2016 American Control Conference (ACC)
Boston Marriott Copley Place
July 6-8, 2016. Boston, MA, USA

Network Invariants for Optimal Input Detection

Rajasekhar Anguluri, Rahul Dhal, Sandip Roy, and Fabio Pasqualetti

Abstract— **This paper studies a detection problem for network systems, where changes in the statistical properties of an input driving certain network nodes has to be detected by sparse and remotely located sensors. We explicitly derive the Maximum A Posteriori (MAP) detector, and characterize its performance as a function of the network parameters, and the location of the sensor nodes. We show that, in the absence of measurement noise, the detection performance obtained when sensors are located on a network cut is not worse than the performance obtained by measuring all nodes of the subnetwork induced by the cut and not containing the input node. Conversely, in the presence of measurement noise, we show that the detection performance may increase or decrease with the graphical distance between the input node and the sensors. We view the propagative properties of the network as an invariant enforced by the structure and weights, and we remark that such invariant properties may be effectively used for the design and operation of secure cyber-physical systems.**

I. Introduction

Cyber-physical systems are extremely vulnerable, as recently demonstrated by attacks on automobiles, medical devices, and the energy grid. Such attacks not only violate confidentiality by releasing personal information but, more importantly, affect the integrity of the system often with life threatening consequences. While numerous mechanisms have been developed to protect the cyber components of a cyber-physical system, protection of the physical layer and of the core inter-layers functionalities remains an outstanding problem. Here we highlight that the fundamental *invariant* relations governing dynamical large-scale and interconnected cyber-physical systems may in fact enable the design and operation of secure and high-assurance cyber-physical systems.

The concept of invariants has a long history across a number of domains, including formal methods for computing systems [1], continuum physics [2], linear algebra [3], and the biological [4] and behavioral sciences [5], [6]. Although definitions for invariants vary, the term is broadly used to describe properties of a system/object that are guaranteed or remain unchanged under a transformation. In addition to serve as descriptive tool, system invariants have been recently used to verify correctness of system operations [7], [8], [9].

This study is concerned with detecting statistical abnormalities in a local stochastic input to a network, using remote time-course measurements of the network dynamics. The basic idea is that the network topological structure enforces invariants or relationships between signals at different locations in the network, which means that changes to the dynamics at one location can potentially be detected using measurement signatures at other locations in the network. Thus, local changes implicate a predictable propagative response across the network, which can in theory be used to identify the change from sparse and remote measurements.

In this work, we consider a network system driven by a stochastic signal, and we focus on the problem of detecting abrupt changes in the input statistics from measurement taken at different network nodes. This specific detection problem is motivated by emerging monitoring needs for cyber-physical networks, where intrusions or abnormalities in cyber and human components may cause subtle changes in stochastic driving or input signals, and ultimately incur significant risk to the network operation. As it may be impractical or impossible to directly monitor these input signals, we exploit network invariants to identify changes in the driving signals. Particularly, we quantify the relation between the detection performance as a function of the network topological properties and the location of the sensor nodes. Our analysis leads, for instance, to the counterintuitive results that, depending on the network weights and structure, the detection performance may improve as the graphical distance between the input signal and the sensor node increases. Our results have immediate applicability for the secure operation of cyber-physical systems. In fact, our results (i) inform the optimal positioning of sensors for the detection of failure of system components or malicious tampering modeled by unknown stochastic inputs, (ii) allow the detection of unexpected modification of the system structure, because such changes would inevitably modify the original detection profile, and (iii) provide network design guidelines to facilitate or prevent measurability of certain network signals.

Conceptual example A car is a prototypical cyber-physical network, comprising diverse mechanical and electrical components, a human operator (driver), and pervasive cyber-systems (control systems, embedded computing, networking features). The reaction of the human driver is a key input to the vehicle's dynamics of motion. Research on fatigue and cognitive performance has shown that moderate fatigue leads to subtle statistical changes in driver reaction times and actions [10] (e.g. steering wheel deflections), including especially an increase in variability in these characteristics. The driver's changed physiological response (e.g. reaction time) is difficult to measure directly in real time, but these subtle changes are reflected in other signals (e.g., the acceleration and steering-wheel deviation profile), sensed distances to other vehicles, etc. Many of these internal signals in the car

This material is based upon work supported in part by ONR award #N00014-14-1-0816. Rajasekhar Anguluri and Fabio Pasqualetti are with the Mechanical Engineering Department, University of California at Riverside, rajasekhar.anguluri@ieee.org, fabiopas@engr.ucr.edu. Sandip Roy is with the School of Electrical Engineering and Computer Science, Washington State University, sroy@eecs.wsu.edu. Rahul Dhal is with Epis Inc., rahuldhal@epis.com.

are sensed, and could be used to detect moderate fatigue. In a similar vein, cyber-attacks may alter driving signals in the network (e.g., change the setpoint or inject variability/delay in the cruise control). Likewise, these attacks could be detected from remote measured signals, by exploiting the essential invariants in the internal workings of the vehicle.

Related work Cyber-physical security has recently emerged as an interdisciplinary research field at the intersection of classic areas, including computer and information security and fault detection and isolation [11], [12], [13], [14], [15]. While a large body of literature focuses on static cyber-physical security, which defines both the attack and defense mechanisms at a single time instance, more recent efforts recognize the importance of dynamic cyber-physical security, for which game theory [16], [17], [18], information theory [19], [20], [21], and control theory [22], [23], [24] may offer more applicable techniques. As a drawback of these approaches, critical assumptions are made on the system dynamics and structure, so that results are usually confined to specific scenarios. Instead, in this work we propose to build security mechanisms based on fundamental invariant relations enforced by the laws of physics or careful design. Although different systems may satisfy different invariants, the core methods remain applicable across various domains.

Contribution The main contribution of this paper is twofold. First, we formulate a detection problem for network systems, where a change in the statistical properties of an input affecting the behavior of some nodes must be detected from sparse and remote measurements. We explicitly characterize the Maximum A Posteriori detector, and quantify its error probability. Second, we study the detector performance as a function of the sensors locations. We prove that, in the absence of measurement noise, the detection performance of a set of sensors forming a cut of the network is as good as the performance obtained by measuring all nodes of the subnetwork identified by the cut and not containing the node affected by the input. Thus, in the absence of measurement noise the detection performance is non-increasing with the graphical distance between the input node and the sensors located on a network cut. Conversely, in the presence of measurement noise, we show that the detection performance may increase or decrease with the graphical distance between the input node and the sensors, depending on the network structure and weights. Our findings suggest that invariants enforced by the network structure and weights can effectively be leveraged to optimize design and operation of networks.

Paper organization The rest of the paper is organized as follows. Section II contains our network model and preliminary notions. The main results of this paper are presented in Section III and Section IV. Finally, Section V contains an illustrative example, and Section VI concludes the paper.

II. Problem setup and preliminary results

To highlight the role of the network topology in detection, we consider a simple linearized network model. Formally, a network with n components or subsystems or nodes, labeled $i = 1, \dots, n$, is considered. Each node $i = 1, \dots, n$ is assumed to have a scalar state $x_i[k]$, which evolves along a discrete time index $k = 0, 1, 2, \dots$. The state vector $\mathbf{x}[k] = [x_1[k] \cdots x_n[k]]^T$ is governed by:

$$\mathbf{x}[k+1] = G(\Gamma)\mathbf{x}[k] + \mathbf{e}_q w[k], \tag{1}$$

where $G(\Gamma)$ is commensurate with the n-vertex digraph Γ but otherwise may be arbitrary, $\mathbf{e}_q$ is a $0-1$ indicator vector where the q-th entry equals to 1, and $w[k]$ is a stochastic drive or input which enters the network through the q-th node.

Our primary focus is on distinguishing whether the stochastic input $w[k]$ is governed by a nominal statistical model (e.g., for an alert driver or operational cyber system in the car example), or alternately a model with altered statistics (e.g., for a moderately fatigued driver or attacked cyber system). Formally, for this study, the stochastic input $w[k]$ is assumed to be governed by one of two possible statistical hypotheses, referred as Hypothesis 1 (H_1) and Hypothesis 2 (H_2). For simplicity, under hypothesis H_1, the stochastic input is assumed to be a stationary white process, with the input $w[k]$ at each time distributed as $\mathcal{N}(m_1, \sigma_1^2)$. Meanwhile, under hypothesis H_2, the stochastic input remains white, but with modified distribution $\mathcal{N}(m_2, \sigma_2^2)$. The hypothesis that is in force is denoted by H, i.e., $H = H_1$ or $H = H_2$, and is assumed to remain the same over the duration of the process. Two special cases of this formulation, where only the means differ ($m_1 \neq m_2$ but $\sigma_1^2 = \sigma_2^2$) or only the variances differ ($m_1 = m_2$ but $\sigma_1^2 \neq \sigma_2^2$), will be considered in some of our analyses. These cases are denoted as the *mean-changed* model and *variance-changed* model, respectively.

The network dynamics are assumed to be measured at m nodes, say $J = \{j_1, \dots, j_m\}$, where J is referred to as the sensor set. Specifically, at each time k, a monitor receives

$$\mathbf{y}_J[k] = \begin{bmatrix} \mathbf{e}_{j_1}^T \\ \vdots \\ \mathbf{e}_{j_m}^T \end{bmatrix} \mathbf{x}[k] + \mathbf{n}[k],$$

where the observation noise $\mathbf{n}[k]$ is an i.i.d gaussian process distributed at each time as $\mathcal{N}(n\mathbf{1}, \sigma_n^2 I)$. Several of our analyses will distinguish the role of measurement noise in detection, hence the noise-free case ($\sigma_n^2 = 0$) will be compared with the noisy case. The monitor aims to use the measurements over the interval $k = 0, \dots, k_f - 1$ to decide which hypothesis is in force. To study the detection problem, we combine the observations over the interval into a full observation vector,

$$Y_J = \begin{bmatrix} \mathbf{y}_J[0] \\ \vdots \\ \mathbf{y}_J[k_f - 1] \end{bmatrix}.$$

The monitor is tasked with using a detector to identify the hypothesis on the input signal, from the observation vector Y_J. Specifically, the detector maps the observation vector Y_J to a detected hypothesis $\widehat{H}$, which may be either H_1 or H_2. In this work, we will primarily consider *maximum a posteriori probability* (MAP) detection of the hypothesis by the

monitor (results with respect to arbitrary optimal detectors will be discussed in Section III). A MAP detector chooses the hypothesis that has more likely given the observation sequence as the detected hypothesis $\widehat{H}$. In our formulation, the MAP detector has been considered in the following way:

$$\Pr(H = H_2|Y_J) \underset{\widehat{H}=H_1}{\overset{\widehat{H}=H_2}{\gtrless}} \Pr(H = H_1|Y_J). \tag{2}$$

The focus of our analysis is on characterizing the structure and performance of the MAP detector in terms of the graph Γ, so as to inform sensor placement for effective detection. In other words, we study how the invariants imposed by the network can be leveraged to identify the correct hypothesis. The performance of the MAP detector is naturally measured by its probability of error, which is the probability that the detected hypothesis $\widehat{H}$ is not equal to the hypothesis in force, H. The probability of error can be computed as

$$\begin{aligned} P_E &= \Pr(\widehat{H} = H_2 \,|\, H = H_1)\Pr(H = H_1) \\ &\quad + \Pr(\widehat{H} = H_1 \,|\, H = H_2)\Pr(H = H_2). \end{aligned} \tag{3}$$

Direct and indirect characterizations of the error probability will be undertaken, to understand the relationship between the graph topology and the detector performance.

III. Algebraic analysis of the MAP detector

Since the measurement signal is a filtration of a white Gaussian input signal, classic results on hypothesis testing using Gaussian observations can be used to obtain algebraic characterizations of the detector and its performance. Then, we use the results to extract structural and graphical insights. Although the analysis holds for sensor set with multiple nodes, for simplicity of presentation we focus on the case where only one node is measured, i.e., $J = j$. We study the MAP detector for the change in mean problem ($\sigma_1^2 = \sigma_2^2 = \sigma^2$). Similar analysis holds for change in variance problem also. From Bayes' rule, Eq (2) can be rewritten as

$$\frac{f_{Y_j|H_2}\Pr(H_2)}{f_{Y_j}} \underset{\widehat{H}=H_1}{\overset{\widehat{H}=H_2}{\gtrless}} \frac{f_{Y_j|H_1}\Pr(H_1)}{f_{Y_j}},$$

where f_{Y_j} is the joint probability density function (pdf) of observation vector Y_j and $f_{Y_j|H_i}$ is the conditional pdf of observation vector Y_j given the hypothesis H_i, with $i \in \{1, 2\}$. Assuming $f_j \geq 0$, it suffices to compare

$$f_{Y_j|H_2}\Pr(H_2) \underset{\widehat{H}=H_1}{\overset{\widehat{H}=H_2}{\gtrless}} f_{Y_j|H_1}\Pr(H_1). \tag{4}$$

The noisy observations captured by the monitor at the desired cut (or node) can be modeled by the following expression:

$$\begin{aligned} y_j[k] &= \mathbf{e}_j^T x[k] + n[k], \\ &= \mathbf{e}_j^T x[k] \left[G^k x[0] + \sum_{i=0}^{k-1} G^{k-(i+1)} \mathbf{e}_q w[i] \right] + n[k], \end{aligned} \tag{5}$$

where $y_J[k]$ and $n[k]$ are the scalar versions of $\mathbf{y}_J[k]$ and $\mathbf{n}[k]$ respectively. We focus on the case where $x[0] = 0$, and obtain

$$Y_j = \underbrace{\begin{bmatrix} \mathbf{e}_j^T\mathbf{e}_q & 0 & \cdots & 0 \\ \mathbf{e}_j^T G\mathbf{e}_q & \mathbf{e}_j^T\mathbf{e}_q & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{e}_j^T G^{k_f-1}\mathbf{e}_q & \mathbf{e}_j^T G^{k_f-2}\mathbf{e}_q & \cdots & \mathbf{e}_j^T\mathbf{e}_q \end{bmatrix}}_{R} \begin{bmatrix} w[0] \\ w[1] \\ \vdots \\ w[k_f-1] \end{bmatrix} + \begin{bmatrix} n[1] \\ n[2] \\ \vdots \\ n[k_f] \end{bmatrix}. \tag{6}$$

The statistics of Y_j can be characterized as follows

$$\begin{aligned} \mathbb{E}[Y_j|H_i] &= \mu_i = m_i R\mathbf{1} + n\mathbf{1}, \\ \mathrm{Cov}[Y_j|H_i] &= \mathbf{\Sigma} = \sigma^2 RR^T + \sigma_n^2 I, \end{aligned} \tag{7}$$

and

$$f_{Y_j|H_i} = \frac{1}{(2\pi)^{k_f/2}\sqrt{|\mathbf{\Sigma}|}} \exp\left(-\frac{(Y_j - \mu_i)^T \mathbf{\Sigma}^{-1} (Y_j - \mu_i)}{2} \right), \tag{8}$$

where $i \in \{1, 2\}$, $|\cdot|$ is matrix determinant operator, $\mathbf{1}$ is column vector of all 1, and I is standard identity matrix.[1]

Lemma 3.1: ***(MAP Detector)*** In the presence of additive, white and Gaussian measurement noise, the decision rule used by MAP detector for the hypotheses H_1 and H_2 is given by:

$$\begin{aligned} (R\mathbf{1})^T \mathbf{\Sigma}^{-1} Y_j \underset{\widehat{H}=H_1}{\overset{\widehat{H}=H_2}{\gtrless}} & \frac{1}{m_2 - m_1} \ln\left(\frac{\Pr(H_1)}{\Pr(H_2)}\right) \\ & + \left(\frac{m_1 + m_2}{2} + n\right) (R\mathbf{1})^T \mathbf{\Sigma}^{-1} R\mathbf{1}. \end{aligned} \tag{9}$$

Proof: By rearranging Eq (4) and taking the natural logarithm on both sides, we obtain the following inequality

$$\mathrm{LLR}(Y_j) = \ln\left(\frac{f_{Y_j|H_2}}{f_{Y_j|H_1}}\right) \underset{\widehat{H}=H_1}{\overset{\widehat{H}=H_2}{\gtrless}} \ln\left(\frac{\Pr(H_1)}{\Pr(H_2)}\right). \tag{10}$$

The left hand side of Eq (10) is often referred as log-likelihood ratio (LLR). By substituting Eq (8) into $\mathrm{LLR}(Y_j)$,

$$\begin{aligned} \mathrm{LLR}(Y_j) = (m_2 - m_1)(R\mathbf{1})^T \mathbf{\Sigma}^{-1} \\ \left[Y_j - \left(\frac{m_1 + m_2}{2} + n\right) R\mathbf{1} \right]. \end{aligned} \tag{11}$$

The claimed statement follows from Eq (10). ■

The detector equation Eq (9) bears an interesting interpretation. To see this, we note that $(R\mathbf{1})^T \Sigma^{-1} Y_j$ is the minimum mean square error (MMSE) estimate of the input sequence $w[k]$ from the observation sequence. The detector can thus be viewed as comparing the mean value of the input sequence estimate with the midpoint of the filtered means for the two hypotheses (with a correction for a priori probabilities). We next characterize the error probability of the MAP detector.

[1] See Remark 1 for a discussion of the invertibility of Σ.

Lemma 3.2: ***(MAP detector's Error Probability)*** The probability of error P_E in the detection process is given by:

$$P_E = \Pr(H_1)Q\left(\frac{\ln\left(\frac{\Pr(H_1)}{\Pr(H_2)}\right)}{2\gamma} + \gamma\right) + \Pr(H_2)Q\left(\frac{\ln\left(\frac{\Pr(H_2)}{\Pr(H_1)}\right)}{2\gamma} + \gamma\right), \quad (12)$$

where

$$\gamma = \frac{(m_2 - m_1)}{2}\sqrt{(R\mathbf{1})^T\mathbf{\Sigma}^{-1}(R\mathbf{1})}, \text{ and}$$

$$Q(x) = \frac{1}{\sqrt{2\pi}}\int_x^\infty \exp\left(-t^2/2\right) dt.$$

The proof of Lemma 3.2 follows because the observations are Gaussian [25], and the details are omitted here.

Remark 1: ***(Invertibility of Σ)*** The matrix Σ is invertible for $\sigma_n^2 \neq 0$. For the case that $\sigma_n = 0$, the matrix Σ may or may not be invertible. It is invertible in the case that the stochastic input and the measurement are collocated in the network ($e_q = e_j$). If the input and output are at different locations, two situations may result. In the atypical case that $e_j^T G^k e_q$ is identically zero for all k, then the stochastic input shows no signature in the measurement. In this case, the MAP detector simply selects the most likely a priori hypothesis, and the probability of error is the same as the a priori. This case corresponds to the circumstance that all modes of the network dynamics are either uncontrollable from the input location or unobservable from the output location. Alternately, if $e_j^T G^k e_q$ is non-zero for some k, then a lower-left block of Σ can be shown to be invertible, which means that a subset of the input sequence $(w[0], \ldots, w[k_f - k])$ can be exactly recovered from the measurement sequence, while the remaining points in the input sequence are not reflected in the measurements. Based on this discussion, a revised expression for the optimal detector can be developed, which permits computation of the detector and its error probability [24], [26]. Detection when measurement noise is negligible is discussed further in Section IV. □

Remark 2: ***(Asymptotic probability of error)*** Provided that G is stable, the probability of error can be shown to follow a dichotomy. Either $e_j^T G^k e_q$ is identically zero for all k, in which case the probability of error is identical to the a priori regardless of the initial condition. Otherwise, the probability of error can be shown to approach 0 exponentially with respect to k_f. Details of the argument are omitted; the reader is referred to [24] for a similar development. □

IV. Network analysis of the MAP detector

The graphical distance from the input location to the monitoring location(s) should influence the MAP detectors performance. Naively, one might expect monitors that are close to the input to be able to better distinguish the input characteristics (specifically, the mean and variance of the input signal). However, previous studies have shown that network structure may amplify a driving input over a distance and implicate nonminimum-phase dynamics [27] and hence, it is far from clear whether spatial proximity indeed yields improved detection. We study this question here.

In this section, we verify that graphical proximity does modulate detection performance, in the following sense. If sensors are placed in a such a way as to partition the graph, then a MAP detector using these sensors outperforms any other detectors that uses sensors in the partition not containing the input. Thus, closeness to the input location necessarily permits more effective detection, provided that the sensors completely separate the input. This result is formalized next; we assume in the analysis that $x[0]$ is independent of the future input sequence $w[0], \ldots, w[k_f - 1]$.

To formalize the result, a sensor set J_1, which forms a node-cutset of the graph Γ is considered. The probability of error of the MAP detector using Y_{J_1} is denoted by E_1. This detector performance is compared to the detection performance of a second sensor set J_2. The sensor set J_2 may contain (i) any nodes in the partition of Γ formed by J_1 that does not contain the input location and (ii) any nodes in J_1. The probability of error for the MAP detector , which uses the observations Y_{J_2} is denoted as E_2.

Lemma 4.1: ***(Error probability with dependent measurements)*** Consider MAP detection of a probabilistic hypothesis using a measurement Y. Alternatively, consider detection using $Z = f(Y) + N$, where $f(.) : R^n \to R^n$ and N is a stochastic signal independent of hypothesis. The probability of error E_Y when Y is used is no bigger than that of probability of error E_Z when Z is used, that is, $E_Y \leq E_Z$.

Lemma 4.1 has been described in standard texts on hypothesis testing [26]. Its proof is evident since Z can be computed from Y and hence the detector for Z can always be used when Y is available. This Lemma will be used to compare the performance of the two detectors.

Theorem 4.2: ***(Error probability for sensors on a network cut)*** The MAP detector error probability for the partitioning sensor set is no bigger than that for the separated sensor set, that is, $E_1 \leq E_2$.

Proof: The main idea of the proof is to show that Y_{J_2} exhibits a functional dependence on Y_{J_1}, where upon Lemma 4.1 can be applied. To show the functional dependence, first notice that for elements in J_2 that are also in J_1, the corresponding entries in $\vec{y}_{J_2}(k)$ are identical to entries in $\vec{y}_{J_1}(k)$ and hence exhibit a functional dependence. It remains to characterize the entries in $\vec{y}_{J_2}(k)$ corresponding to elements in J_2 that are not in J_1, for $k = 0, \ldots, k_f$.

To do so, consider the union of the partitions of Γ by J_1 that do not include the input location, which we label as $\widehat{J}$. Let $\widehat{x}[k]$ contain the states of the nodes in these partitions i.e., in $\widehat{J}$ at time k. We claim that $\widehat{x}[k]$, $k = 0, \ldots, k_f$ can be computed exactly from $\widehat{x}[0]$ and $\vec{y}_{J_1}[k]$, $k = 0, \ldots, k_f$. To see this, notice that when $\vec{y}_{J_1}[k]$ is known for $k = 0, \ldots, k_f$, $\widehat{x}[k]$ can be computed via a linear state equation where in $\vec{y}_{J_1}[k]$ serves as an input, without knowledge of the remaining states. Specifically, $\widehat{x}$ evolves according to

$$\widehat{x}[k+1] = \widehat{G}\widehat{x}[k] + G_a\vec{y}_{J_1}[k],$$

where, $\widehat{G}$ is the principal sub matrix of G corresponding to the nodes in $\widehat{J}$, and G_a is the sub matrix of G whose rows correspond to $\widehat{J}$ and whose columns correspond to J_1. Thus, it follows that

$$\widehat{x}[k] = \widehat{G}^k \widehat{x}[0] + \sum_{i=0}^{k-1} \widehat{G}^{k-(i+1)} G_a \widehat{y}_{J_1}[i]. \tag{13}$$

Finally, we notice that the entries in $\vec{y}_{J_2}[k]$ that do not correspond to J_1 are in $\widehat{x}[k]$. Thus, from Eq (13) it follows that $\vec{y}_{J_2}[k]$. Hence Y_{J_2} can be written as linear function of Y_{J_1}:

$$Y_{J_2} = MY_{J_1} + L\widehat{x}[0] \tag{14}$$

Since $\widehat{x}[0]$ is independent of the hypothesis, Lemma 4.1 can be applied, and $E_1 \leq E_2$. This completes the proof. ■

The above analysis indicates that, when the measurement noise is negligible, sensor sets nearer the disruption location always achieve a lower detection error probability compared to those far away from the disruption, for any specified measurement horizon. Interestingly, however, in the asymptote of a long measurement horizon, the detector performance becomes comparable for any measurement location. To see this for the mean-changed model, let us first compare the detector performance for a particular measurement location with the performance of a MAP *oracle detector* which can directly access the stochastic input sequence $w[0], \ldots, w[k_f - 1]$. The error probability of the MAP oracle detector is well-known to approach $Q(\frac{m_2 - m_1}{2\sigma}\sqrt{k_f})$ asymptotically as k_f becomes large, provided that there is some a priori probability of each hypothesis (see [26]). From the asymptotic characterization of the Gaussian cumulative distribution, the error probability of the oracle detector is asymptotically exponential in k_f, specifically given by $Ce^{\frac{-k_f(m_2-m_1)}{4\sigma}}$ for a fixed constant C.

Let us now consider MAP detection from noiseless measurements at any sensor set J. Clearly, the detection error probability cannot be less than that of the MAP oracle detector, since the measurements are a function of the stochastic input sequence. Also, the detection error probability is less than the MAP error if only one node $j \in J$ is measured. In section III, we have already established that either: 1) the detector cannot improve on the a priori error probability (for the case that $e_j^T G^z e_q$ is identically 0 for z), or 2) the detector asymptotically achieves a zero error probability.Let us refer to nodes for which the second case is in force as *valid measurement nodes*. For a valid measurement node j, let us label the smallest nonnegative integer z such $e_j^T G^z e_q \neq 0$ as $\widehat{z}_j$. From Eq (7) and assuming no measurement noise, it immediately follows that part of the stochastic input sequence, specifically $w[0], \ldots, w[k_f - \widehat{z}_j - 1]$, can be recovered exactly from Y_j. Since these samples of the stochastic input sequence can be recovered, the error probability should be no bigger than that of an *almost* oracle detector that directly uses $w[0], \ldots, w[k_f - \widehat{z}_j - 1]$ for detection. Thus, it immediately follows that the error probability when the valid node j is measured is upper bounded by $Ce^{\frac{-(k_f-\widehat{z})(m_2-m_1)}{4\sigma}}$.

Based on the above discussion we have that:

(i) For any sensor set J containing a valid measurement node, the detection error probability is asymptotically upper bounded and lower bounded by an exponential function of k_f with rate constant $\frac{m_2-m_1}{4\sigma}$.

(ii) For any sensor set J containing a valid measurement node j, the detection error probability is upper bounded by a multiple of the oracle detector's error in the asymptote. The scale factor is at most $e^{\frac{\widehat{z}_j(m_2-m_1)}{4\sigma}}$.

Remark 3: ***(Bounds on the error probability)*** From the Cayley-Hamilton theorem, it follows that $\widehat{z}_j$ is upper bounded by the dimension of the network n for any valid measurement node, and hence the ratio between the error probability for any two measurement locations is asymptotically upper bounded by $e^{\frac{n(m_2-m_1)}{4\sigma}}$. □

Remark 4: ***(Error probability for Metzler networks)*** In the special case that G is an irreducible Metzler matrix, then all nodes are valid measurement nodes. Further, $\widehat{z}_j$ is equal to the graphical distance between node j and node q in the network. Thus, the ratio between the error probability for any two measurement locations is asymptotically upper bounded by $e^{\frac{D(m_2-m_1)}{4\sigma}}$, where D is the diameter of the network. □

V. An illustrative example

In this section we discuss an illustrative example to validate our analysis. Consider a Toeplitz line network with 20 nodes and network matrix

$$G = \begin{bmatrix} a & b & 0 & \cdots & 0 & 0 \\ c & a & b & \cdots & 0 & 0 \\ 0 & c & a & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & a & b \\ 0 & 0 & 0 & \cdots & c & a \end{bmatrix}. \tag{15}$$

Let the stochastic input affect the first network node, that is, $q = 1$ in Eq (1). Let $m_1 = 0.25$ and $m_2 = 0.75$ be the two possible means of the input, with variance $\sigma_1^2 = \sigma_2^2 = 0.5$. Let the measurement noise have zero mean and variance $\sigma_n^2 = 1$. Finally, let the prior probabilities of each hypothesis be $1/2$, and the observation vector contain 100 samples. To have a better understanding of the detector performance we considered different set of weights a, b and c. Our results are illustrated in Fig. 1 and 2. It should be observed that different network weights ensure different invariant properties for the propagation of signals and detection performance. In particular, from Fig. 2, the probability of error may either decrease or increase withe the graphical distance between the input and the sensor. We regard this property as a *network invariant*, which is enforced by the network structure and weights, and that can be exploited to detect network alterations by comparing measurements at different network nodes. While this work shows the existence of network invariants that may be used for change detection, detailed characterization and study are left for future research.

VI. Conclusion

This paper studies a detection problem for networks, where changes in the statistics of an input affecting certain

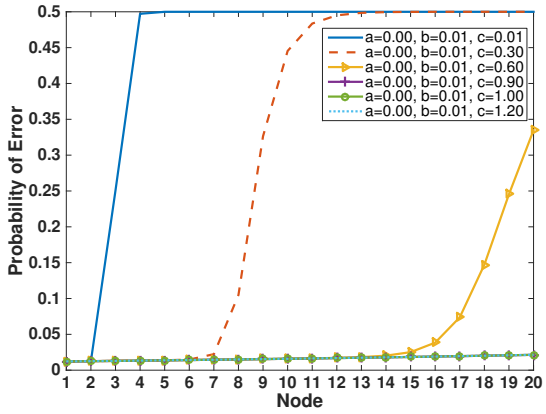

Fig. 1. In the absence of measurement noise the probability of error increases monotonically with the graphical distance between the input and the sensor, independently of the network weights. This result is consistent with Theorem 4.2, because each node is in fact a vertex cut.

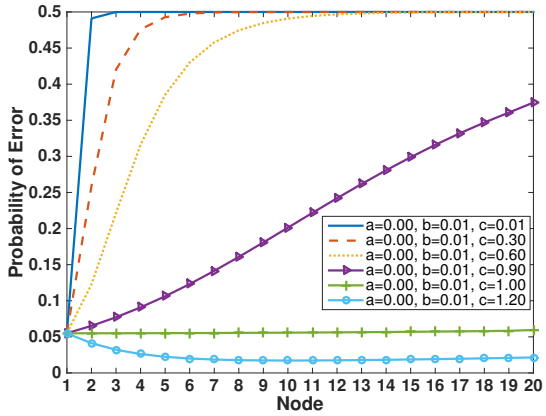

Fig. 2. In the presence of measurement noise, the probability of error may increase or decrease with the graphical distance between the input and the sensor. This property can be considered a network invariant, and is consistent with Lemma 3.1, where we see that the network dynamics and the measurement noise covariance matrix both appear in the detector.

nodes have to be detected from remotely located sensors. We provide analytic expressions for the MAP detector, and we analyze the relation between the detector performance, the location of the sensor nodes, and the network weights. We show that, in the absence of measurement noise and when sensors are positioned on a network cut, the detector performance is a nondecreasing function of the graphical distance between the input node and the sensors. Instead, in the presence of measurement noise, we show that the detector performance may improve or deteriorate when the sensor node is further away from the input node.

References

[1] Michael D Ernst, Jake Cockrell, William G Griswold, and David Notkin. Dynamically discovering likely program invariants to support program evolution. *IEEE Transactions on Software Engineering*, 27(2):99–123, 2001.

[2] AJM Spencer. Part iii. theory of invariants. *Continuum physics*, 1:239–353, 1971.

[3] Morris W Hirsch, Michael Shub, and Charles C Pugh. *Invariant manifolds*. Springer, 1977.

[4] Cyrus Chothia. Structural invariants in protein folding. *Nature*, 254(5498):304–308, 1975.

[5] Herbert A Simon. Invariants of human behavior. *Annual review of psychology*, 41(1):1–20, 1990.

[6] Mario Negrello. *Invariants of behavior: constancy and variability in neural systems*. Springer, 2011.

[7] Ashish Choudhari, Harini Ramaprasad, Tamal Paul, Jonathan W Kimball, Maciej J Zawodniok, Bruce M McMillin, and Sriram Chellappan. Stability of a cyber-physical smart grid system using cooperating invariants. In *IEEE International Computer Software and Applications Conference*, pages 760–769, 2013.

[8] Taylor T Johnson and Sayan Mitra. Invariant synthesis for verification of parameterized cyber-physical systems with applications to aerospace systems. In *AIAA Infotech at Aerospace Conference*, Boston, MA, USA, August 2013.

[9] Sayan Mitra, Tichakorn Wongpiromsarn, and Richard M Murray. Verifying cyber-physical interactions in safety-critical systems. *Security & Privacy*, 11(4):28–37, 2013.

[10] Pia M Forsman, Bryan J Vila, Robert A Short, Christopher G Mott, and Hans PA Van Dongen. Efficient driver drowsiness detection at moderate levels of drowsiness. *Accident Analysis & Prevention*, 50:341–350, 2013.

[11] A. A. Cárdenas, S. Amin, and S. S. Sastry. Research challenges for the security of control systems. In *Conference on Hot Topics in Security*, pages 1–6, Berkeley, CA, USA, 2008.

[12] S. Sridhar, A. Hahn, and M. Govindarasu. Cyber–physical system security for the electric power grid. *Proceedings of the IEEE*, 99(1):1–15, 2012.

[13] Clifford Neuman. Challenges in security for cyber-physical systems. In *DHS Workshop on Future Directions in Cyber-Physical Systems Security*, pages 22–24, 2009.

[14] Ayan Banerjee, Krishna K Venkatasubramanian, Tridib Mukherjee, and Sandeep KS Gupta. Ensuring safety, security, and sustainability of mission-critical cyber–physical systems. *Proceedings of the IEEE*, 100(1):283–299, 2012.

[15] Insup Lee, Oleg Sokolsky, Sanjian Chen, John Hatcliff, Eunkyoung Jee, BaekGyu Kim, Andrew King, Margaret Mullen-Fortino, Soojin Park, Alex Roederer, and Krishna K. Venkatasubramanian. Challenges and research directions in medical cyber–physical systems. *Proceedings of the IEEE*, 100(1):75–90, 2012.

[16] M. Zhu and S. Martínez. Stackelberg-game analysis of correlated attacks in cyber-physical systems. In *American Control Conference*, pages 4063–4068, San Francisco, CA, USA, July 2011.

[17] Sourabh Bhattacharya and Tamer Başar. Differential game-theoretic approach to a spatial jamming problem. In *Advances in Dynamic Games*, pages 245–268. Springer, 2013.

[18] Sabita Maharjan, Quanyan Zhu, Yan Zhang, Stein Gjessing, and Tamer Başar. Dependable demand response management in the smart grid: A stackelberg game approach. *IEEE Trans. Smart Grid*, 4(1):120–132, 2013.

[19] Cheng-Zong Bai, F. Pasqualetti, and V. Gupta. Security in stochastic control systems: Fundamental limitations and performance bounds. In *American Control Conference*, pages 195–200, Chicago, Il, July 2015.

[20] F. Hamza, P. Tabuada, and S. Diggavi. Secure state-estimation for dynamical systems under active adversaries. In *Allerton Conf. on Communications, Control and Computing*, September 2011.

[21] Hamza Fawzi, Paulo Tabuada, and Suhas Diggavi. Secure estimation and control for cyber-physical systems under adversarial attacks. *arXiv preprint arXiv:1205.5073*, 2012.

[22] F. Pasqualetti, F. Dörfler, and F. Bullo. Attack detection and identification in cyber-physical systems. *IEEE Transactions on Automatic Control*, 58(11):2715–2729, 2013.

[23] S. Sundaram and C. Hadjicostis. Distributed function calculation via linear iterative strategies in the presence of malicious agents. *IEEE Transactions on Automatic Control*, 56(7):1495–1508, 2011.

[24] R Dhal, J Abad Torres, and S Roy. Detecting link failures in complex network processes using remote monitoring. *Physica A: Statistical Mechanics and its Applications*, 437:36 – 54, 2015.

[25] Robert G Gallager. *Stochastic processes: theory for applications*. Cambridge University Press, 2013.

[26] Athanasios Papoulis and S Unnikrishna Pillai. *Probability, random variables, and stochastic processes*. Tata McGraw-Hill Education, 2002.

[27] Jackeline Abad Torres and Sandip Roy. Graph-theoretic analysis of network input–output processes: Zero structure and its implications on remote feedback control. *Automatica*, 61:73–79, 2015.