Distributed Detection and Isolation of Topology Attacks in Power Networks

James Weimer
ACCESS Linnaeus Centre
School of Electrical Engineering
KTH Royal Institute of Technology
Stockholm, Sweden 10044
weimerj@kth.se

Soummya Kar
Department of Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh, PA. USA 15213
soummyak@cmu.edu

Karl Henrik Johansson
ACCESS Linnaeus Centre
School of Electrical Engineering
KTH Royal Institute of Technology
Stockholm, Sweden 10044
kallej@kth.se

ABSTRACT

This paper addresses the issue of detecting and isolating topology attacks in power networks. A topology attack, unlike a data attack and power injection attack, alters the physical dynamics of the power network by removing bus interconnections. These attacks can manifest as both cyber and physical attacks. A physical topology attack occurs when a bus interconnection is physically broken, while a cyber topology attack occurs when incorrect information about the network topology is transmitted to the system estimator and incorporated as the truth. To detect topology attacks, a stochastic hypothesis testing problem is considered assuming noisy measurements are obtained by periodically sampling a dynamic process described by the networked swing equation dynamics, modified to assume stochastic power injections. A centralized approach to network topology detection and isolation is introduced as a two-part scheme consisting of topology detection followed by topology isolation, assuming a topology attack exists. To address the complexity issues arising with performing centralized detection in large-scale power networks, a decentralized approach is presented that uses only local measurements to detect the presence of a topology attack. Simulation results illustrate that both the centralized and decentralized approaches accurately detect and isolate topology attacks.

Categories and Subject Descriptors

G.3 [**Mathematics of Computing**]: Probability and Statistics

General Terms

Theory, Security

HiCoNS'12, A pril 17–18, 2012, Beijing, China.

Keywords

Distributed hypothesis testing, Distributed fault detection, Power networks

1. INTRODUCTION

As systems become more integrated into their physical environment and grow in size and complexity, the automatic detection of system faults has become increasingly important. While most of the model-based fault detection and isolation (FDI) literature focuses on centralized systems where the FDI scheme has access to all the available information [1–3], in large-scale systems centralized approaches may be prohibitively costly due to the required communication infrastructure, energy requirements, and computational complexity. For large-scale interconnected dynamical systems, such as power networks, distributed control and monitoring is more suitable [4].

Power networks are large-scale spatially distributed systems. As a critical infrastructure, they possess strict safety and reliability constraints [5]. State monitoring of the system is essential to guarantee safety, and is typically implemented in a centralized control center through a single state estimator. The core methodology for state estimation of power systems dates back to 1970, [6, 7]. Due to the low sampling frequency of the sensors in these systems a steady state approach is taken, which only allows for an over-constrained operation of the system to ensure reliability. Furthermore, faults are handled mainly by hardware devices deployed in the field, so local events leading to cascade failures may pass undetected, since the global state of the system is not taken into account. In recent years, measurement units with higher sampling rate have been developed, e.g. Phaser Measurement Units (PMU), opening the way to dynamic state estimators and observer-based fault detection schemes taking in account the dynamics of the system. Such centralized FDI schemes have been proposed in the recent literature, see [8, 9]. And more recently, extensions to distributed FDI methods has been proposed [10–13].

While there has been much research on model-based FDI within power systems for data attacks and power injection attacks, all these dynamic approaches assume an underlying model based on bus interconnections. In this work, we consider the problem of detecting topological attacks on the network occurring when bus interconnections are broke. These attacks include physical attacks, such as physically destroying an inter-bus connection, and cyber attacks on the transmission of information regarding the topological condition of the network, such as transmitting the wrong network topology configuration. The problem of detecting topological attacks is formulated for a linearized stochastic representation of the swing-equation. Both centralized and decentralized detectors are developed and evaluated using a simulated electrical power grid.

The remainder of this paper is organized as follows. In the following section we present an overview of binary hypothesis testing. Section 3 introduces the topology detection problem considered in this work. Centralized and decentralized solutions to the topology detection problem are described in Section 4 and Section 5, respectively. Section 6 provides simulation results using both 9-bus and 118-bus power networks. The final section provides a discussion and outlines future work.

2. PRELIMINARIES

This section provides a brief summary of the classical test for accepting the null hypothesis in a binary hypothesis testing problem developed by [14]. A binary hypothesis testing problem between a simple null hypothesis, H_0, and a simple alternative hypothesis, H_1, is written as

$$H_0 : \tilde{z} \sim f_0(z) \quad \text{vs.} \quad H_1 : \tilde{z} \sim f_1(z), \tag{1}$$

where $f_0(z)$ and $f_1(z)$ are the distributions of the observation random variables, $\tilde{z}$, under the null and alternative hypotheses, respectively. Given an observation, z, a test for deciding between the null and alternative hypotheses, $\phi(z) \in \{H_0, H_1\}$, is required to satisfy a performance constraint on the probability of false alarm, namely

$$P\left[\phi(z) = H_1 | H_0\right] \leq \alpha. \tag{2}$$

where $P[x|y]$ denotes the probability of x occurring conditioned on y being true. A test for rejecting the null hypothesis (accepting the alternative hypothesis), such that the performance constraint is satisfied, results from a worst case analysis of Wald's approximation (as discussed in [14]) where

$$l(z) \geq -\ln\alpha \Longrightarrow P\left[\phi(z) = H_1 | H_0\right] \leq \alpha. \tag{3}$$

and

$$l(z) = \ln f_1(z) - \ln f_0(z). \tag{4}$$

is the log-likelihood ratio. Applying Wald's approximation results in a conservative test that over-constrains the performance and is commonly employed when testing hypotheses containing complex distributions since the threshold for the resulting test is independent of the underlying distributions. The above test for rejecting H_0 will be used in this paper to develop a sequential test for detecting network topology faults.

3. PROBLEM FORMULATION

We consider an electrical power network of N interconnected buses. We assume there exists an underlying interconnection graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, between the N buses, where $\mathcal{V} \triangleq \{i\}_1^N$ is the vertex set, with $i \in \mathcal{V}$ corresponding to bus i, and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the fault-free edge set of the graph. The undirected edge $\{i, j\}$ is incident on vertices i and j if buses i and j are assumed to share an interconnection. We introduce a parameter $\theta \subseteq \mathcal{E}$, such that

$$\mathcal{N}_{i,\theta} = \{j \in \mathcal{V} : \{i, j\} \in \theta\} \tag{5}$$

defines the neighborhood of bus i assuming the interconnection specified by the edge set θ. The phase angle associated with each bus $i \in \mathcal{V}$, δ_i, has continuous-time double integrator dynamics given by the so-called swing equation [15]

$$m_i\ddot{\delta}_i(t) + d_i\dot{\delta}_i(t) = u_i(t) + w_i(t) - \sum_{j\in\mathcal{N}_{i,\Theta}} P_{ij}(t) \tag{6}$$

where $\delta_i(0) = \ddot{\delta}_i(0) = 0$ is the initial condition, m_i and d_i are the inertia and damping coefficients, respectively, P_{ij} is the active power flow from bus i to j, $w_i(t)$ is the scalar zero-mean Gaussian process noise with covariance W_i, and $u_i(t)$ is the piecewise constant mechanical input power such that for a sampling period of T_s

$$u_i(t) = u_{k,i} \quad \forall \quad T_s k \leq t \leq T_s(k+1) \tag{7}$$

Each bus is sampled periodically at the same rate of the mechanical input power using noisy sensors according to

$$\begin{aligned} y_{k,i,1} &= \delta_i(T_s k) + v_{k,i,1} \\ y_{k,i,2} &= \dot{\delta}_i(T_s k) + v_{k,i,2} \end{aligned} \tag{8}$$

where $v_{k,i,j}$ is a scalar zero-mean Gaussian measurement noise with variance $V_{i,j}$.

We assume there are no power losses nor ground admittances and let $V_i = |V_i|\, e^{j\delta_i}$ be the complex voltage of bus i such that the active power flow between bus i and bus l, P_{il}, is

$$P_{il}(t) = k_{il}\sin(\delta_i(t) - \delta_l(t)) \tag{9}$$

where $k_{il} = |V_i|\,|V_l|\, b_{il}$ and b_{il} is the susceptance of the power line connecting buses i and l. In the remainder of this work, we revert to the standard $\{i, j\}$ indexing convention where j represents an index and not the complex operator. Assuming the phase angle differences between interconnected buses are small, then

$$\sin(\delta_i(t) - \delta_j(t)) \approx \delta_i(t) - \delta_j(t) \tag{10}$$

and (6) can be written in its linear form as

$$m_i\ddot{\delta}_i(t)+d_i\dot{\delta}_i(t) = u_i(t)+w_i(t)-\sum_{j\in N_{i,\theta}} k_{ij}(\delta_i(t)-\delta_j(t)) \tag{11}$$

The linear process in (11) can be written using a continuous-time state-space formulation as

$$\dot{x}(t) = A_c(\theta)x(t) + B_c\left(u(t) + w(t)\right), \tag{12}$$

where,

$$\begin{aligned} x(t) &= \begin{bmatrix} \delta_1(t) & \cdots & \delta_N(t) & \dot{\delta}_1(t) & \cdots & \dot{\delta}_N(t) \end{bmatrix}^\top \\ u(t) &= \begin{bmatrix} u_{k,1} & \cdots & u_{k,N} \end{bmatrix}^\top \\ w(t) &= \begin{bmatrix} w_1(t) & \cdots & w_N(t) \end{bmatrix}^\top \sim N\left[0, W_c\right] \end{aligned} , \tag{13}$$

and, by defining $\mathcal{L}_\theta$ to be the graph Laplacian assuming the edges specified by θ,

$$\begin{aligned} A_c(\theta) &= \begin{bmatrix} 0_N & I_N \\ -\bar{M}\mathcal{L}_\theta & -\bar{M}\bar{D} \end{bmatrix} \\ B_c &= \begin{bmatrix} 0 \\ \bar{M} \end{bmatrix} \\ \bar{M} &= \operatorname{diag}\left(\frac{1}{m_1}, \cdots, \frac{1}{m_N}\right) \\ \bar{D} &= \operatorname{diag}\left(d_1, \cdots, d_N\right) \\ W_c &= \operatorname{diag}\left(W_1, \cdots, W_N\right) \end{aligned} \tag{14}$$

By discretizing the continuous-time state-space model in (12), a discrete-time state-space model, including the measurement model in (8), can be written as

$$\begin{aligned} x_{k+1} &= A_\theta x_k + B_\theta u_k + w_{k,\theta} \\ y_k &= x_k + v_k \end{aligned}, \tag{15}$$

where,

$$\begin{aligned} x_k &= x(T_s k) \\ u_k &= u(T_s k) \\ y_k &= \begin{bmatrix} y_{k,1,1} & \cdots & y_{k,N,1} & y_{k,1,2} & \cdots & y_{k,N,2} \end{bmatrix}^\top \\ v_k &= \begin{bmatrix} v_{k,1,1} & \cdots & v_{k,N,1} & v_{k,1,2} & \cdots & v_{k,N,2} \end{bmatrix}^\top, \\ v_k &\sim N\left[0, V\right] \\ w_{k,\theta} &\sim N\left[0, W_\theta\right] \end{aligned} \tag{16}$$

and

$$\begin{aligned} A_\theta &= \mathrm{e}^{A_c(\theta)T_s} \\ B_\theta &= \left(\int_0^{T_s} \mathrm{e}^{A_c(\theta)\tau} d\tau\right) B \\ W_\theta &= \int_0^{T_s} \mathrm{e}^{A_c(\theta)\tau} B W_c B^T \mathrm{e}^{A_c(\theta)\tau} d\tau \\ V &= \operatorname{diag}\left(V_{1,1}, \cdots, V_{N,1}, V_{1,2}, \cdots, V_{N,2}\right) \end{aligned} \tag{17}$$

For the above formulation, many researchers have considered the detection and isolation of data attacks on the input, u_k, and output, y_k [10]. However, these approaches all assume the underlying dynamical model is accurate. In this work, we consider the problem of detecting topological attacks on the network. These attacks include physical attacks on an inter-bus connection (such as physically destroying an inter-bus connection) and cyber attacks on the transmission of information regarding the topological condition of the network (such as transmitting the wrong network topology configuration). To simplify the following discussion, we assume that at most one edge has been removed from the fault-free network topology and write distribution of the sensor measurements, conditioned on the previous measurements and parameterized by $\Theta \subseteq \mathcal{E}$, as

$$f_\theta(y_k) = N\left[\mu_{k,\theta}, \Sigma_{k,\theta}\right]. \tag{18}$$

where

$$\begin{aligned} \mu_{k,\theta} &= m_{k|k-1,\theta} \\ \Sigma_{k,\theta} &= S_{k|k-1,\theta} + V \end{aligned} \tag{19}$$

and $m_{k|k-1,\theta}$ and $S_{k|k-1,\theta}$ are the *a priori* mean and *a priori* covariance, respectively, of the minimum mean-squared error state estimate, calculated recursively using a Kalman filter as

$$\begin{aligned} m_{k+1|k,\theta} &= A_\theta m_{k|k,\theta} + B_\theta u_k \\ S_{k+1|k,\theta} &= A_\theta S_{k|k,\theta} A_\theta^T + W_\theta \\ K_{k,\theta} &= S_{k|k-1,\theta}\left[S_{k|k-1,\theta} + V\right]^{-1} \\ m_{k|k,\theta} &= \left(I - K_{k,\theta}\right) m_{k|k-1,\theta} + K_{k,\theta} r_k \\ S_{k|k,\theta} &= \left(I - K_{k,\theta}\right) S_{k|k-1,\theta} \end{aligned} \tag{20}$$

To formulate a test for network topology attacks, we define $\Theta_{ij} \triangleq \mathcal{E}\backslash\{i,j\}$ to be a potential edge set and write the set of all possible edge sets as $\Theta = \{\Theta_{ij} | \{i,j\} \in \mathcal{E}\}$. Applying this notation, we consider the following M-ary hypothesis testing problem

$$\begin{aligned} H_0 &: \theta = \mathcal{E} \\ H_{1,2} &: \theta = \Theta_{1,2} \\ &\vdots \\ H_{N,N-1} &: \theta = \Theta_{N,N-1} \end{aligned} \tag{21}$$

where the hypothesis testing problem simultaneously tests the null hypothesis, H_0, which assumes no edges have been removed, against all possible alternative hypotheses, H_{ij}, representing the removal of a single edge. We note that if an edge is not contained in the fault-free network topology, $\{i,j\} \notin \mathcal{E}$, then $\mathcal{E} \equiv \Theta_{N,N-1}$ and hypothesis H_{ij} is excluded from the set of alternative hypotheses since it is statistically equivalent to the null hypothesis, H_0. To decide between the hypotheses, we introduce a test on the measurements,

$$\phi(\vec{y}_k) \in \{H_0\} \cup \{H_{ij} | \{i,j\} \in \mathcal{E}\} \tag{22}$$

where

$$\vec{y}_k \triangleq \left[y_0^T, \ldots, y_k^T\right]^T \tag{23}$$

is the time-concatenated vector of the measurements. The test $\phi(\vec{y}_k)$ can either accept the null hypothesis ($\phi(\vec{y}_k) = H_0$), or reject the null hypothesis and accept one of the alternative hypotheses ($\phi(\vec{y}_k) \in \{H_{ij} | \{i,j\} \in \mathcal{E}\}$). The decision to accept or reject the null hypothesis is made according to a design constraint on the probability of *false alarm* such that

$$P\left[\phi(\vec{y}_k) = H_{ij} | H_0\right] \leq \alpha \quad \forall \{i,j\} \in \mathcal{E} \tag{24}$$

where α is the maximum probability of false alarm. In words, we require that the probability of accepting an alternative hypothesis when the null hypothesis is correct must be less than the maximum probability of false alarm. The following sections consider the topology detection problem formulated in (21)-(24) for centralized and distributed approaches, respectively.

4. CENTRALIZED TOPOLOGICAL DETECTION

In this section we present a centralized approach to performing the hypothesis testing problem in (21) assuming the performance constraints in (24). Under this assumption, we present a two-part method for detecting and isolating network topology attacks consisting of *topology attack detection* and *topology attack isolation*. Topology attack detection is concerned with identifying whether any attack or no attack is present in the network, while topology attack isolation identifies the most likely topology attack. For topology attack detection, we write the M-ary hypothesis testing problem in (21) as a binary hypothesis testing problem between a simple null hypothesis, H_0, and a composite alternative, $\neg H_0$, as

$$H_0 : \theta = \mathcal{E} \quad \text{vs.} \quad \neg H_0 : \theta \in \Theta \tag{25}$$

Performing the test for (25) involves solving $|\mathcal{E}|$ independent binary hypothesis testing problems simultaneously, where a topology fault is detected if any of the hypothesis testing problems accept the alternative hypothesis (rejects the null hypothesis). The distribution of the measurements under each hypothesis in (25) is parameterized by the assumed network topology of the hypothesis. From (18), we observe that the hypothesis testing problem in (25) is a test on the measurements whose distributions for each parameter under each hypothesis are Gaussian with different means and different covariances. Determining a test threshold for a binary test between Gaussian distributions with different means and covariances is known to be complicated [16], thus we apply the results from Section 2 and introduce the following test for testing H_0 vs. $\neg H_0$

$$\hat{\phi}(\vec{y}_k) = \begin{cases} H_0 & \text{if} \quad \max_{\theta\in\Theta} \hat{l}_\theta(\vec{y}_k) \leq -\ln\alpha \\ \neg H_0 & \text{else} \end{cases} \tag{26}$$

where, $\hat{l}_\theta(\vec{y}_k)$ is a recursively-bounded log-likelihood ratio written as

$$\begin{aligned} \hat{l}_\theta(\vec{y}_k) =& \max\left\{0, \bar{l}_\theta(\vec{y}_k)\right\} \\ \bar{l}_\theta(\vec{y}_k) =& \hat{l}_\theta(\vec{y}_{k-1}) + \ln|\Sigma_{k,\mathcal{E}}| - \ln|\Sigma_{k,\theta}| \\ &+ (y_k - \mu_{k,\mathcal{E}})^T \Sigma_{k,\mathcal{E}}^{-1} (y_k - \mu_{k,\mathcal{E}}) \\ &- (y_k - \mu_{k,\theta})^T \Sigma_{k,\theta}^{-1} (y_k - \mu_{k,\theta}) \end{aligned} \tag{27}$$

The recursively-bounded log-likelihood ratio, $\hat{l}_\theta(\vec{y}_k)$, employs a sequential change-detection approach to test whether to accept the null hypothesis or accept the alternative hypothesis based on whether the recent measurements indicate the alternative hypothesis is more likely. If any of the binary test, $\hat{\phi}(\vec{y}_k)$, rejects the null hypothesis (accepts the alternative hypothesis), then topology isolation is performed to isolate the most likely attack. This is performed by maximizing the recursively-bounded log-likelihood ratio over the alternative hypotheses, namely

$$\max_{\theta\in\Theta} \hat{l}_\theta(\vec{y}_k) > -\ln\alpha \Longrightarrow \phi(\vec{y}_k) = H_{ij} \tag{28}$$

In words, upon rejecting the null hypothesis, the topology attack is isolated by selecting the attack that best explains the measurements. In this section, we introduced a sequential test for detecting network topology attacks using all the available measurements. In the following section we develop a distributed approach which uses only local measurements to detect topology attacks.

5. DISTRIBUTED TOPOLOGICAL DETECTION

For large-scale systems, such as power transmission networks, communication and computational constraints can prohibit the real-time implementation of a centralized detection and isolation scheme. In this section, we present a distributed approach for detecting and isolating topology attacks based strictly on local information. To formulate the distributed topology detection scheme, we introduce

$$y_{k,\Theta_{ij}} = Q_{\Theta_{ij}} y_k \tag{29}$$

where $Q_{\Theta_{ij}}$ is the binary selection matrix with full row rank and columns which contain a unit entry if and only if the corresponding element of y_k is contained in the combined neighborhood of the vertices of the edge $\{i, j\}$. For each $\theta \subset \mathcal{E}$, the distribution of the local measurements, independent of the previously local measurements, is written as

$$f_\theta(y_{k,\theta}) = N\left[\hat{\mu}_{k,\theta}, \hat{\Sigma}_{k,\theta}\right]. \tag{30}$$

where

$$\begin{aligned} \hat{\mu}_{k,\theta} &= Q_\theta \hat{m}_{k|k-1,\theta} \\ \Sigma_{k,\theta} &= Q_\theta \left(\hat{S}_{k|k-1,\theta} + V\right) Q_\theta^T \end{aligned} \tag{31}$$

and $\hat{m}_{k|k-1,\theta}$ and $\hat{S}_{k|k-1,\theta}$ are the *a priori* state mean and covariance, respectively, and are calculated recursively using a Kalman filter based on only local measurements as

$$\begin{aligned} \hat{m}_{k+1|k,\theta} =& A_\theta \hat{m}_{k|k,\theta} + B_\theta u_k \\ \hat{S}_{k+1|k,\theta} =& A_\theta \hat{S}_{k|k,\theta} A_\theta^T + W_\theta \\ \hat{K}_{k,\theta} =& \hat{S}_{k|k-1,\theta} Q_\theta^T \left[Q_\theta \left(\hat{S}_{k|k-1,\theta} + V\right) Q_\theta^T\right]^{-1} \\ \hat{m}_{k|k,\theta} =& \left(I - \hat{K}_{k,\theta} Q_\theta\right) \hat{m}_{k|k-1,\theta} + \hat{K}_{k,\theta} y_{k,\theta} \\ \hat{S}_{k|k,\theta} =& \left(I - \hat{K}_{k,\theta} Q_\theta\right) \hat{S}_{k|k-1,\theta} \end{aligned} \tag{32}$$

Additionally, we define the time concatenation of the local measurements as

$$\vec{y}_{k,\theta} \triangleq \left[y_{0,\theta}^T, \ldots, y_{k,\theta}^T\right]^T. \tag{33}$$

To distribute topology attack detection as presented in the previous section, we implement the composite binary hypothesis testing problem in (25) using a collection of distributed simple binary hypothesis testing problems between the null hypothesis, H_0, and a unique alternative hypothesis, H_{ij}, each of which is implemented using only local measurements, $\vec{y}_{k,\theta}$, and written as

$$H_0 : \theta = \mathcal{E} \quad \text{vs.} \quad H_{ij} : \theta = \Theta_{ij} \tag{34}$$

In the distributed scheme, there are $|\mathcal{E}|$ independent binary hypothesis testing problems performed simultaneously and a topology fault is detected if any of the hypothesis testing problems accept the alternative hypothesis (rejects the null hypothesis). Following the same reasoning as in the centralized scheme, we introduce a test for the hypothesis testing problem in (34) as

$$\hat{\phi}_{\Theta_{ij}}(\vec{y}_{k,\Theta ij}) = \begin{cases} H_0 & \text{if} \quad \hat{l}_{\Theta_{ij}}(\vec{y}_{k,\Theta ij}) \leq -\ln\alpha \\ H_{ij} & \text{else} \end{cases} \tag{35}$$

Upon detecting a topology attack using local measurements, the corresponding recursive log-likelihood ratio value is transmitted to the central estimator and compared with the other (if any) recursive log-likelihood ratios associated with any other binary tests that reject the null hypothesis. As a heuristic, the topology isolation is performed according to

$$\max_{\Theta_{ij}\in\Theta} \hat{l}_{\Theta ij}(\vec{y}_{k,\Theta_{ij}}) > -\ln\alpha \Longrightarrow \phi(\vec{y}_k) = H_{ij} \tag{36}$$

In words, the topology attack is isolated selecting the attack that best explains the local measurements. The complexity involved in computing the distributed statistics is significantly less since the matrix inverse involved in performing the Kalman filter is proportional to the number of sensor measurement. Thus, using only local sensor measurements in the distributed case results in significant computational savings. In this section, we introduced a distributed sequential test for detecting network topology attacks using only local measurements.

6. SIMULATION RESULTS

In this section we evaluate the centralized and decentralized approaches to topology detection using both 9-bus and 118-bus power networks. For comparison between the centralized and decentralized performance, this section first considers a 9-bus power network as defined by the MATPOWER toolbox [17]. The connectivity graph of the 9-bus system is illustrated in Fig. 1. The power injections and

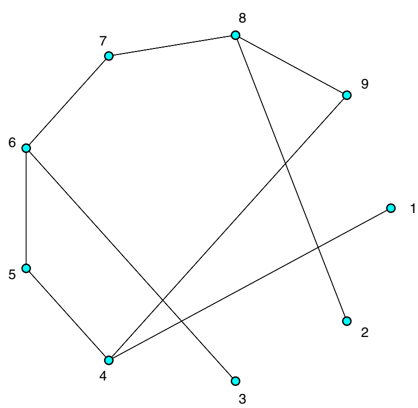

Figure 1: Connectivity graph of 9-bus power network

loading are assumed to have an expected value equivalent to the DC steady-state solution as defined by the MATPOWER toolbox, namely

$$u(t) = \begin{bmatrix} 67 & 163 & 85 & 0 & -90 & 0 & -100 & 0 & -125 \end{bmatrix}^\top \tag{37}$$

with process noise covariance, W, and measurement noise covariance, V, defined as

$$W = 10I \quad \text{and} \quad V = 0.1I \tag{38}$$

We assume a sampling period of 0.02 seconds, which is comparable with current PMU measurement technologies and assume the inertia and damping coefficients, $m = \{m_1, \ldots, m_9\}$ and $d = \{d_1, \ldots, d_9\}$, respectively, are

$$\begin{aligned} m &= \left\{ \begin{array}{r} 8.9304, 8.8462, 8.5269, 0.0008, 0.0008, \\ 0.0008, 0.0008, 0.0009, 0.0009 \end{array} \right\} \\ d &= \left\{ \begin{array}{r} 3.0920, 3.6539, 3.4160, 0.0004, 0.0004, \\ 0.0003, 0.0004, 0.0004, 0.0003 \end{array} \right\} \end{aligned} \tag{39}$$

where the first three buses contain generators which have significantly greater inertia than the non-generation buses as captured by their respective inertia and damping coefficient magnitudes [15]. The voltage phase angle trajectory assuming edge $\{4, 5\}$ is removed at time $t = 0.5s$ is shown in Fig. 2 From the state trajectory we observe that the phase

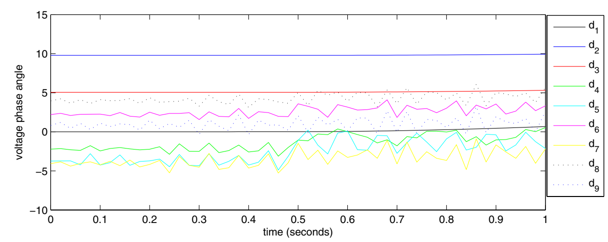

Figure 2: 9-bus voltage phase trajectories (in degrees) versus time (in seconds).

change caused by removing the edge between bus 4 and bus 5 does not result in a safety critical situation since the DC voltage phase angle between buses remains small (between 45 degrees). We specifically consider this situation when evaluating detection and isolation capabilities since drastic changes in phase are much easier to detect and isolate.

For testing purposes, we assume a maximum probability of false alarm as $\alpha = 10^{-4}$ and plot the time evolution of the test statistics as calculated using the centralized approach (assuming all measurements) and the decentralized approach (assuming only local measurements) against the decision threshold in Fig. 3. In Fig. 3, the centralized test statistic trajectories are displayed in the upper subplot while the decentralized test statistic trajectories are shown in the lower subplot. For both plots, the statistic corresponding the actual error (removal of the edge between bus 4 and bus 5) is denoted by the solid black line, while the dotted and dashed lines correspond to all other test statistic trajectories.

We observe from the results in Fig. 3 that detection in both the centralized and decentralized approaches occurs within one sample of the error, indicating detection of a topology attack is quick and accurate. In terms of isolation, we observe that in the centralized case, isolation is performed correctly and remains correct throughout the entire statistic trajectory. In the distributed case, isolation is initially correct, but occasionally leads to incorrect isolation as the statistics evolve. We assert that accurate initial detection and isolation is most important since action would be taken immediately to remedy the situation. Moreover, once the voltage phase angle reaches a steady state value, DC analysis on the phase change can be applied to identify the topology change. Thus, we are primarily concerned with the transient detection and isolation of topology attacks, and observe that both the centralized and decentralized approaches yield promising results.

In terms of the comparative performance between the decentralized and centralized approaches, we observe that the centralized approach is more accurate than the decentralized, which is expected since the centralized approach uses much more information to both detect and isolate. It is clear in Fig. 3 that the decentralized approach is prone to a higher probability of false alarm than the centralized approach as

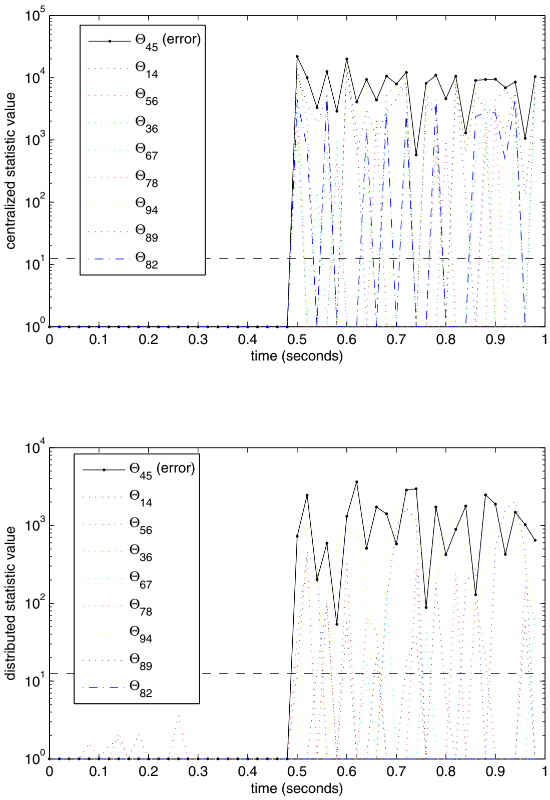

Figure 3: 9-bus power network : centralized test statistics (top) and decentralized test statistics (bottom) vs. time step.

indicated by the peaks in the statistic trajectories in the decentralized approach during the period when no errors are present (time zero to time 0.5 seconds). Although both approaches use the same threshold (which ensures the maximum probability of false alarm is bounded by α), the difference in their actual probability of false alarm is attributed to the conservative nature of Wald's approximation.

To evaluate performance of the distributed approach in a large-scale power network, we consider consider a 118-bus network as defined by the MATPOWER toolbox [17]. In this simulation, and similar to the 9-bus network, we assume a DC-operating point as defined by the MATPOWER toolbox while measuring and perturbing the system using the same noise profiles (but of larger dimension). In this simulation, we assume a topology attack on the edge between bus 5 and bus 6. Figure 4 indicates the trajectory of the statistic corresponding to the actual topology attack (solid black line) versus the maximum of all other statistic trajectories corresponding to all other topology attacks (dotted red line). In Fig. 4, we observe that the distributed approach to detection and isolation accurately detects and isolates the attack quickly. Moreover, the computational complexity of performing decentralized testing is greatly reduced since only local sensor measurements are selected for inclusion in the local tests.

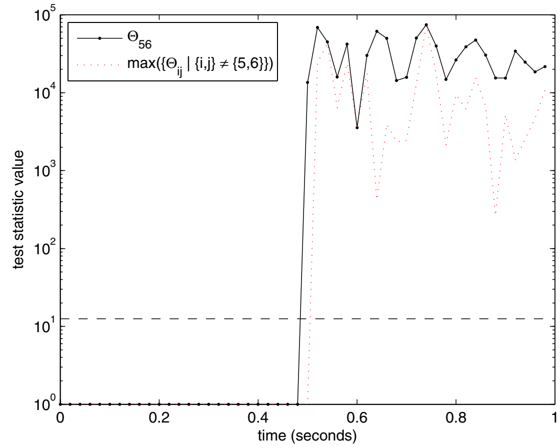

Figure 4: 118-bus power network: decentralized test statistics vs. time step.

7. DISCUSSION AND FUTURE WORK

This work considers the detection of topology attacks in second-order network systems, with specific application to power networks. A DC-power analysis is employed to generate the well known swing-equation dynamics that when linearized is dependent upon the power network connectivity graph. A decentralized approach to detection and isolation is proposed that is evaluated using the state-of-the-art MATPOWER simulation toolbox. Results indicate the the decentralized approach is well suited for large-scale detection and isolation in networked systems.

Future work on this topic include applying principles of invariance to develop test statistics that evolve independent of non-hypothesized topology attacks. These approaches have shown promise in power networks when isolating power injection attacks; however, topology attacks present much more difficult scenarios when discrete sensing is assumed since plant discretization results in systems that are heavily dependent on the assumed network laplacian. Additionally, we envision future distributed detection schemes to incorporate collaboration between different agents (tests) where agents iteratively exchange information with each other over a communication graph (possibly different from the physical graph). These approaches have been shown to yield promising results in distributed gossip based approaches for collaborative hypothesis testing [18].

8. ACKNOWLEDGMENTS

This work is supported by the Swedish Energy Agency, the Swedish Governmental Agency for Innovation Systems (VINNOVA), the Swedish Foundation for Strategic Research (SSF), and the Knut and Alice Wallenberg Foundation.

9. REFERENCES

[1] R. Isermann, "Model-based fault detection and diagnosis: status and applications," in *Proceedings of the 16th IFAC Symposium on Automatic Control in*

Aerospace, St. Petersburg, Russia, June 2004, pp. 71–85.

[2] S. X. Ding, *Model-based Fault Diagnosis Techniques: Design Schemes*. Springer Verlag, 2008.

[3] J. Chen and R. J. Patton, *Robust Model-Based Fault Diagnosis for Dynamic Systems*. Kluwer Academic Publishers, 1999.

[4] D. D. Siljak, *Decentralized control of complex systems*. Academic Press, 1991.

[5] M. Shahidehpour, W. F. Tinney, and Y. Fu, "Impact of security on power systems operation," *Proceedings of the IEEE*, vol. 93, no. 11, pp. 2013–2025, Nov. 2005.

[6] F. C. Schweppe and J. Wildes, "Power system static-state estimation, part I: Exact model," *IEEE Transactions on Power Apparatus and Systems*, vol. 89, no. 1, pp. 120–125, January 1970.

[7] A. Abur and A. Exposito, *Power System State Estimation: Theory and Implementation*. Marcel-Dekker, 2004.

[8] E. Scholtz and B. Lesieutre, "Graphical observer design suitable for large-scale DAE power systems," in *Proceedings of the IEEE Conf. on Decision and Control*, Cancun, Dec. 2008, pp. 2955–2960.

[9] M. Aldeen and F. Crusca, "Observer-based fault detection and identification scheme for power systems," in *IEE Proceedings - Generation, Transmission and Distribution*, vol. 153, no. 1, Jan. 2006, pp. 71–79.

[10] I. Shames, A. M. H. Teixeira, H. Sandberg, and K. H. Johansson, "Distributed fault detection for interconnected second-order systems with applications to power networks," in *IN FIRST WORKSHOP ON SECURE CONTROL SYSTEMS*, 2010.

[11] S. X. Ding, P. Zhang, C. Chihaia, W. Li, Y. Wang, and E. L. Ding, "Advanced design scheme for fault tolerant distributed networked control systems," in *Proceedings of the 17th IFAC World Congress*, Seoul, Korea, July 2008, pp. 13 569 – 13 574.

[12] W. H. Chung, J. L. Speyer, and R. H. Chen, "A decentralized fault detection filter," *Journal of Dynamic Systems, Measurement, and Control*, vol. 123, no. 2, pp. 237–247, 2001. [Online]. Available: http://link.aip.org/link/?JDS/123/237/1

[13] F. Pasqualetti, A. Bicchi, and F. Bullo, "Consensus computation in unreliable networks: A system theoretic approach," *IEEE Transactions on Automatic Control*, 2010, submitted, available online at http://www.fabiopas.it/papers/FP-AB-FB-10a.pdf.

[14] A. Wald, *Sequential Analysis*. John Wiley & Sons, Inc., New York, 1947.

[15] P. Kundur, *Power System Stability and Control*. McGraw-Hill Professional, 1994.

[16] L. L. Scharf, *Statistical Signal Processing, Detection, Estimation, and Time Series Analysis*. Addison-Welsley Publishing Company Inc., Reading, Massachusetts, 1991.

[17] R. D. Zimmerman, Carlos, and D. . Gan, "MATPOWER: A MATLAB Power System Simulation Package, Version 3.1b2, User's Manual," Power Systems Engineering Research Center, Tech. Rep., 2006. [Online]. Available: http://www.pserc.cornell.edu/matpower/

[18] S. Kar, S. Aldosari, and J. M. F. Moura, "Topology for distributed inference on graphs," *Jun. 2006 [Online]. Available: http://www.arxiv. org/abs/cs/0606052*.