Automatica 50 (2014) 852–857

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Security concepts for the dynamics of autonomous vehicle networks[☆]

Mengran Xue[a], Wei Wang[b], Sandip Roy[b,1]

[a] University of Michigan, Ann Arbor, MI 48109, United States
[b] Washington State University, Pullman, WA 99164, United States

ARTICLE INFO

Article history:
Received 23 January 2012
Received in revised form
22 September 2013
Accepted 30 October 2013
Available online 1 February 2014

Keywords:
Network security
Autonomous vehicle teams
Network dynamics
Estimation
Control theory
Graph theory

ABSTRACT

The secure operation of autonomous vehicle networks in the presence of adversarial observation is examined, in the context of a canonical double-integrator-network (DIN) model. Specifically, we study the ability of a sentient adversary to estimate the full network's state, from noisy local measurements of vehicle motions. Algebraic, spectral, and graphical characterizations are provided, which indicate the critical role of the inter-vehicle communication topology and control scheme in achieving security.

1. Introduction

Threat assessment and mitigation challenges are increasingly arising in networks with tightly intertwined physical- and cyber-components. Dynamical-network estimation and control tools appear to be well-suited to address these challenges, in that they permit modeling and design of network dynamics involving both complex physical processes and information–communication/processing capabilities. In consequence, there has been a growing focus in the network-controls community on defining and analyzing security and vulnerability notions in a range of cyber–physical networks, e.g. Pasqualetti, Dorfler, and Bullo (2011), Xue, Roy, Wan, and Das (2011) and Zhu and Martinez (2011). In this article, motivated by autonomous-vehicle-network (AVN) applications, we examine one interesting cyber–physical threat-assessment or security-analysis problem. Specifically, we focus on the canonical double-integrator-network (DIN) model for an AVN engaged in a tracking task, which provides an abstract representation of both the network's physical dynamics and its cyber- (communication/control) capabilities. Here, we enrich the DIN to represent an adversary, which can obtain noisy measurements of some vehicles' local dynamics but is not able to actuate the dynamics in any way. We then define notions of security, which capture whether or not an adversary can discover or estimate the initial state of the DIN from its measurements, and also describe the fidelity of these estimates when observations are noisy.[2] Starting from an observability analysis of the closed-loop dynamics, several spectral and graphical characterizations of these security notions are obtained (for both the noise-free and noisy cases), that indicate the role of the network's communication and control architecture in protecting the network dynamics from discovery.

The security analysis pursued here is related to several current research thrusts in the network-controls and cyber-physical-systems literature. First, our work is closely tied to several fundamental studies of network dynamics and structure estimation from local observations (including via distributed estimation), e.g. Pasqualetti, Bicchi, and Bullo (2012), Pasqualetti et al. (2011), Roy, Xue, and Das (2012), Sou, Sandberg, and Johansson (2012), Sundaram and Hadjicostis (2011) and Wan and Roy (2009). We

[☆] The material in this paper was partially presented at the 51st IEEE Conference on Decision and Control (CDC), December 10–13, 2012, Maui, Hawaii. This paper was recommended for publication in revised form by Associate Editor Antonis Papachristodoulou under the direction of Editor Frank Allgöwer.

E-mail addresses: mxue@umich.edu (M. Xue), wwang1@eecs.wsu.edu (W. Wang), sroy@eecs.wsu.edu (S. Roy).

[1] Tel.: +1 509 335 2448; fax: +1 509 335 3818.

[2] The term "security" is also used to describe a system's ability to thwart an active attack. In accordance with our earlier work (Xue et al., 2011), we here use the term to describe the protection of information from a measuring adversary, but acknowledge the varying definitions. It is also worth stressing that the security notion we consider here is closely connected to the emerging concept of privacy or anonymity in distributed algorithms/controls, see e.g. Le Ny and Pappas (2012) and Telerius, Varagnolo, Baquero, and Johansson (2013). From this perspective, our results can be viewed as characterizing the level of privacy among the agents/vehicles in a DIN.

0005-1098/$ – see front matter © 2014 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.automatica.2013.12.001

particularly point out that several of these efforts are also concerned with local monitoring of network dynamics, whether by system planners to detect an attack or by adversaries to detect nominal network dynamics. As a dual to these graph-theoretic estimation/observability analyses, graph-theoretic viewpoints on controllability have also been developed that have a similar flavor (Rahmani, Ji, Mesbahi, & Egerstedt, 2010). Second, this work is complementary to control-theoretic modeling of attacks in cyber-networks and networked control systems (e.g., Alpcan & Basar, 2003, Amin, Cardenas, & Sastry, 2009, Liu, Ning, & Reiter, 2011, Mo & Sinopoli, 2009 and Texiera, Sandberg, & Johansson, 2010). While these efforts mostly are concerned with active adversaries, many of these works also obtain graphical results on adversarial conduct in networks, as in our work. We also note the connection of our work to recent efforts on fault and event detection in networks, including for systems with second-order local dynamics (Roy & Chen, 2013; Shames, Texiera, Sandberg, & Johansson, 2010). More broadly, this study is related to efforts to define security and vulnerability concepts for cyber–physical networks (e.g., Pasqualetti et al., 2011 and Xue et al., 2011); it also enriches the extensive study of AVN control in the control community (e.g., Ren & Beard, 2005 and Roy, Saberi, & Herlugson, 2004), toward performance design to achieve security. While our explorations here are connected to these research thrusts, we stress that a particular focus of our analyses is to distinguish the roles of the AVN's physical dynamics (motion), communication, and control, and of the adversary's (sparse) measurement capabilities, in network security.

In brief, the particular contributions described in this article are the following:

- We extend the DIN model to capture an adversary with local observation capabilities, and introduce attendant notions of security concerned with estimation of the full network's initial state.
- We characterize a binary notion of security (an observability notion) when the adversary's observations are not noisy, in terms of the network's graph matrix, its spectrum, and the graph topology itself.
- We characterize security levels (which codify estimation error) in the noisy-measurement case, in terms of the graph matrix and its spectrum.

2. Problem formulation

In this section, the DIN model is reviewed, the adversary is modeled, and security notions are formally defined.

The DIN. We consider a team of n vehicles, labeled as $i = 1, \dots, n$. For convenience, we assume that each vehicle is moving in a single dimension: the multi-dimensional case can be transformed to a single-dimensional model with more vehicles, see Roy et al. (2004). The vehicles' positions satisfy the differential equation $\ddot{\mathbf{x}}(t) = \mathbf{u}(t)$, where the *full position vector* $\mathbf{x}(t) = \begin{bmatrix} x_1(t), & \dots, & x_n(t) \end{bmatrix}^T$ contains the positions $x_i(t)$ of each vehicle, and the *full input vector* $\mathbf{u}(t) = \begin{bmatrix} u_1(t), & \dots, & u_n(t) \end{bmatrix}^T$ specifies the control input $u_i(t)$ for each vehicle. We call $h_i(t) \triangleq \dot{x}_i(t)$ the *velocity* of vehicle i, and call $\mathbf{h}(t) = \begin{bmatrix} h_1(t), & \dots, & h_n(t) \end{bmatrix}^T$ the *full velocity vector*.

Each vehicle uses information that is sensed and/or communicated from other vehicles, as well as (possibly) information about a target location, to set its control inputs. Formally, each vehicle i is assumed to have available a p_i-component *position-observation vector* $\mathbf{y}_{pi}(t) = G_i\mathbf{x}(t)$, where the $p_i \times n$-dimensional matrix G_i specifies agent i's capability to observe the vehicle network's full state and hence is called agent i's *observation matrix* (see Roy et al., 2004 for many illustrative examples). The vehicle is also assumed to have commensurate velocity observations—specifically, that p_i-component *velocity-observation vector* $\mathbf{y}_{vi}(t) = G_i\mathbf{h}(t)$—either obtained through direct measurement, or as the derivative of the position measurement.[3] The vehicles' internal dynamics together with their observation capabilities nominally specify the DIN. For notational convenience, we also stack the observation matrices into a single matrix: $G \triangleq \begin{bmatrix} G_1^T, & \dots, & G_n^T \end{bmatrix}^T$, which we call the *full observation matrix*. We assemble $\mathbf{y}_{pi}$ and $\mathbf{y}_{vi}$ as: $\mathbf{y}_p \triangleq \begin{bmatrix} \mathbf{y}_{p1}^T, & \dots, & \mathbf{y}_{pn}^T \end{bmatrix}^T$, and $\mathbf{y}_v \triangleq \begin{bmatrix} \mathbf{y}_{v1}^T, & \dots, & \mathbf{y}_{vn}^T \end{bmatrix}^T$.

Tracking control in the DIN. Decentralized controller designs have been obtained for the DIN, for several coordinated-motion tasks. Here, we focus on a fixed-target-tracking problem. Specifically, the vehicles in the team are all tasked with moving from initial locations or *home bases* $x_i(0)$ to a specified scalar *target location* $\bar{s}$. It is assumed that the target location is distributed to the vehicles as needed prior to the tracking task. We stress here that the vehicles may not have measurements of absolute position information in the reference frame of the target, and hence they depend on sensing or communication to complete target tracking (Roy et al., 2004). It turns out that vehicles whose measurements are all relative positions/velocities (i.e., each row of G_i sums to 0) do not require knowledge of the target (Roy et al., 2004).

Here, we consider using a memoryless linear decentralized controller architecture to achieve tracking. Specifically, each vehicle i is assumed to use a controller of the form $u_i(t) = K_i\mathbf{y}_{vi}(t) + \alpha K_i(\mathbf{y}_{pi}(t) - G_i\mathbf{s}^*)$, where the 1-by-$p_i$ *gain matrix* K_i weights the observations in computing the concurrent input, α is a scalar gain factor that is common for all vehicles, and $\mathbf{s}^* \triangleq \begin{bmatrix} \bar{s}, & \dots, & \bar{s} \end{bmatrix}^T$. It is convenient to assemble the gain matrices as the diagonal blocks of a *full gain matrix*: $K = \operatorname{diag}(K_i)$. In this notation, the controllers of all n vehicles are captured by the equation $\mathbf{u}(t) = K\mathbf{y}_v(t) + \alpha K(\mathbf{y}_p(t) - G\mathbf{s}^*)$, or equivalently $\mathbf{u}(t) = KG\mathbf{h}(t) + \alpha K(G\mathbf{x}(t) - G\mathbf{s}^*)$. A couple notes about the control architecture are worthwhile. First, we note that $G_i\mathbf{s}^* = 0$ for any vehicle i whose observations are relative positions/velocities, which verifies that these vehicles do not require information about the tracking task. Also, we stress that at least one vehicle must make measurements in an absolute frame (and hence must have information about the target location) for tracking to be possible: otherwise, G will not have full rank and stabilization is impossible.

Controllers of this architecture have been designed to achieve tracking, for a broad class of network topologies G, see Roy et al. (2004). Specifically, under broad conditions, the full gain matrix K can be designed so that KG has real, negative, and distinct eigenvalues. In turn, by choosing α to be small (specifically, less than 1/4 of the minimum eigenvalue of KG), the tracking task is achieved. Such memoryless low-gain controllers can be designed to be robust to actuator saturation, communication delay, and topological variation (Locatelli & Schiavoni, 2012; Roy et al., 2004). Beyond this class of special low-gain controllers, an even broader family of designs has been obtained to place the eigenvalues of KG in the *open left half plane* (OLHP), and in turn to achieve tracking. Many of our results here encompass any tracking controller of this type, while others are focused on the special low-gain control.

Closed-loop dynamics: matrix representation. Let us define an *extended state vector* of dimension $2n + 1$ that stacks the positions and velocities of the vehicles as well as the target location: $\boldsymbol{\varphi}(t) = \begin{bmatrix} \mathbf{x}^T(t) & \mathbf{h}^T(t) & s(t) \end{bmatrix}^T$, where $s(t)$ is a scalar that is fixed at the target location (i.e., $s(t) = \bar{s}$ for all t). This state vector $\boldsymbol{\varphi}(t)$, in the closed loop, evolves according to:

$$\dot{\boldsymbol{\varphi}} = A\boldsymbol{\varphi}, \tag{1}$$

[3] Derivative computations magnify high-frequency noise, so must be low-pass filtered. Implementations of derivatives used in feedback have been widely studied, including for the DIN, so we omit the details and assume that an accurate derivative observation is available.