49th IEEE Conference on Decision and Control
December 15-17, 2010
Hilton Atlanta Hotel, Atlanta, GA, USA

The Wireless Control Network: Monitoring for Malicious Behavior

Shreyas Sundaram, Miroslav Pajic, Christoforos N. Hadjicostis, Rahul Mangharam, and George J. Pappas

Abstract— **We consider the problem of stabilizing a plant with a network of resource constrained wireless nodes. In a companion paper, we developed a protocol where each node repeatedly transmits an appropriate (stabilizing) linear combination of the values in its neighborhood. In this paper, we design an Intrusion Detection System (IDS) for this control scheme, which observes the transmissions of certain nodes and uses that information to (a) recover the plant outputs (for data-logging and diagnostic purposes) and (b) identify malicious behavior by any of the wireless nodes in the network. We show that if the connectivity of the network is sufficiently high, the IDS only needs to observe a subset of the nodes in the network in order to achieve this objective. Our approach provides a characterization of the set of nodes that should be observed, a systematic procedure for the IDS to use to identify the malicious nodes and recover the outputs of the plant, and an upper bound on the delay required to obtain the necessary information.**

I. Introduction

The advent of low-cost and reliable wireless networks holds great promise for large, spatially distributed industrial control systems. The topic of control over networks (wireless or otherwise) has been intensively studied by researchers over the past decade, leading to design procedures for controllers that are tolerant to network imperfections such as packet dropouts and transmission delays [1], [2], [3]. In the companion paper [4], we introduced the *Wireless Control Network* (WCN), where the network *itself* acts as the controller (instead of having a specially designated node performing this task). Specifically, we considered a wireless network consisting of simple nodes that are able to exchange information only with their direct neighbors. We devised a protocol where each node transmits, at each time-step, a single value that is a linear combination of the values in its neighborhood. This scheme causes the wireless network to behave as a linear system with sparsity constraints on the system matrices (corresponding to the topology of the network). We provided a numerical design procedure (based on linear matrix inequalities) to determine the appropriate linear combinations for each node to use in order to stabilize the plant. As discussed in [4], this scheme has several benefits, including simple scheduling, low computational requirements, and the ability to handle geographically dispersed sensing and actuation points.

Recently, the need for a rigorous theory of *security* in industrial control systems has started to gain attention [5], [6], [7], [8]. In domains such as chemical process industries, aviation and critical infrastructure, attacks on the control systems could have disastrous consequences. The report [9] makes several key recommendations for designing secure control systems, including the need to maintain accurate logs of plant behavior, and to analyze this information to quickly detect and isolate anomalies. In traditional (data) networks, this is performed with an *Intrusion Detection System* (IDS), which raises an alarm if the observed traffic flow in the network deviates from expected patterns [10]. The paper [11] suggests an IDS for wireless networks in process control industries, capturing (at a policy level) attacks such as jamming, denial of service attacks, and corruptions in the formatting of data transmitted by certain nodes.

A more dangerous (and difficult to detect) attack in control networks is that of data *modification*, where malicious nodes subtly change the contents of messages that they are passing through the network, but otherwise follow the normal rules of transmission. In this paper, we describe how to design an IDS to detect data modification attacks in the control scheme proposed in [4]. The IDS will be responsible for observing the transmissions of certain nodes in the network in order to (a) recover the outputs of the plant (e.g., for data-logging purposes), and (b) identify data modification attacks; the overall architecture of the plant, control network and IDS is shown in Fig. 1. We show that the wireless control scheme from [4] allows malicious behavior to be identified by examining the transmissions of only a *subset* of the nodes in the network, provided that the network topology satisfies certain conditions. We provide an explicit characterization of the subset of nodes that needs to be monitored, along with a procedure for the IDS to follow in order to extract the required information from the transmissions of these nodes.

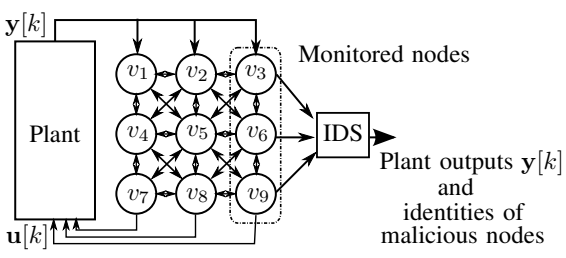

Fig. 1. Architecture of the wireless control network with an IDS.

This research has been supported by NSF CNS-0931239, DARPA MuSyC, and an NSERC Discovery Grant.

S. Sundaram is with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, N2L 3G1. Email: ssundara@uwaterloo.ca. M. Pajic, R. Mangharam and G. J. Pappas are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA 19014. Email: {pajic, rahulm, pappasg}@seas.upenn.edu. C. N. Hadjicostis is with the Department of Electrical and Computer Engineering, University of Cyprus. Email: chadjic@ucy.ac.cy.

II. Notation and Background on Graph Theory

We use $\mathbf{e}_i$ to denote the column vector (of appropriate size) with a 1 in its i-th position and 0's elsewhere, and $\mathbf{1}$ to denote the column vector (of appropriate size) consisting of all 1's. The symbol $\mathbf{I}_N$ denotes the $N \times N$ identity matrix, and $\mathbf{A}'$ indicates the transpose of matrix $\mathbf{A}$. The cardinality of a set $\mathcal{S}$ is denoted by $|\mathcal{S}|$, and for two sets $\mathcal{S}$ and $\mathcal{R}$, we use $\mathcal{S} \setminus \mathcal{R}$ to denote the set of elements in $\mathcal{S}$ that are not in $\mathcal{R}$. The set of nonnegative integers is denoted by $\mathbb{N}$.

A graph is an ordered pair $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where $\mathcal{V} = \{v_1, v_2, \ldots, v_N\}$ is a set of vertices (or nodes), and $\mathcal{E}$ is a set of ordered pairs of different vertices, called directed edges. The vertices in the set $\mathcal{N}_{v_i} = \{v_j | (v_j, v_i) \in \mathcal{E}\}$ are the neighbors of vertex v_i. A *subgraph* of $\mathcal{G}$ is a graph $\mathcal{H} = \{\bar{\mathcal{V}}, \bar{\mathcal{E}}\}$, with $\bar{\mathcal{V}} \subseteq \mathcal{V}$ and $\bar{\mathcal{E}} \subseteq \mathcal{E}$ (where all edges in $\bar{\mathcal{E}}$ are between vertices in $\bar{\mathcal{V}}$).

A *path* P from vertex v_{i_0} to vertex v_{i_t} is a sequence of vertices $v_{i_0} v_{i_1} \cdots v_{i_t}$ such that $(v_{i_j}, v_{i_{j+1}}) \in \mathcal{E}$ for $0 \leq j \leq t-1$. The nonnegative integer t is the *length* of the path. We will call a graph *disconnected* if there exists at least one pair of vertices $v_i, v_j \in \mathcal{V}$ such that there is no path from v_j to v_i. The *connectivity* of the graph is the smallest number of vertices that must be removed to disconnect the graph, and is denoted by κ. A set of paths $P_1, P_2, \ldots, P_r$ are vertex disjoint if no vertex appears in more than one path. Given two subsets $\mathcal{V}_1, \mathcal{V}_2 \subset \mathcal{V}$, an *r-linking* from $\mathcal{V}_1$ to $\mathcal{V}_2$ is a set of r vertex disjoint paths, each with start vertex in $\mathcal{V}_1$ and end vertex in $\mathcal{V}_2$. Note that if $\mathcal{V}_1$ and $\mathcal{V}_2$ are not disjoint, we will take their common vertices to be vertex disjoint paths between $\mathcal{V}_1$ and $\mathcal{V}_2$ of length zero.

Lemma 1 ([12]): Let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ have connectivity κ, and let $\mathcal{V}_1$ and $\mathcal{V}_2$ be subsets of $\mathcal{V}$, each of size at least κ. Then there is a κ-linking from $\mathcal{V}_1$ to $\mathcal{V}_2$ (and vice versa).

III. The Wireless Control Network

Consider a plant of the form:

$$\mathbf{x}[k+1] = \mathbf{A}\mathbf{x}[k] + \mathbf{B}\mathbf{u}[k], \quad \mathbf{y}[k] = \mathbf{C}\mathbf{x}[k], \qquad (1)$$

with $\mathbf{A} \in \mathbb{R}^{n \times n}, \mathbf{B} \in \mathbb{R}^{n \times m}$ and $\mathbf{C} \in \mathbb{R}^{p \times n}$. The output vector $\mathbf{y}[k] = \begin{bmatrix} y_1[k] & y_2[k] & \ldots & y_p[k] \end{bmatrix}'$ contains measurements of the plant state vector $\mathbf{x}[k]$ provided by the sensors $s_1, \ldots, s_p$. The input vector $\mathbf{u}[k] = \begin{bmatrix} u_1[k] & u_2[k] & \ldots & u_m[k] \end{bmatrix}'$ corresponds to the signals applied to the plant by actuators $a_1, \ldots, a_m$.

The plant is to be controlled using a wireless network consisting of a set of nodes that interact with each other and with the sensors and actuators installed on the plant. The network is described by a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where $\mathcal{V} = \{v_1, v_2, \ldots, v_N\}$ is the set of N nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ represents the radio connectivity (communication topology) in the network (i.e., edge $(v_j, v_i) \in \mathcal{E}$ if node v_i can receive information directly from node v_j). In addition, we define $\mathcal{V}_S \subset \mathcal{V}$ as the set of nodes that can receive information directly from at least one sensor, and $\mathcal{V}_A \subset \mathcal{V}$ as the set of nodes whose transmissions can be heard by at least one actuator. We will refer to $\mathcal{V}_S$ as the *source* nodes in the network. We will also assume that there are some *malicious nodes* in the network, given by the set $\mathcal{F} \subset \mathcal{V}$. These malicious nodes will transmit false values (perhaps by conspiring with each other) in an attempt to damage the system in some way. Note that the set $\mathcal{F}$ is unknown *a priori*.

We will find it convenient to consider a new graph $\bar{\mathcal{G}}$ that captures how the plant outputs enter into the wireless control network. This graph is obtained by taking the network graph $\mathcal{G}$ and adding p new vertices $\mathcal{S} = \{s_1, s_2, \ldots, s_p\}$, which correspond to the sensors on the plant. Define the edge set

$$\mathcal{E}_I = \left\{ (s_l, v_j) \;\middle|\; \begin{array}{c} s_l \in \mathcal{S}, v_j \in \mathcal{V}_S, \\ s_l\text{'s value is available to node } v_j \end{array} \right\}.$$

We then obtain $\bar{\mathcal{G}} = \{\mathcal{V} \cup \mathcal{S}, \mathcal{E} \cup \mathcal{E}_I\}$.

The proposed control scheme (introduced in [4]) consists of having each node in the network update its value to be a linear combination of its previous value and the values of its neighbors. In addition, each source node will include a linear combination of the sensor measurements (i.e., plant outputs) that it receives at each time-step. Finally, the malicious nodes will update their values arbitrarily at each time-step. Mathematically, if we let $z_i[k]$ denote node v_i's value at time-step k, we obtain the update equations:[1]

$$z_i[k+1] = \qquad (2)$$
$$\begin{cases} w_{ii} z_i[k] + \sum_{v_j \in \mathcal{N}_{v_i}} w_{ij} z_j[k] & \\ \quad + \sum_{s_j \in \mathcal{N}_{v_i}} h_{ij} y_j[k] & \text{if } v_i \in \mathcal{V}_S \setminus \mathcal{F}, \\ w_{ii} z_i[k] + \sum_{v_j \in \mathcal{N}_{v_i}} w_{ij} z_j[k] & \\ \quad + \sum_{s_j \in \mathcal{N}_{v_i}} h_{ij} y_j[k] + f_i[k] & \text{if } v_i \in \mathcal{V}_S \cap \mathcal{F}, \\ w_{ii} z_i[k] + \sum_{v_j \in \mathcal{N}_{v_i}} w_{ij} z_j[k] + f_i[k] & \text{if } v_i \in \mathcal{F} \setminus \mathcal{V}_S, \\ w_{ii} z_i[k] + \sum_{v_j \in \mathcal{N}_{v_i}} w_{ij} z_j[k] & \text{if } v_i \notin \mathcal{V}_S \cup \mathcal{F}. \end{cases}$$

The scalars w_{ij} and h_{ij} specify the linear combinations that are computed by each node in the network. The scalar $f_i[k]$ is an (arbitrary) additive error committed by node v_i at time-step k if it is malicious. If we let $\mathcal{F} = \{v_{j_1}, v_{j_2}, \ldots, v_{j_{|\mathcal{F}|}}\}$ denote the set of malicious nodes, and aggregate the values transmitted by all nodes at time-step k into the value vector $\mathbf{z}[k] = \begin{bmatrix} z_1[k] & z_2[k] & \cdots & z_N[k] \end{bmatrix}'$, the transmission strategy for the entire system can be represented as

$$\mathbf{z}[k+1] = \mathbf{W}\mathbf{z}[k] + \underbrace{\begin{bmatrix} \mathbf{H} & \mathbf{E}_{\mathcal{F}} \end{bmatrix}}_{\mathbf{B}_{\mathcal{F}}} \underbrace{\begin{bmatrix} \mathbf{y}[k] \\ \mathbf{f}[k] \end{bmatrix}}_{\mathbf{v}[k]}, \qquad (3)$$

for all $k \in \mathbb{N}$. In the above equation, $\mathbf{E}_{\mathcal{F}} = \begin{bmatrix} \mathbf{e}_{j_1} & \mathbf{e}_{j_2} & \cdots & \mathbf{e}_{j_{|\mathcal{F}|}} \end{bmatrix}$, and the vector $\mathbf{f}[k]$ contains all of the additive errors injected by the malicious nodes at time-step k. Furthermore, the (i,j) entry of $\mathbf{W}$ satisfies $w_{ij} = 0$ if $v_j \notin \mathcal{N}_{v_i} \cup \{v_i\}$, and the (i,j) entry of $\mathbf{H}$ satisfies $h_{ij} = 0$ if $s_j \notin \mathcal{N}_{v_i}$. We assume that $\mathbf{z}[0]$ (i.e., the initial state of the WCN) is known to the IDS.

At each actuator $l \in \{1, 2, \ldots, m\}$, we apply the input $u_l[k] = \mathbf{g}_l \mathbf{z}[k]$, where $\mathbf{g}_l$ is a vector that specifies a linear combination of the values transmitted by the nodes in $\mathcal{V}_A$

[1] The neighborhood $\mathcal{N}_v$ of a vertex v is with respect to the graph $\bar{\mathcal{G}}$.

that can be heard by that actuator. Thus, the entire input applied to the system can be written as $\mathbf{u}[k] = \mathbf{G}\mathbf{z}[k]$, where the sparsity pattern in $\mathbf{G}$ adheres to the network topology. When there are no malicious nodes (i.e., $\mathcal{F} = \emptyset$), the overall closed loop system evolves as:

$$\begin{bmatrix}\mathbf{x}[k+1]\\ \mathbf{z}[k+1]\end{bmatrix} = \begin{bmatrix}\mathbf{A} & \mathbf{B}\mathbf{G}\\ \mathbf{H}\mathbf{C} & \mathbf{W}\end{bmatrix}\begin{bmatrix}\mathbf{x}[k]\\ \mathbf{z}[k]\end{bmatrix} \triangleq \hat{\mathbf{A}}\begin{bmatrix}\mathbf{x}[k]\\ \mathbf{z}[k]\end{bmatrix}.$$

Let Ψ_s denote the set of all tuples $(\mathbf{W}, \mathbf{H}, \mathbf{G})$ that satisfy the required sparsity patterns and that cause the matrix $\hat{\mathbf{A}}$ to be stable. In [4], a numerical procedure was provided to find an element of Ψ_s (if one exists).

In this paper, we consider the problem of data collection and analysis in this network for the purpose of identifying malicious behavior by a nonempty subset $\mathcal{F}$ of nodes. Specifically, we will describe the design of an IDS whose task is to collect data from the network in order to (a) recover the plant outputs[2] $\mathbf{y}[k]$ and (b) detect and isolate anomalous behavior in the WCN. Clearly, one trivial option would be for the IDS to listen to the transmissions of *every* node and sensor in the network, and double-check that all nodes are computing the proper linear combinations at each time-step. However, this is not a satisfactory solution, since the entire point of the WCN is to avoid the communication infrastructure required for a centralized solution of this kind. Instead, we will show that it is possible to identify the malicious nodes and obtain the plant outputs by monitoring the transmissions of just a *subset* $\mathcal{T} \subset \mathcal{V}$ of the nodes (provided that the network topology satisfies certain conditions).

IV. Analysis Algorithm for the Intrusion Detection System

For any set $\mathcal{T} \subset \mathcal{V}$, denote the vector of transmissions of the nodes in that set at time-step k by $\mathbf{t}[k]$. We can write

$$\mathbf{t}[k] = \mathbf{T}\mathbf{z}[k] \ , \tag{4}$$

where $\mathbf{T}$ is a $|\mathcal{T}| \times N$ matrix with a single 1 in each row capturing the positions of the vector $\mathbf{z}[k]$ that are in the set $\mathcal{T}$, and zeros elsewhere. We will also find it useful to consider a slightly more general version of the system model (3). For any subset $\mathcal{Q} = \{v_{q_1}, v_{q_2}, \ldots, v_{q_{|\mathcal{Q}|}}\} \subset \mathcal{V}$ of nodes, let $\mathbf{E}_\mathcal{Q} = \begin{bmatrix}\mathbf{e}_{q_1} & \mathbf{e}_{q_2} & \cdots & \mathbf{e}_{q_{|\mathcal{Q}|}}\end{bmatrix}$, and define $\mathbf{B}_\mathcal{Q} = \begin{bmatrix}\mathbf{H} & \mathbf{E}_\mathcal{Q}\end{bmatrix}$ (where $\mathbf{H}$ is the matrix from (3) specifying the linear combinations of the plant outputs that are used by the source nodes). Note that $\mathbf{B}_\mathcal{Q}$ has $p + |\mathcal{Q}|$ columns. The values transmitted by the monitored nodes $\mathcal{T}$ over $L+1$ time-steps (for some nonnegative integer L) for the system

$$\mathbf{z}[k+1] = \mathbf{W}\mathbf{z}[k] + \mathbf{B}_\mathcal{Q}\mathbf{v}[k], \quad \mathbf{t}[k] = \mathbf{T}\mathbf{z}[k] \tag{5}$$

are given by

$$\underbrace{\begin{bmatrix}\mathbf{t}[k]\\ \mathbf{t}[k+1]\\ \mathbf{t}[k+2]\\ \vdots\\ \mathbf{t}[k+L]\end{bmatrix}}_{\mathbf{t}[k:k+L]} = \underbrace{\begin{bmatrix}\mathbf{T}\\ \mathbf{T}\mathbf{W}\\ \mathbf{T}\mathbf{W}^2\\ \vdots\\ \mathbf{T}\mathbf{W}^L\end{bmatrix}}_{\mathbf{\Theta}_L}\mathbf{z}[k] + \mathbf{M}_L^\mathcal{Q}\underbrace{\begin{bmatrix}\mathbf{v}[k]\\ \mathbf{v}[k+1]\\ \mathbf{v}[k+2]\\ \vdots\\ \mathbf{v}[k+L-1]\end{bmatrix}}_{\mathbf{v}[k:k+L-1]}, \tag{6}$$

[2]This information can be used for plant monitoring and data logging.

where

$$\mathbf{M}_L^\mathcal{Q} \triangleq \begin{bmatrix}\mathbf{0} & \mathbf{0} & \cdots & \mathbf{0}\\ \mathbf{T}\mathbf{B}_\mathcal{Q} & \mathbf{0} & \cdots & \mathbf{0}\\ \mathbf{T}\mathbf{W}\mathbf{B}_\mathcal{Q} & \mathbf{T}\mathbf{B}_\mathcal{Q} & \cdots & \mathbf{0}\\ \vdots & \vdots & \ddots & \vdots\\ \mathbf{T}\mathbf{W}^{L-1}\mathbf{B}_\mathcal{Q} & \mathbf{T}\mathbf{W}^{L-2}\mathbf{B}_\mathcal{Q} & \cdots & \mathbf{T}\mathbf{B}_\mathcal{Q}\end{bmatrix}. \tag{7}$$

The following theorem shows that the IDS can recover the desired quantities from the transmissions of nodes in $\mathcal{T}$, provided that a certain algebraic condition holds. We will later relate this algebraic condition to conditions on the network topology and choices of the monitored nodes $\mathcal{T}$.

Theorem 1: Suppose that there exists an integer D such that, for all possible sets $\mathcal{Q}$ of $2f$ nodes, the matrix $\mathbf{M}_D^\mathcal{Q}$ satisfies

$$\text{rank}\left(\mathbf{M}_D^\mathcal{Q}\right) = p + |\mathcal{Q}| + \text{rank}\left(\mathbf{M}_{D-1}^\mathcal{Q}\right) \ . \tag{8}$$

Then, as long as there are no more than f malicious nodes in the network during any set of D contiguous time-steps, the IDS can uniquely recover the plant outputs $\mathbf{y}[k]$ and identify all of the malicious nodes with a delay of D time-steps, based on the transmissions of the nodes in $\mathcal{T}$. □

Before proceeding with the proof of the above theorem, we provide a more detailed explanation of condition (8). Specifically, note from (7) that for any set $\mathcal{Q}$, the last $(L-1)$ block-columns of $\mathbf{M}_L^\mathcal{Q}$ have the form $\begin{bmatrix}\mathbf{0}\\ \mathbf{M}_{L-1}^\mathcal{Q}\end{bmatrix}$, and thus have rank equal to the rank of $\mathbf{M}_{L-1}^\mathcal{Q}$. Condition (8) is therefore equivalent to saying that the first $p + |\mathcal{Q}|$ columns of $\mathbf{M}_D^\mathcal{Q}$ must be linearly independent of each other, and of all other columns in $\mathbf{M}_D^\mathcal{Q}$. With this interpretation in hand, we are now ready to continue with the proof of Theorem 1.

Proof: [Theorem 1] Consider time-steps $k = 0, 1, \ldots, D$, and suppose that the malicious nodes during this period are a subset of the set $\mathcal{F} = \{v_{j_1}, v_{j_2}, \ldots, v_{j_f}\}$. From (3), (4) and (6), the values seen by the IDS over these time-steps are given by

$$\mathbf{t}[0:D] = \mathbf{\Theta}_D\mathbf{z}[0] + \mathbf{M}_D^\mathcal{F}\mathbf{v}[0:D-1] \ , \tag{9}$$

where $\mathbf{v}[k] = \begin{bmatrix}\mathbf{y}'[k] & \mathbf{f}'[k]\end{bmatrix}'$. Note that the IDS knows the quantities $\mathbf{t}[0:D]$ and $\mathbf{\Theta}_D\mathbf{z}[0]$, but it does not know the set $\mathcal{F}$ or the values $\mathbf{v}[0:D-1]$. The IDS will try to identify these unknown parameters based on the known quantities.

Let $\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_{\binom{N}{f}} \subset \mathcal{V}$ denote all possible sets of f nodes, and let $\mathbf{M}_D^{\mathcal{F}_1}, \mathbf{M}_D^{\mathcal{F}_2}, \ldots, \mathbf{M}_D^{\mathcal{F}_{\binom{N}{f}}}$ denote the input matrices corresponding to these sets. With these matrices in hand, suppose that the IDS finds the first $j \in \{1, 2, \ldots, \binom{N}{f}\}$ such that the vector $\mathbf{t}[0:D] - \mathbf{\Theta}_D\mathbf{z}[0]$ is in the column space of the matrix $\mathbf{M}_D^{\mathcal{F}_j}$. This means that the IDS can find a vector $\bar{\mathbf{v}}[0:D-1]$ such that

$$\mathbf{M}_D^{\mathcal{F}_j}\bar{\mathbf{v}}[0:D-1] = \mathbf{t}[0:D] - \mathbf{\Theta}_D\mathbf{z}[0].$$

The vector $\bar{\mathbf{v}}[0:D-1]$ is the IDS's *estimate* of the value of $\mathbf{v}[0:D-1]$ (note that the value $\bar{\mathbf{v}}[k] = \begin{bmatrix}\bar{\mathbf{y}}'[k] & \bar{\mathbf{f}}'[k]\end{bmatrix}'$ contains estimates of the plant outputs and the malicious errors

at time-step k). Substituting (9) into the above expression and rearranging, we have

$$\mathbf{M}_D^{\mathcal{F}}\mathbf{v}[0:D-1] - \mathbf{M}_D^{\mathcal{F}_j}\bar{\mathbf{v}}[0:D-1] = \mathbf{0} \ .$$

Let $\{\mathcal{F}, \mathcal{F}_j\}$ denote the set that is obtained by concatenating sets $\mathcal{F}$ and $\mathcal{F}_j$ (i.e., it is the union of the two sets, with duplications allowed). Exploiting the form of matrix $\mathbf{M}_D^{\mathcal{Q}}$ shown in (7), the above expression can be written as

$$\underbrace{\begin{bmatrix} \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{TB}_{\{\mathcal{F},\mathcal{F}_j\}} & \cdots & \mathbf{0} \\ \mathbf{TWB}_{\{\mathcal{F},\mathcal{F}_j\}} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{TW}^{D-1}\mathbf{B}_{\{\mathcal{F},\mathcal{F}_j\}} & \cdots & \mathbf{TB}_{\{\mathcal{F},\mathcal{F}_j\}} \end{bmatrix}}_{\mathbf{M}_D^{\{\mathcal{F},\mathcal{F}_j\}}} \begin{bmatrix} \tilde{\mathbf{v}}[0] \\ \tilde{\mathbf{v}}[1] \\ \tilde{\mathbf{v}}[2] \\ \vdots \\ \tilde{\mathbf{v}}[D-1] \end{bmatrix} = \mathbf{0} \tag{10}$$

where $\mathbf{B}_{\{\mathcal{F},\mathcal{F}_j\}} = \begin{bmatrix} \mathbf{H} & \mathbf{E}_{\mathcal{F}} & \mathbf{E}_{\mathcal{F}_j} \end{bmatrix}$ and

$$\tilde{\mathbf{v}}[k] = \begin{bmatrix} \mathbf{y}[k] - \bar{\mathbf{y}}[k] \\ \mathbf{f}[k] \\ -\bar{\mathbf{f}}[k] \end{bmatrix} \ .$$

Now consider the matrix $\mathbf{M}_D^{\mathcal{F}\cup\mathcal{F}_j}$. Since $\mathcal{F}\cup\mathcal{F}_j$ has at most $2f$ nodes, equation (8) indicates that the first $p + |\mathcal{F}\cup\mathcal{F}_j|$ columns of the matrix $\mathbf{M}_D^{\mathcal{F}\cup\mathcal{F}_j}$ are linearly independent of each other, and of all other columns of the matrix. Now, note that the matrix $\mathbf{M}_D^{\{\mathcal{F},\mathcal{F}_j\}}$ is obtained from matrix $\mathbf{M}_D^{\mathcal{F}\cup\mathcal{F}_j}$ simply by duplicating certain columns (namely, the columns corresponding to nodes that appear in both $\mathcal{F}$ and $\mathcal{F}_j$). Consider a node $v_l \in \mathcal{F}$. If $v_l \notin \mathcal{F}_j$, then the column corresponding to v_l within the first $p + 2f$ columns of $\mathbf{M}_D^{\{\mathcal{F},\mathcal{F}_j\}}$ will be linearly independent of all other columns in $\mathbf{M}_D^{\{\mathcal{F},\mathcal{F}_j\}}$ (since this column will also appear in the first $p + |\mathcal{F}\cup\mathcal{F}_j|$ columns of $\mathbf{M}_D^{\mathcal{F}\cup\mathcal{F}_j}$). This means that equation (10) can be satisfied only if $f_l[0] = 0$. On the other hand, if $f_l[0] \neq 0$, the only way for equation (10) to be satisfied is if $v_l \in \mathcal{F}_j$ and $\bar{f}_l[0] = f_l[0]$. In other words, if equation (8) is satisfied, any malicious node that commits an error during the first time-step will appear in set $\mathcal{F}_j$, and its additive error can be found by the IDS.

Next, note from (8) that the first p columns of $\mathbf{M}_D^{\{\mathcal{F},\mathcal{F}_j\}}$ will be linearly independent of each other and of all other columns in that matrix (since these columns also appear in $\mathbf{M}_D^{\mathcal{F}\cup\mathcal{F}_j}$ and are not duplicated in $\mathbf{M}_D^{\{\mathcal{F},\mathcal{F}_j\}}$). This means that the only way for equation (10) to be satisfied is if $\bar{\mathbf{y}}[0] = \mathbf{y}[0]$. Thus, the IDS has also recovered the outputs of the plant that were injected into the network at time-step $k = 0$.

At this point, the IDS knows $\mathbf{y}[0]$ and the identities of those nodes in $\mathcal{F}$ that committed errors during time-step 0, along with the exact values of their additive errors. The IDS can then use (3) to obtain the transmitted values of all nodes at time-step $k = 1$ as

$$\mathbf{z}[1] = \mathbf{W}\mathbf{z}[0] + \mathbf{H}\mathbf{y}[0] + \mathbf{B}_{\mathcal{F}_j}\bar{\mathbf{f}}[0] \ .$$

Now, using the identity

$$\mathbf{t}[1:D+1] = \mathbf{\Theta}_D\mathbf{z}[1] + \mathbf{M}_D^{\mathcal{F}}\mathbf{v}[1:D] \ ,$$

the IDS can repeat the above process to find the values of $\mathbf{y}[1]$ along with the identities of the nodes that are malicious during time-step $k = 1$. By repeating the above procedure for all positive values of k, the IDS can obtain the identities of all malicious nodes and the errors that they commit, along with the source streams $\mathbf{y}[k]$ for all k, simply by listening to the transmissions of the nodes in $\mathcal{T}$. ■

V. Network Topology Conditions for Misbehavior Identification and Data Recovery

Theorem 1 provides a decoding procedure for the IDS provided that condition (8) is true. In this section, we will relate this condition to the topology of the network.

A. System Inversion

The quantities $\mathbf{y}[k]$ and $\mathbf{f}[k]$ in (3) are unknown to the IDS, and so linear systems of this type are termed *linear systems with unknown inputs*.[3] For such systems, it is often of interest to "invert" the system in order to reconstruct some or all of the unknown inputs, and this problem has been studied under the moniker of *dynamic system inversion* [13]. We will now apply some pertinent results on system inversion to the problem of detecting and identifying malicious nodes in the wireless control network. First, recall that for any set $\mathcal{Q} \subseteq \mathcal{V}$, the transfer function of the linear system (5) is $\mathbf{P}(z) = \mathbf{T}\left(z\mathbf{I} - \mathbf{W}\right)^{-1}\mathbf{B}_{\mathcal{Q}}$, which is a $|\mathcal{T}|\times(p+|\mathcal{Q}|)$ matrix of rational functions of z.

Definition 1: The system (5) is said to have an L-delay inverse if there exists a system with transfer function $\widehat{\mathbf{P}}(z)$ such that $\widehat{\mathbf{P}}(z)\mathbf{P}(z) = z^{-L}\mathbf{I}_{p+|\mathcal{Q}|}$. The system is invertible if it has an L-delay inverse for some finite L. The least integer L for which an L-delay inverse exists is called the inherent delay of the system. □

In order for the system to be invertible, its transfer function must have rank $p + |\mathcal{Q}|$ over the field of rational functions in z. The following result from [13] and [14] provides a test for invertibility in terms of the matrices $\mathbf{W}, \mathbf{B}_{\mathcal{Q}}$ and $\mathbf{T}$.

Theorem 2 ([13], [14]): For any nonnegative integer L,

$$\text{rank}(\mathbf{M}_L^{\mathcal{Q}}) \leq p + |\mathcal{Q}| + \text{rank}(\mathbf{M}_{L-1}^{\mathcal{Q}}) \tag{11}$$

with equality if and only if the system has an L-delay inverse (note that $\text{rank}(M_{-1}^{\mathcal{Q}})$ is defined to be zero). If the system is invertible, its inherent delay will not exceed $L = N - p - |\mathcal{Q}| + 1$. □

We will now relate the algebraic test from the above theorem to a *graph-theoretic* test for invertibility.

B. Structured Systems

A linear system of the form (5) is said to be *structured* if each entry of the matrices $\mathbf{W}, \mathbf{B}_{\mathcal{Q}}$ and $\mathbf{T}$ is either a fixed zero or an independent free parameter [15]. Interestingly, such systems have certain properties that can be inferred purely from the zero/nonzero structure of the system matrices; these properties are *generic*, meaning that they will

[3] In our case, the set $\mathcal{F}$ (and thus the matrix $\mathbf{B}_{\mathcal{F}}$) is also unknown to the IDS, so the system given by (3) and (4) is more general than the linear systems with unknown inputs commonly considered in the literature.

hold for almost any choice of free parameters (i.e., the set of parameters for which the property does not hold has Lebesgue measure zero [15]). Of particular relevance to this paper is the *generic normal rank* of the transfer function matrix of a structured system, which is the maximum rank (over the field of rational functions in z) of the transfer function matrix over all possible choices of free parameters.

To analyze structural properties of linear systems of the form (5), one associates a graph $\mathcal{H}$ with the structured set $(\mathbf{W}, \mathbf{B}_\mathcal{Q}, \mathbf{T})$ as follows. The vertex set of $\mathcal{H}$ is given by $\mathcal{V} \cup \mathcal{I} \cup \mathcal{O}$, where $\mathcal{V} = \{v_1, v_2, \ldots, v_N\}$ is the set of state vertices, $\mathcal{I} = \{i_1, i_2, \ldots, i_{p+|\mathcal{Q}|}\}$ is the set of input vertices, and $\mathcal{O} = \{o_1, o_2, \ldots, o_{|\mathcal{T}|}\}$ is the set of output vertices. The edge set of $\mathcal{H}$ is given by $\mathcal{E}_{vv} \cup \mathcal{E}_{iv} \cup \mathcal{E}_{vo}$, where $\mathcal{E}_{vv} = \{(v_j, v_l) \mid \mathbf{W}_{lj} \neq 0\}$, $\mathcal{E}_{iv} = \{(i_j, v_l) \mid \mathbf{B}_{\mathcal{Q},lj} \neq 0\}$, and $\mathcal{E}_{vo} = \{(v_j, o_l) \mid \mathbf{T}_{lj} \neq 0\}$ (where $\mathbf{W}_{lj}$ indicates entry (l, j) of matrix $\mathbf{W}$, and so forth). The following theorem characterizes the generic normal rank of the transfer function of a structured linear system in terms of the graph $\mathcal{H}$.

Theorem 3 ([15], [16]): Let the graph of a structured linear system be given by $\mathcal{H}$. Then the generic normal rank of the transfer function of the system is equal to the maximal size of a linking in $\mathcal{H}$ from $\mathcal{I}$ to $\mathcal{O}$. □

The above result says that if the graph of the structured system (5) has $p + |\mathcal{Q}|$ vertex disjoint paths from the inputs to the outputs, then for almost any choice of free parameters in $\mathbf{W}$, $\mathbf{B}_\mathcal{Q}$ and $\mathbf{T}$, the transfer function matrix $\mathbf{T}(z\mathbf{I} - \mathbf{W})^{-1}\mathbf{B}_\mathcal{Q}$ will have full column rank. Based on Theorem 2, this will mean that the first $p + |\mathcal{Q}|$ columns of the matrix $\mathbf{M}^\mathcal{Q}_{N-p-|\mathcal{Q}|+1}$ will be linearly independent of all other columns in $\mathbf{M}^\mathcal{Q}_{N-p-|\mathcal{Q}|+1}$.

We now have a graph-theoretic characterization of the invertibility of linear structured systems, and are in place to apply this to the problem of identifying malicious behavior and recovering the plant outputs in the WCN.

C. Topological Conditions for Identifying Malicious Nodes

From Theorem 1 and Theorem 2, the IDS can identify up to f malicious nodes if the linear system given by the tuple $(\mathbf{W}, \mathbf{B}_\mathcal{Q}, \mathbf{T})$ is invertible for every set $\mathcal{Q} \subset \mathcal{V}$ of up to $2f$ nodes. To verify that this property holds, note that for any given set $\mathcal{Q}$, the tuple $(\mathbf{W}, \mathbf{B}_\mathcal{Q}, \mathbf{T})$ essentially defines a structured linear system, with the only exception being that the nonzero entries in the matrices $\mathbf{E}_\mathcal{Q}$ (where $\mathbf{B}_\mathcal{Q} = \begin{bmatrix}\mathbf{H} & \mathbf{E}_\mathcal{Q}\end{bmatrix}$) and $\mathbf{T}$ are taken to be "1", rather than free parameters. However, this is of no consequence, since each nonzero entry in those matrices appears in a row and column by itself, and thus can essentially be "scaled" to a free parameter by an appropriate redefinition of the inputs and outputs (e.g., see [17]). Thus, we can proceed with applying the above results on structured system theory to the tuple $(\mathbf{W}, \mathbf{B}_\mathcal{Q}, \mathbf{T})$, which brings us to the following result.

Theorem 4: Let $\bar{\mathcal{G}} = \{\mathcal{V} \cup \mathcal{S}, \mathcal{E} \cup \mathcal{E}_I\}$ denote the graph of the wireless control network $\mathcal{G}$ augmented with the sensor vertices $\mathcal{S}$ and the corresponding edges. Let $\mathcal{T} \subset \mathcal{V}$ denote the set of monitored nodes. Suppose that for every possible set $\mathcal{Q} \subset \mathcal{V}$ of $2f$ nodes, the graph $\bar{\mathcal{G}}$ contains a $(p + 2f)$–linking from $\mathcal{S} \cup \mathcal{Q}$ to $\mathcal{T}$. Then, for almost any element $(\mathbf{W}, \mathbf{H}, \mathbf{G}) \in \Psi_s$ (if it is nonempty), there exists an integer $D \leq N - p - 2f + 1$ such that the IDS can recover the outputs of the plant and identify all malicious nodes with a delay of at most D time-steps, as long as there are no more than f malicious nodes in any set of D contiguous time-steps. □

Proof: For any set $\mathcal{Q} \subset \mathcal{V}$ of $2f$ nodes, consider the graph[4] $\mathcal{H}_\mathcal{Q}$ associated with the structured set $(\mathbf{W}, \mathbf{B}_\mathcal{Q}, \mathbf{T})$. To obtain this graph, start by taking the graph of the network $\mathcal{G}$. To this graph, add $p + 2f$ input vertices (denoted by $\mathcal{I}$) which will connect to the nodes in the graph according to the structure of the input matrix $\mathbf{B}_\mathcal{Q}$. Specifically, p of these input vertices correspond to the plant sensors $\mathcal{S}$ (which produce $\mathbf{y}[k]$), and each of these has outgoing edges to the nodes in $\mathcal{V}_S$ (specified by the structure of matrix $\mathbf{H}$). The other $2f$ input vertices each have a single outgoing edge to a node in $\mathcal{Q}$ (corresponding to the single 1 in each column of $\mathbf{E}_\mathcal{Q}$). Next, add $|\mathcal{T}|$ output vertices (denoted by the set $\mathcal{O}$), and place a single edge from each node in the set $\mathcal{T}$ to a node in $\mathcal{O}$, corresponding to the single nonzero entry in each row of the matrix $\mathbf{T}$. Furthermore, add a self loop to every state vertex corresponding to the nonzero entries on the diagonal of the matrix $\mathbf{W}$.

From the statement of the theorem, note that graph $\bar{\mathcal{G}}$ contains a linking of size $p + 2f$ from $\mathcal{S} \cup \mathcal{Q}$ to $\mathcal{T}$, for any set $\mathcal{Q}$ of $2f$ nodes. This linking also exists in the graph $\mathcal{H}_\mathcal{Q}$, since $\bar{\mathcal{G}}$ is a subgraph of $\mathcal{H}_\mathcal{Q}$.[5] This linking can be extended to a linking from the entire set $\mathcal{I}$ to $\mathcal{T}$ in $\mathcal{H}_\mathcal{Q}$ simply by including the edges from the set $\mathcal{I} \setminus \mathcal{S}$ to the set $\mathcal{Q}$. Finally, this linking can be further extended to a linking from $\mathcal{I}$ to $\mathcal{O}$ simply by including the edges from each vertex in $\mathcal{T}$ to the corresponding output vertex in $\mathcal{O}$. From Theorem 3, we see that the system $(\mathbf{W}, \mathbf{B}_\mathcal{Q}, \mathbf{T})$ will be invertible for almost any choice of matrices $\mathbf{W}$ and $\mathbf{H}$ (subject to the required sparsity patterns). This genericness implies that invertibility will hold simultaneously for all of the sets $(\mathbf{W}, \mathbf{B}_\mathcal{Q}, \mathbf{T})$ for every set $\mathcal{Q}$ of $2f$ nodes with almost any choice of free parameters in the matrices $\mathbf{W}$ and $\mathbf{H}$. From Theorem 2, the first $p + |\mathcal{Q}|$ columns of the matrix $\mathbf{M}^\mathcal{Q}_{N-p-2f+1}$ will be linearly independent of each other and of all other columns in $\mathbf{M}^\mathcal{Q}_{N-p-2f+1}$. Thus, condition (8) in Theorem 1 is satisfied, and the IDS can uniquely determine the identities of the malicious nodes, as well as the values of the plant outputs, based on the transmissions of the nodes in $\mathcal{T}$, with a delay of at most $N - p - 2f + 1$ time-steps.

Finally, we show that there is a tuple $(\mathbf{W}, \mathbf{H}, \mathbf{G})$ in the set Ψ_s (which contains all stabilizing structured matrices for the plant and is assumed to be nonempty) that allows the IDS to recover the desired information. This is done by noting that the set of matrices for which the system is stable has nonzero measure in $\mathbb{R}^r$ (where r is the number of free parameters in the matrices $\mathbf{W}$ and $\mathbf{H}$). More precisely, if $\lambda \in \mathbb{R}^r$ denotes

[4]The notation $\mathcal{H}_\mathcal{Q}$ is used to denote the fact that this graph is associated with the structured set $(\mathbf{W}, \mathbf{B}_\mathcal{Q}, \mathbf{T})$, for a particular set $\mathcal{Q}$ of $2f$ nodes.

[5]Specifically, it is the graph obtained by dropping the output vertices and the $2f$ input vertices connecting to the set $\mathcal{Q}$ in $\mathcal{H}_\mathcal{Q}$.

a numerical vector of free parameters in $\mathbf{W}$ and $\mathbf{H}$ that produces stability (e.g., obtained from the procedure in [4]), the closed loop system will remain stable for any parameter vectors λ^* satisfying the component-wise inequalities $\lambda - \epsilon\mathbf{1} \leq \lambda^* \leq \lambda + \epsilon\mathbf{1}$, for sufficiently small $\epsilon > 0$; this is because the eigenvalues of a matrix vary continuously with the parameters in that matrix. Thus, the set of stabilizing parameters has measure at least $(2\epsilon)^r > 0$, whereas the set of parameters for which the system is not invertible has measure zero. Thus, for almost any $(\mathbf{W}, \mathbf{H}, \mathbf{G}) \in \Psi_s$, the system is stable *and* allows the IDS to recover the plant outputs and identify malicious behavior. ■

Theorem 4 characterizes the set of nodes $\mathcal{T}$ that the IDS should observe in order to achieve its objectives. However, the fact that the theorem is framed in terms of *all possible* sets $\mathcal{Q}$ of $2f$ nodes makes it somewhat unwieldy. One can come up with a more compact condition when the entire network is sufficiently well connected, as follows.

Corollary 1: Suppose that the network $\mathcal{G}$ has connectivity at least $p+2f$, and that each sensor in $\mathcal{S}$ connects to at least $p+2f$ source nodes. Let $\mathcal{T} \subseteq \mathcal{V}$ be any set of at least $p+2f$ nodes. Then, for almost any element $(\mathbf{W}, \mathbf{H}, \mathbf{G}) \in \Psi_s$ (if it is nonempty), there exists an integer $D \leq N - p - 2f + 1$ such that the IDS can recover the outputs of the plant and identify all malicious nodes with a delay of D time-steps, as long as there are no more than f malicious nodes in any set of D contiguous time-steps.

Proof: For any set $\mathcal{Q} \subset \mathcal{V}$ of $2f$ nodes, note that $\mathcal{V}_S \setminus \mathcal{Q}$ has at least p nodes (since $|\mathcal{V}_S| \geq p + 2f$). Since each sensor in $\mathcal{S}$ connects to at least $p + 2f$ source nodes, each sensor will connect to at least p nodes in the set $\mathcal{V}_S \setminus \mathcal{Q}$. By Hall's Theorem (e.g., see [12]), there is a linking of size p from $\mathcal{S}$ to $\mathcal{V}_S \setminus \mathcal{Q}$ (this is also called a *matching*). Let $\bar{\mathcal{V}}_S \subset \mathcal{V}_S$ be the set of p source nodes contained in this matching, and note that $\bar{\mathcal{V}}_S \cup \mathcal{Q}$ has $p+2f$ nodes. Since the network has connectivity $p+2f$, Lemma 1 shows that there is a linking of size $p + 2f$ from $\bar{\mathcal{V}}_S \cup \mathcal{Q}$ to $\mathcal{T}$. Thus, the graph $\bar{\mathcal{G}} = \{\mathcal{V} \cup \mathcal{S}, \mathcal{E} \cup \mathcal{E}_I\}$ contains a linking of size $p+2f$ from $\mathcal{S} \cup \mathcal{Q}$ to $\mathcal{T}$ for any set $\mathcal{Q}$ of $2f$ nodes. Theorem 4 is thus satisfied, from which the result follows. ■

Note that the above corollary indicates that in networks with connectivity $p + 2f$ or higher, *any* set of $p + 2f$ nodes can be chosen to be observed by the IDS in order to recover the desired information about the system. For example, consider the wireless control network shown in Fig. 1. The source nodes $\mathcal{V}_S = \{v_1, v_2, v_3\}$ have access to the plant's (scalar) output $y[k]$ at each time-step. Note that the connectivity of the network is $\kappa = 3$, and since there is a single sensor on the plant ($p = 1$) that connects to three nodes, Corollary 1 indicates that the IDS can detect and identify up to $f = \lfloor \frac{\kappa - p}{2} \rfloor = 1$ malicious node, simply by monitoring the transmissions of any $p + 2f$ nodes (e.g., the set $\mathcal{T} = \{v_3, v_6, v_9\}$).

Remark 1: The approach presented in this paper can be viewed as a form of *network coding* that has been recently proposed by the communications community, where nodes in a network transmit linear combinations of incoming packets rather than simply routing them (e.g., see [18], [19] and the references therein). It was shown in [18], [19] that the capacity of the network must be at least $p + 2f$ in order to transmit information from p sources to a set of sinks, despite the presence of f attackers. Our work shows that this bound holds for the WCN proposed in [4] using a purely linear system theoretic approach, even though the linear combinations have been chosen to obtain stability. Further investigation of the applicability of resilient network coding to control network design is a ripe avenue for research. □

REFERENCES

[1] L. Schenato, B. Sinopoli, M. Franceschetti, K. Poolla, and S. S. Sastry, "Foundations of control and estimation over lossy networks," *Proc. of the IEEE*, vol. 95, no. 1, pp. 163–187, Jan. 2007.

[2] J. P. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proc. of the IEEE*, vol. 95, no. 1, pp. 138–162, Jan. 2007.

[3] V. Gupta, A. F. Dana, J. Hespanha, R. M. Murray, and B. Hassibi, "Data transmission over networks for estimation and control," *IEEE Transactions on Automatic Control*, vol. 54, no. 8, pp. 1807–1819, Aug. 2009.

[4] M. Pajic, S. Sundaram, J. Le Ny, R. Mangharam, and G. J. Pappas, "The Wireless Control Network: Synthesis and Robustness," in *Proc. of the 49th IEEE Conference on Decision and Control*, 2010.

[5] F. Pasqualetti, A. Bicchi, and F. Bullo, "Distributed intrusion detection for secure consensus computations," in *Proc. of the 46th IEEE Conference on Decision and Control*, 2007, pp. 5594–5599.

[6] A. Giani, G. Karsai, T. Roosta, A. Shah, B. Sinopoli, and J. Wiley, "A testbed for secure and robust SCADA systems," in *Proc. of the 14th IEEE Real-time and Embedded Technology and Applications Symposium*, 2008, pp. 1–4.

[7] S. Sundaram and C. N. Hadjicostis, "Distributed function calculation via linear iterations in the presence of malicious agents," in *Proc. of the American Control Conference*, 2008, pp. 1350–1361.

[8] S. Amin, A. A. Cárdenas, and S. S. Sastry, "Safe and secure networked control systems under denial-of-service attacks," in *HSCC '09: Proc. of the 12th International Conference on Hybrid Systems: Computation and Control*. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 31–45.

[9] K. Stouffer, J. Falco, and K. Scarfone, "Guide to industrial control systems (ICS) security," National Institute of Standards and Technology, Tech. Rep. 800-82, Sep. 2008.

[10] S. Northcutt and J. Novak, *Network Intrusion Detection*, 3rd ed. New Riders Publishing, 2003.

[11] T. Roosta, D. K. Nilsson, U. Lindqvist, and A. Valdes, "An intrusion detection system for wireless process control systems," in *Proc. of the 5th IEEE International Conference on Mobile Ad Hoc and Sensor Systems*, 2008, pp. 866–872.

[12] D. B. West, *Introduction to Graph Theory*. Prentice-Hall Inc., Upper Saddle River, New Jersey, 2001.

[13] M. K. Sain and J. L. Massey, "Invertibility of linear time-invariant dynamical systems," *IEEE Transactions on Automatic Control*, vol. AC-14, no. 2, pp. 141–149, Apr. 1969.

[14] A. S. Willsky, "On the invertibility of linear systems," *IEEE Transactions on Automatic Control*, vol. 19, no. 2, pp. 272–274, June 1974.

[15] J.-M. Dion, C. Commault, and J. van der Woude, "Generic properties and control of linear structured systems: a survey," *Automatica*, vol. 39, no. 7, pp. 1125–1144, July 2003.

[16] J. W. van der Woude, "A graph-theoretic characterization for the rank of the transfer matrix of a structured system," *Mathematics of Control, Signals and Systems*, vol. 4, no. 1, pp. 33–40, Mar. 1991.

[17] S. Sundaram and C. N. Hadjicostis, "Distributed function calculation and consensus using linear iterative strategies," *IEEE Journal on Selected Areas in Communications: Special Issue on Control and Communications*, vol. 26, no. 4, pp. 650–660, May 2008.

[18] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, M. Médard, and M. Effros, "Resilient network coding in the presence of Byzantine adverseries," *IEEE Transactions on Information Theory*, vol. 54, no. 6, pp. 2596–2603, June 2008.

[19] R. Koetter and F. R. Kschischang, "Coding for errors and erasures in random network coding," *IEEE Transactions on Information Theory*, vol. 54, no. 8, pp. 3579–3591, Aug. 2008.