$$
$$\text{subject to } \|\hat{D}\| \leqslant \Delta,$$
$$Y = G_x \hat{x}_0 + G_d \hat{D} + G_a \hat{A}. \qquad (7)$$

When the min and max in (7) commute, we say that P_{U}'s policy is a *saddle-point*. □

The definition below addresses the case where the optima in (6) and (7) are not global.

Definition 2 (Strict local Nash equilibrium): When the optima in (6) and (7) are strict local optima[1] (rather than global optima), we say that the pair (A, U) is a *strict local Nash equilibrium* (rather than a Nash equilibrium). When the strict local min and max in (7) commute, we say that U is a *strict local saddle-point for* P_{U}. □

[1]We recall that a vector x is a *strict local maximum* to

$$\max_{x \in \mathbb{R}^n,\, g(x) \leqslant 0,\, h(x) = 0} f(x).$$

if there exists an neighborhood $\mathcal{X} \subset \mathbb{R}^n$ of x such that $f(x) > f(\bar{x})$ for every vector $\bar{x} \in \mathcal{X} \backslash \{x\}$ such that $g(\bar{x}) \leqslant 0$ and $h(\bar{x}) = 0$. *Strict local minimum* is defined in an analogous fashion.

Although the two players have opposite goals, the definition of Nash equilibrium in Definitions 1 and 2 do not correspond to a saddle-point equilibria (i.e., they do not correspond to the commutation of min and max operations) because of the information asymmetry.

Remark 1 (Security policy for P_{U}.): Since player P_{U} has access to all the information needed to perform the optimization in (7), this optimization implicitly defines a *policy for* P_{U}, in the sense that it maps each observation Y to an action U. Moreover, since this optimization considers worst-case assumptions on all the variables x_0, D, and A unknown to P_{U}, it results in a *security policy for* P_{U} with *security value* equal to $J_u(Y)$, in the sense that the true cost $\|Z\|$ is never worse than $J_u(Y)$, regardless of the action A selected by P_{A} and the true values of the initial state x_0 and disturbance D. In contrast, (6) does not really define a policy for P_{A} because it only permits the computation of the optimal A once we know the value of U, which is not known a priori by P_{A}. □

Remark 2 (Special case: Bounded initial state): The case of $x_0 \in X \subset \mathbb{R}^n$, where X is a compact set, becomes a special case of the problem formulation. In this case, we can define a single vector $\mathbf{D} = [x_0' \ D']'$ and consider it as the disturbance instead. In this case, the corresponding G_x and F_x will be zero. Further, the optimal solution will always be finite.

C. LQ finite-horizon control of LTI systems

Consider the finite-horizon control of the LTI system

$$x(t+1) = A_x x(t) + B\big(u(t) + d_{\text{input}}(t) + a_{\text{act}}\big), \quad \forall t \in \{-L, \ldots, 0, \ldots, T-1\}, \qquad (8a)$$
$$y(t) = C_y x(t) + d_{\text{noise}}(t) + D_y a_{\text{decep}}(t), \quad \forall t \in \{-L, -L+1, \ldots, 0\}, \qquad (8b)$$
$$z(t) = C_z x(t) + D_z u(t), \quad \forall t \in \{1, 2, \ldots, T\}, \qquad (8c)$$

where $x(t)$ denotes the systems state, $u(t)$ the control input, $d_{\text{input}}(t)$ an (unmeasured) input disturbance, $z(t)$ the output to be regulated, $y(t)$ the measured output, $d_{\text{noise}}(t)$ measurement noise, $a_{\text{decep}}(t)$ a deception attack, and a_{act} a constant actuation attack. The measurements $y(t)$ are collected over a (past) horizon $t \in \{-L, -L+1, \ldots, 0\}$ and we want to minimize the 2-norm of a controlled output $z(t)$ over a (future) horizon $t \in \{1, 2, \ldots, T\}$.

The control of (8) can be mapped to the problem formulation (4)–(5) in Section II-B by defining

$$U \coloneqq \begin{bmatrix} u(0) \\ \vdots \\ u(T) \end{bmatrix}, D \coloneqq \begin{bmatrix} d_{\text{input}}(-L) \\ \vdots \\ d_{\text{input}}(T-1) \\ d_{\text{noise}}(-L) \\ \vdots \\ d_{\text{noise}}(0) \end{bmatrix},$$

$$A \coloneqq \begin{bmatrix} a_{\text{act}} \\ a_{\text{decep}}(-L) \\ \vdots \\ a_{\text{decep}}(0) \end{bmatrix}, x_0 \coloneqq x(-L),$$

$$Y \coloneqq \begin{bmatrix} y(-L) - y_{\text{h}}(-L) \\ \vdots \\ y(0) - y_{\text{h}}(0) \end{bmatrix}, Z \coloneqq \begin{bmatrix} z(1) \\ \vdots \\ z(T) \end{bmatrix}$$

where $y_{\text{h}}(t)$ denotes the value of the measured output $y(t)$ in the absence of a control input. We then obtain the matrices $F_x, F_u, F_d, F_a, G_x, G_d, G_a$ using the variation of constants formula:

$$\begin{aligned} &y(t) - y_{\text{h}}(t) = C_y A_x^{t+L} x(-L) + d_{\text{noise}}(t) + D_y a_{\text{decep}}(t) \\ &+ \sum_{\tau=-L}^{t-1} C_y A_x^{t-1-\tau} B\big(d_{\text{input}}(\tau) + a_{\text{act}}\big) \\ &z(t) = C_z A_x^{t+L} x(-L) + D_z u(t) \\ &+ \sum_{\tau=-L}^{t-1} C_z A_x^{t-1-\tau} B\big(u(\tau) + d_{\text{input}}(\tau) + a_{\text{act}}\big), \end{aligned}$$

which lead to

$$G_x = \begin{bmatrix} C_y \\ C_y A_x \\ \vdots \\ C_y A_x^L \end{bmatrix}, Z_0 = \begin{bmatrix} \sum_{\tau=-L}^{-1} C_z A_x^{-\tau} B u(\tau) \\ \sum_{\tau=-L}^{-1} C_z A_x^{1-\tau} B u(\tau) \\ \vdots \\ \sum_{\tau=-L}^{-1} C_z A_x^{T-1-\tau} B u(\tau) \end{bmatrix},$$

$$G_d = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 & I & 0 & \cdots & 0 \\ C_y B & 0 & 0 & \cdots & 0 & 0 & I & \cdots & 0 \\ C_y A_x B & C_y B & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ C_y A_x^{L-1} B & C_y A_x^{L-2} B & C_y A_x^{L-3} B & \cdots & 0 & 0 & 0 & \cdots & I \end{bmatrix},$$

$$G_a = \begin{bmatrix} 0 & D_y & 0 & 0 & \cdots & 0 \\ C_y B & 0 & D_y & 0 & \cdots & 0 \\ \sum_{\tau=0}^{1} C_y A_x^\tau B & 0 & 0 & D_y & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots \\ \sum_{\tau=0}^{L-1} C_y A_x^\tau B & 0 & 0 & 0 & \cdots & D_y \end{bmatrix},$$

$$F_d = \begin{bmatrix} C_z A_x^L B & C_z A_x^{L-1} B & C_z A_x^{L-2} B & \cdots & 0 & 0 & \cdots & 0 \\ C_z A_x^{L+1} B & C_z A_x^L B & C_z A_x^{L-1} B & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & & \vdots \\ C_z A_x^{L+T-1} B & C_z A_x^{L+T-2} B & C_z A_x^{L+T-3} B & \cdots & C_z B & 0 & \cdots & 0 \end{bmatrix},$$

$$F_a = \begin{bmatrix} \sum_{\tau=0}^{L} C_z A_x^\tau B & 0 & \cdots & 0 \\ \sum_{\tau=0}^{L+1} C_z A_x^\tau B & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ \sum_{\tau=0}^{L+T-1} C_z A_x^\tau B & 0 & \cdots & 0 \end{bmatrix}, F_x = \begin{bmatrix} C_z A_x^{L+1} \\ C_z A_x^{L+2} \\ \vdots \\ C_z A_x^{L+T} \end{bmatrix},$$

$$F_u = \begin{bmatrix} C_z B & D_z & 0 & \cdots & 0 \\ C_z A_x B & C_z B & D_z & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ C_z A_x^{T-1} B & C_z A_x^{T-2} B & C_z A_x^{T-3} B & \cdots & D_z \end{bmatrix}.$$

III. Existence of Nash equilibrium

This section provides sufficient conditions for the existence of a strict local Nash equilibrium.

A. P_{U}'s saddle-point

We start with a result that provides a sufficient condition for the existence of a saddle-point for P_{U}, assuming that P_{A} has selected an attack A that results in a measured output Y that is compatible with the attack-free model (2).

Theorem 1 (P_{U}'s optimization): Suppose that there exist vectors $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u$ and a scalar $\lambda_u > 0$ such that

$$\begin{bmatrix} 2F_a' F_a & 2F_a' F_d & 2F_a' F_x & 2F_a' F_u & G_a' \\ 2F_d' F_a & 2F_d' F_d - 2\lambda_u I_{n_d} & 2F_d' F_x & 2F_d' F_u & G_d' \\ 2F_x' F_a & 2F_x' F_d & 2F_x' F_x & 2F_x' F_u & G_x' \\ 2F_u' F_a & 2F_u' F_d & 2F_u' F_x & 2F_u' F_u & 0 \\ G_a & G_d & G_x & 0 & 0 \end{bmatrix}$$
$$\times \begin{bmatrix} \hat{A} \\ \hat{D} \\ \hat{x}_0 \\ U \\ \nu_u \end{bmatrix} = \begin{bmatrix} -2F_a' Z_0 \\ -2F_d' Z_0 \\ -2F_x' Z_0 \\ -2F_u' Z_0 \\ Y \end{bmatrix}, \tag{9a}$$

and

$$V_u' \left(\begin{bmatrix} F_a' \\ F_d' \\ F_x' \end{bmatrix} \begin{bmatrix} F_a & F_d & F_x \end{bmatrix} - \lambda_u \begin{bmatrix} 0 & 0 & 0 \\ 0 & I_{n_d} & 0 \\ 0 & 0 & 0 \end{bmatrix} \right) V_u < 0, \tag{9b}$$

where V_u is any full column rank matrix such that

$$\operatorname{Im} V_u = \ker \begin{bmatrix} 0 & 2\hat{D}' & 0 \\ G_a & G_d & G_x \end{bmatrix}.$$

Then U is a strict local saddle-point for P_{U} with

$$\Delta = \|\hat{D}\|. \tag{9c}$$

The assumptions of the theorem guarantee that all these conditions are simultaneously satisfied. □

Remark 3: In view of Lemma **??**, a necessary and sufficient condition for the existence of some $\lambda_u > 0$ for which (9b) holds is that

$$\ker \begin{bmatrix} 0 & 2\hat{D}' & 0 \\ G_a & G_d & G_x \end{bmatrix} \cap \ker \begin{bmatrix} 0 & I_{n_d} & 0 \end{bmatrix} = \{0\},$$

or, equivalently,

$$G_a A + G_x x_0 \quad \Rightarrow \quad A = 0,\ x_0 = 0,$$

which means that only the zero attack/initial state pair (A, x_0) can have no net effect on Y. □

Proof of Theorem 1. The vectors $U, \hat{x}_0, \hat{D}, \hat{A}$ are a local strict saddle-point equilibrium for (7) if they are strict local optima for the following coupled optimizations:

$$\|Z_0 + F_x\hat{x}_0 + F_u U + F_d\hat{D} + F_a\hat{A}\|^2 = \min_{\bar{U}\in\mathbb{R}^{n_u}} \|Z_0 + F_x\hat{x}_0 + F_u\bar{U} + F_d\hat{D} + F_a\hat{A}\|^2 \tag{10}$$

$$\|Z_0 + F_x\hat{x}_0 + F_u U + F_d\hat{D} + F_a\hat{A}\|^2 = \max_{\substack{\bar{x}_0\in\mathbb{R}^{n_x}, \bar{D}\in\mathbb{R}^{n_d},\\ \bar{A}\in\mathbb{R}^{n_a}}} \|Z_0 + F_x\bar{x}_0 + F_u U + F_d\bar{D} + F_a\bar{A}\|^2.$$

$$\text{subject to } \|\bar{D}\| \leqslant \Delta,$$

$$Y = G_x\bar{x}_0 + G_d\bar{D} + G_a\bar{A}, \tag{11}$$

and satisfy the constraints $\|\hat{D}\| \leqslant \Delta$, $Y = G_x\hat{x}_0 + G_d\hat{D}$. The control U is a minimum for the unconstrained convex optimization (10) if and only if it satisfies the 1st order optimality condition

$$2F_u'(Z_0 + F_x\hat{x}_0 + F_u U + F_d\hat{D} + F_a\hat{A}) = 0. \tag{12}$$

Defining

$$L_u(\hat{A}, \hat{D}, \hat{x}_0, \lambda_u, \nu_u) \coloneqq \|Z_0 + F_x\hat{x}_0 + F_u U + F_d\hat{D} + F_a\hat{A}\|^2 - \lambda_u(\|\hat{D}\|^2 - \Delta^2) - \nu_u'\big(Y - G_x\hat{x}_0 - G_d\hat{D} - G_a\hat{A}\big),$$

$$g_u(\hat{A}, \hat{D}, \hat{x}_0) \coloneqq \|\hat{D}\|^2 - \Delta^2,$$

$$h_u(\hat{A}, \hat{D}, \hat{x}_0) \coloneqq Y - G_x\hat{x}_0 - G_d\hat{D} - G_a\hat{A},$$

we conclude from [19][Theorem 2.21] that the triple $(\hat{A}, \hat{x}_0, \hat{D})$ is a strict local maximum to (11) if there exist $\lambda_u > 0$, $\nu_u \in \mathbb{R}^{n_y}$ such that

$$\frac{\partial L_u(\hat{A}, \hat{D}, \hat{x}_0, \lambda_u, \nu_u)}{\partial \hat{A}} = G_a'\nu_u + 2F_a'(Z_0 + F_x\hat{x}_0 + F_u U + F_d\hat{D} + F_a\hat{A}) = 0, \tag{13a}$$

$$\frac{\partial L_u(\hat{A}, \hat{D}, \hat{x}_0, \lambda_u, \nu_u)}{\partial \hat{D}} = -2\lambda_u\hat{D} + G_d'\nu_u + 2F_d'(Z_0 + F_x\hat{x}_0 + F_u U + F_d\hat{D} + F_a\hat{A}) = 0, \tag{13b}$$

$$\frac{\partial L_u(\hat{A}, \hat{D}, \hat{x}_0, \lambda_u, \nu_u)}{\partial \hat{x}_0} = G_x'\nu_u + 2F_x'(Z_0 + F_x\hat{x}_0 + F_u U + F_d\hat{D} + F_a\hat{A}) = 0, \tag{13c}$$

$$G_x\hat{x}_0 + G_d\hat{D} + G_a\hat{A} = Y, \tag{13d}$$

$$\|\hat{D}\|^2 = \Delta^2, \tag{13e}$$

$$V_u'\left(\begin{bmatrix} F_a' \\ F_d' \\ F_x' \end{bmatrix}\begin{bmatrix} F_a & F_d & F_x \end{bmatrix} - \lambda_u\begin{bmatrix} 0 & 0 & 0 \\ 0 & I_{n_d} & 0 \\ 0 & 0 & 0 \end{bmatrix}\right)V_u < 0. \tag{13f}$$

Equation (9a) guarantees that (12) and (13a)–(13d) hold, condition (9b) implies that (13f) holds, and condition (9c) guarantees that (13e) holds. ■

B. Nash equilibrium

Building upon the result from the section above, we now state a result that provides a sufficient condition for the existence of a Nash equilibrium.

Theorem 2 (Sufficient condition): Suppose that there exist vectors $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u, A, \tilde{D}, \tilde{x}_0, \nu_a$ and scalars $\lambda_u > 0, \lambda_a > 0$ such that

$$\begin{bmatrix}
2F_a'F_a & 2F_a'F_d & 2F_a'F_x & 2F_a'F_u & 0 & 0 & 0 & G_a' & 0 \\
2F_d'F_a & 2F_d'F_d - 2\lambda_u I_{n_d} & 2F_d'F_x & 2F_d'F_u & 0 & 0 & 0 & G_d' & 0 \\
2F_x'F_a & 2F_x'F_d & 2F_x'F_x & 2F_x'F_u & 0 & 0 & 0 & G_x' & 0 \\
2F_u'F_a & 2F_u'F_d & 2F_u'F_x & 2F_u'F_u & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 2F_a'F_u & 2F_a'F_a & 0 & 0 & 0 & -G_a' \\
0 & 0 & 0 & 0 & 0 & -2\lambda_a I_{n_d} & 0 & 0 & G_d' \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & G_x' \\
G_a & G_d & G_x & 0 & -G_a & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -G_a & G_d & G_x & 0 & 0
\end{bmatrix}
\begin{bmatrix} \hat{A} \\ \hat{D} \\ \hat{x}_0 \\ U \\ A \\ \tilde{D} \\ \tilde{x}_0 \\ \nu_u \\ \nu_a \end{bmatrix}
=
\begin{bmatrix} -2F_a'Z_0 \\ -2F_d'Z_0 \\ -2F_x'Z_0 \\ -2F_u'Z_0 \\ -2F_a'(Z_0 + F_x x_0 + F_d D) \\ 0 \\ 0 \\ G_x x_0 + G_d D \\ G_x x_0 + G_d D \end{bmatrix} \tag{14a}$$

and

$$\|\hat{D}\| = \|\tilde{D}\| = \Delta, \tag{14b}$$

$$V_u'\left(\begin{bmatrix} F_a' \\ F_d' \\ F_x' \end{bmatrix}\begin{bmatrix} F_a & F_d & F_x \end{bmatrix} - \lambda_u\begin{bmatrix} 0 & 0 & 0 \\ 0 & I_{n_d} & 0 \\ 0 & 0 & 0 \end{bmatrix}\right)V_u < 0, \tag{14c}$$

$$V_a'\left(\begin{bmatrix} F_a' \\ 0 \\ 0 \end{bmatrix}\begin{bmatrix} F_a & 0 & 0 \end{bmatrix} - \lambda_a\begin{bmatrix} 0 & 0 & 0 \\ 0 & I_{n_d} & 0 \\ 0 & 0 & 0 \end{bmatrix}\right)V_a < 0, \tag{14d}$$

where V_u and V_a are full column rank matrices and

$$\operatorname{Im} V_u = \ker\begin{bmatrix} 0 & 2\hat{D}' & 0 \\ G_a & G_d & G_x \end{bmatrix},$$

$$\operatorname{Im} V_a = \ker\begin{bmatrix} 0 & 2\tilde{D}' & 0 \\ G_a & -G_d & -G_x \end{bmatrix}.$$

Then, (A, U) is a strict local Nash equilibrium and U is a strict local saddle-point for P_U. □

Proof of Theorem 2. Defining

$$
\begin{aligned}
&L_a(A, \tilde{D}, \tilde{x}_0, \lambda_a, \nu_a) \coloneqq \\
&\|Z_0 + F_x x_0 + F_u U + F_d D + F_a A\|^2 - \lambda_a(\|\tilde{D}\|^2 - \Delta^2) \\
&- \nu_a'\big(G_x(x_0 - \tilde{x}_0) + G_d(D - \tilde{D}) + G_a A\big), \\
&g_a(A, \tilde{D}, \tilde{x}_0) \coloneqq \|\tilde{D}\|^2 - \Delta^2, \\
&h_a(A, \tilde{D}, \tilde{x}_0) \coloneqq G_x(x_0 - \tilde{x}_0) + G_d(D - \tilde{D}) + G_a A,
\end{aligned}
$$

we conclude from [19][Theorem 2.21] that the triple $(A, \tilde{D}, \tilde{x}_0)$ is a strict local maximum to (6) if there exist $\lambda_a > 0$, $\nu_a \in \mathbb{R}^{n_y}$ such that

$$
\begin{aligned}
&\frac{\partial L_a(A, \tilde{D}, \tilde{x}_0, \lambda_a, \nu_a)}{\partial A} = -G_a'\nu_a + \\
&2F_a'(Z_0 + F_x x_0 + F_u U + F_d D + F_a A) = 0 && \text{(15a)} \\
&\frac{\partial L_a(A, \tilde{D}, \tilde{x}_0, \lambda_a, \nu_a)}{\partial \tilde{D}} = -2\lambda_a \tilde{D} + G_d'\nu_a = 0 && \text{(15b)} \\
&\frac{\partial L_a(A, \tilde{D}, \tilde{x}_0, \lambda_a, \nu_a)}{\partial \tilde{x}_0} = G_x'\nu_a = 0 && \text{(15c)} \\
&G_x(x_0 - \tilde{x}_0) + G_d(D - \tilde{D}) + G_a A = 0 && \text{(15d)} \\
&\|\tilde{D}\|^2 = \Delta^2, && \text{(15e)} \\
&V_a'\left(\begin{bmatrix} F_a' \\ 0 \\ 0 \end{bmatrix}\begin{bmatrix} F_a & 0 & 0 \end{bmatrix} - \lambda_a \begin{bmatrix} 0 & 0 & 0 \\ 0 & I_{n_d} & 0 \\ 0 & 0 & 0 \end{bmatrix}\right) V_a < 0. && \text{(15f)}
\end{aligned}
$$

On the other hand, we saw in the proof of Theorem 1 that $(U, \hat{x}_0, \hat{D}, \hat{A})$ is a local strict saddle-point equilibrium for (7) if (9a)–(9c) hold. The theorem's assumptions guarantee that all these conditions are simultaneously satisfied, with Y in (9a) given by (5). ■

IV. Computation of Nash equilibrium

This section builds upon the existence results in Theorems 1 and 2 and the following lemma, the proof of which is presented in the technical report [18][Appendix].

Lemma 1: If $F_u'F_u$ is a nonsingular matrix and

$$\operatorname{Im} G_a \cap \operatorname{Im} G_x = \{0\}. \tag{16}$$

Then, there exists scalars $\lambda_u^* > 0$, $\lambda_a^* > 0$ such that for every $\lambda_u \geqslant \lambda_u^*$ and $\lambda_a \geqslant \lambda_a^*$,

1) the inequalities (9b), (14c) and (14d) hold;
2) the system of equations (9a) have a solution for every Z_0 and every Y consistent with (5); and
3) the system of equations (14a) have a solution for every Z_0, x_0, D. □

Condition (16) essentially means that the equation

$$G_a A = G_x x_0$$

can only hold for $A = 0$ and $x_0 = 0$ because otherwise there would be a nonzero vector simultaneously in the images of G_a and G_x. This means that we cannot have any (nonzero) attack A whose effect on Y can be precisely equal to the effect on Y of an initial condition x_0. We should thus view (16) as an observability condition that is needed to guarantee that the attacker cannot completely "hide" behind a legitimate initial condition.

A. Computation of P_{U}'s action

In practice, the player P_{U} cannot compute the Nash equilibrium because knowledge of x_0 and D is needed to solve (14a) and P_{U} does not have this information. However, in view of Theorem 1, P_{U} can compute its saddle-point U by solving (9) for the unknowns $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u, \lambda_u$, because these equations only depend on the measurement Y.

While the nonlinear system of equalities/inequalities in (9) is generally difficult to solve for all the unknowns $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u, \lambda_u$, in view of Lemma 1, for sufficiently large values of λ_u the inequality (9b) always holds and the linear system of equations (9a) always has a solution for $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u$. This provides a computationally simple procedure to find the saddle point:

1) For each $\lambda_u \in [\lambda_u^*, \infty)$ solve the linear system of equations (9a) for $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u$.
2) Any value of λ_u for which the solution leads to $\|\hat{D}\| = \Delta$ corresponds to a strict local saddle-point for P_{U}.

In practice, one does not really need to solve (9a) over the whole range of value $\lambda_u \in [\lambda_u^*, \infty)$, as it suffices to solve this equation over a fine grid $\{\lambda_u^1, \lambda_u^2, \dots, \lambda_u^K\}$ and then zone in on any interval $[\lambda_u^k, \lambda_u^{k+1}]$ where $\|\hat{D}\|$ goes from a value above Δ to a value below delta or vice versa. Within such a small interval, one can use either a bisection method of Newton's method to find the precise value of λ_u for which we have $\|\hat{D}\| = \Delta$.

The only limitation of this method is that if we are unable to find a value $\lambda_u \in [\lambda_u^*, \infty)$ for which $\|\hat{D}\| = \Delta$, we cannot be sure that a saddle-point does not exist for $\lambda_u < \lambda_u^*$. However, the method will always succeed for sufficiently small Δ because the solution to (9a) leads to $\hat{D} \to 0$ as we make $\lambda_u \to \infty$.

B. Computation of P_{A}'s action

For player P_{A} to compute its action A, it needs to compute the Nash equilibrium, which means solving the nonlinear system of equalities/inequalities in (14) for all the unknowns $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u, A, \tilde{D}, \tilde{x}_0, \nu_a, \lambda_u, \lambda_a$. However, also here Lemma 1 guarantees that for sufficiently large values of λ_u, λ_a the inequalities (14c)–(14d) always hold and the linear system of equations (14a) always has a solution for $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u, A, \tilde{D}, \tilde{x}_0, \nu_a$. This provides another computationally simple procedure to find saddle points:

1) For each pair $(\lambda_u, \lambda_a) \in [\lambda_u^*, \infty) \times [\lambda_a^*, \infty)$ solve the linear system of equations (9a) for $U, \hat{A}, \hat{D}, \hat{x}_0, \nu_u, A, \tilde{D}, \tilde{x}_0, \nu_a$.
2) Any pair (λ_u, λ_a) for which the solution leads to $\|\hat{D}\| = \|\tilde{D}\| = \Delta$ corresponds to a strict local Nash equilibrium and provide the action A for P_A.

This procedure is computationally more intensive than the one proposed to compute U because we now need to range over two scalar variables. However, gridding two scalar variables is still computationally quite simple and, once we find a small rectangle in $[\lambda_u^*, \infty) \times [\lambda_a^*, \infty)$, where the desired pair (λ_u, λ_a) is suspected to lie, we can use Newton's method to get $\|\hat{D}\| = \|\tilde{D}\| = \Delta$.

Denoting by $\hat{D}(\lambda_u, \lambda_a)$ and $\tilde{D}(\lambda_u, \lambda_a)$ the solutions to (14a) corresponding to a specific pair (λ_u, λ_a), our goal is to get

$$\begin{bmatrix} \hat{D}'(\lambda_u, \lambda_a)\hat{D}(\lambda_u, \lambda_a) - \Delta^2 \\ \tilde{D}'(\lambda_u, \lambda_a)\tilde{D}(\lambda_u, \lambda_a) - \Delta^2 \end{bmatrix} = 0$$

and Newton's method iteration to solve this equation is of the form

$$\begin{bmatrix} \lambda_u^{k+1} \\ \lambda_a^{k+1} \end{bmatrix} = \begin{bmatrix} \lambda_u^{k} \\ \lambda_a^{k} \end{bmatrix} - \alpha J^{-1} \begin{bmatrix} \hat{D}'(\lambda_u, \lambda_a)\hat{D}(\lambda_u, \lambda_a) - \Delta^2 \\ \tilde{D}'(\lambda_u, \lambda_a)\tilde{D}(\lambda_u, \lambda_a) - \Delta^2 \end{bmatrix},$$

where $\alpha \in [0, 1]$ is the step size, to be chosen as close to one as possible, while making sure that λ_u^k and λ_a^k remain in a desired region and

$$J := \begin{bmatrix} \frac{\partial \hat{D}'(\lambda_u,\lambda_a)\hat{D}(\lambda_u,\lambda_a)}{\partial \lambda_u} & \frac{\partial \hat{D}'(\lambda_u,\lambda_a)\hat{D}(\lambda_u,\lambda_a)}{\partial \lambda_a} \\ \frac{\partial \tilde{D}'(\lambda_u,\lambda_a)\tilde{D}(\lambda_u,\lambda_a)}{\partial \lambda_u} & \frac{\partial \tilde{D}'(\lambda_u,\lambda_a)\tilde{D}(\lambda_u,\lambda_a)}{\partial \lambda_a} \end{bmatrix}$$

The partial derivatives of $\hat{D}$ and $\tilde{D}$ with respect to λ_u, λ_a can be obtained from (14a) using the Implicit Function Theorem, which leads to a system of linear equations in the partial derivatives.

V. Example

For the position control of a double integrator with a quadratic penalty in the control, based on a noisy measurement vector $y(t)$ with position and velocity information and the velocity measurement subject to a deception attack, we have a model like (8) with

$$A_x := \begin{bmatrix} 1 & h \\ 0 & 1 \end{bmatrix}, \quad B := \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad C_y := \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$
$$C_z := \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad D_y := \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad D_z = \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

where h is the sampling interval and $A_x^k = \left[\begin{smallmatrix} 1 & kh \\ 0 & 1 \end{smallmatrix}\right]$. This leads to

$$G_x = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & h \\ 0 & 1 \\ 1 & 2h \\ 0 & 1 \\ \vdots & \\ 1 & Lh \\ 0 & 1 \end{bmatrix}, F_x = \begin{bmatrix} 1 & (L+1)h \\ 1 & (L+2)h \\ 1 & (L+3)h \\ \vdots & \\ 1 & (L+T)h \end{bmatrix},$$

$$G_a = \begin{bmatrix} 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 \\ 1 & 0 & 1 & 0 & \cdots & 0 \\ \sum_{\tau=0}^{1} \tau h & 0 & 0 & 0 & \cdots & 0 \\ 2 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots \\ \sum_{\tau=0}^{L-1} \tau h & 0 & 0 & 0 & \cdots & 0 \\ L & 0 & 0 & 0 & \cdots & 1 \end{bmatrix},$$

$$F_a = \begin{bmatrix} \sum_{\tau=0}^{L} \tau h & 0 & \cdots & 0 \\ \sum_{\tau=0}^{L+1} \tau h & 0 & \cdots & 0 \\ \sum_{\tau=0}^{L+2} \tau h & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ \sum_{\tau=0}^{L+T-1} \tau h & 0 & \cdots & 0 \end{bmatrix},$$

and it is straightforward to see that for any $L \geqslant 1$,

$$G_x x_0 = G_a A \quad \Rightarrow \quad x_0 = 0$$

and for $L \geqslant 2$ we further have that

$$G_x x_0 = G_a A \quad \Rightarrow \quad x_0 = 0, \ F_a A = 0,$$

which means that (16) holds. This would not be the case for $D_y := \left[\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right]$, which would correspond to a deception attack on the position measurement. This is not surprising because in this case the deception attack would compromise the only measurement for which we have observability of the state.

Figure 1 shows the results of the optimal control and attacks for different values of the upper bound Δ on the norm of D. As Δ increases, the attacker has more freedom to choose A without being detected and consequently the cost for P_U increases.However, the true cost $\|Z\|$ for P_U never exceeds its security value $J_u(Y)$. □

VI. Conclusion and Future Directions

This paper addressed the output-feedback robust control of a linear system subject to disturbances and noise

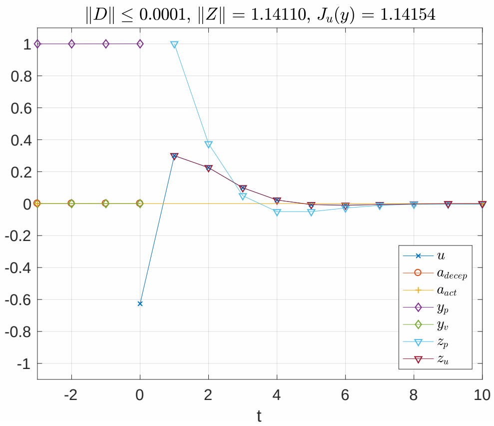

(a) $\|D\| \leqslant .0001$

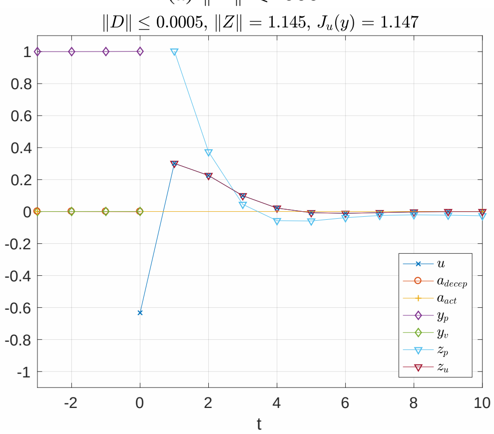

(b) $\|D\| \leqslant .0005$

Fig. 1: Signals for the double integrator example with $L = 3$, $T = 10$, $x(-L) = [1\ 0]'$, and two distinct values for the upper bound on the norm of D. In these plots we actually have $D = 0$, but the attacker manipulates Y to remain consistent with the upper bounds on D.

that cannot be measured in presence of an attacker that can corrupt the measured output (deception attack) and can introduce perturbations to the control signal (actuation attack). We formulated this problem as a zero-sum game between a defender that selects the control signal based on a measured output and an attacker that selects the attack signals. The formulation included asymmetric information in terms of what is known to the players. The main contributions were (i) sufficient conditions for the existence of a Nash equilibrium corresponding to a saddle-point for the defender and (ii) a computationally efficient procedure to compute a pair of policies that form a Nash equilibrium for the game. We applied this theory to finite horizon control of a linear system with quadratic cost.

Future directions include a receding horizon extension of this framework and extensions to more general information structures for the players.

References

[1] J. A. B. Bacet, "Inside the cunning, unprecedented hack of Ukraine's power grid," available online. [Online]. Available: http://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/

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

[3] J. Slay and M. Miller, "Lessons learned from the Maroochy water breach," *Critical infrastructure protection*, pp. 73–82, 2007.

[4] R. B. Bobba, K. M. Rogers, Q. Wang, H. Khurana, K. Nahrstedt, and T. J. Overbye, "Detecting false data injection attacks on DC state estimation," in *Preprints of the First Workshop on Secure Control Systems, CPSWEEK*, 2010.

[5] Y. Liu, M. K. Reiter, and P. Ning, "False data injection attacks against state estimation in electric power grids," in *ACM conference on Computer and communications security*, 2009.

[6] Y. Mo and B. Sinopoli, "False data injection attacks in control systems," in *First Workshop on Secure Control Systems, CPS Week*, 2010.

[7] A. Teixeira, I. Shames, H. Sandberg, and K. H. Johansson, "Revealing stealthy attacks in control systems," in *50th Annual Allerton Conf. on Communication, Control and Comp.*, 2012.

[8] S. D. Bopardikar and A. Speranzon, "On analysis and design of stealth-resilient control systems," in *6th International Symposium on Resilient Control Systems*, 2013, pp. 48–53.

[9] S. D. Bopardikar, A. Speranzon, and J. P. Hespanha, "An h-infinity approach to stealth-resilient control design," in *Resilience Week (RWS), 2016*. IEEE, 2016, pp. 56–61.

[10] F. Pasqualetti, F. Dorfler, and F. Bullo, "Attack detection and identification in cyber-physical systems," *IEEE Transaction on Automatic and Control*, vol. 58, no. 11, pp. 2715–2729, 2013.

[11] M. S. Chong, M. Wakaiki, and J. P. Hespanha, "Observability of linear systems under adversarial attacks," in *American Control Conference (ACC), 2015*. IEEE, 2015, pp. 2439–2444.

[12] S. Mishra, Y. Shoukry, N. Karamchandani, S. N. Diggavi, and P. Tabuada, "Secure state estimation against sensor attacks in the presence of noise," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 1, pp. 49–59, 2017.

[13] T. Başar and P. Bernhard, *H-infinity optimal control and related minimax design problems: a dynamic game approach*. Springer Science & Business Media, 2008.

[14] F. Farokhi, A. M. H. Teixeira, and C. Langbort, "Estimation with strategic sensors," *IEEE Transactions on Automatic Control*, vol. 62, no. 2, pp. 724–739, 2017.

[15] S. Sarıtaş, S. Yüksel, and S. Gezici, "Quadratic multi-dimensional signaling games and affine equilibria," *IEEE Transactions on Automatic Control*, vol. 62, no. 2, pp. 605–619, 2017.

[16] I. Yaesh and U. Shaked, "Game theory approach to state estimation of linear discrete-time processes and its relation to h-infinity optimal estimation," *International Journal of Control*, vol. 55, no. 6, pp. 1443–1452, 1992.

[17] D. Simon, "A game theory approach to constrained minimax state estimation," *IEEE Transactions on Signal Processing*, vol. 54, no. 2, pp. 405–412, 2006.

[18] J. P. Hespanha and S. D. Bopardikar, "Output-feedback linear quadratic robust control under actuation and deception attacks," *arXiv preprint arXiv:* **To be included**, 2018.

[19] R. H. W. Hope, "Optimization I: Chapter 2 –Theory of Constrained Optimization," available online. [Online]. Available: https://www.math.uh.edu/~rohop/fall_06/Chapter2.pdf