Physica A 437 (2015) 36–54

Contents lists available at ScienceDirect

Physica A

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

Detecting link failures in complex network processes using remote monitoring

R. Dhal[a,*], J. Abad Torres[b], S. Roy[c,*]

[a] EPIS Inc., 13535 72nd Ave., Suite 165, Tigard, OR 97223, United States
[b] Escuela Politécnica Nacional, Ladrón de Guevara E11253, PO. Box 17-01-2759, Quito, Ecuador
[c] EME 402, PO BOX 642752, Washington State University, Pullman, WA, 99164-2752, United States

HIGHLIGHTS

- We study detection of network changes from remote noisy time-series measurements.
- A Maximum A-Posteriori Probability hypothesis testing scheme is employed.
- Relationships between the network topology and MAP detector performance are developed.
- Detector performance depends on presence of certain paths in the network.
- Simulations demonstrate the analytical results developed.

ARTICLE INFO

Article history:
Received 12 September 2014
Received in revised form 29 March 2015
Available online 27 May 2015

Keywords:
Estimation in networks
Decision/estimation theory
Fault detection/diagnosis
Synchronization

ABSTRACT

We study whether local structural changes in a complex network can be distinguished from passive remote time-course measurements of the network's dynamics. Specifically the detection of link failures in a network synchronization process from noisy measurements at a single network component is considered. By phrasing the detection task as a Maximum *A Posteriori* Probability hypothesis testing problem, we are able to obtain conditions under which the detection is (1) improved over the *a priori* and (2) asymptotically perfect, in terms of the network spectrum and graph. We find that, in the case where the detector has knowledge of the network's state, perfect detection is possible under general connectivity conditions regardless of the measurement location. When the detector does not have state knowledge, a remote signature permits improved but not perfect detection, under the same connectivity conditions. At its essence, detectability is achieved because of the close connection between a network's topology, its eigenvalues and local response characteristics.

1. Introduction

Network synchronization models – which describe coordination or equalization of network components' states via local interactions – are descriptive of both interesting physical-world processes and distributed consensus/agreement algorithms [1–4]. These synchronization models have proved influential in several disciplines, including in characterizing emergent behaviors in nature [5], analyzing traditional engineered networks (e.g., power, transportation and Internet [6,7]), and designing algorithms for distributed cyber-systems such as wireless sensor networks [8].

---

[*] Corresponding authors.
E-mail addresses: rahuldhal@epis.com (R. Dhal), jabadtor@eecs.wsu.edu (J. Abad Torres), sroy@eecs.wsu.edu (S. Roy).

http://dx.doi.org/10.1016/j.physa.2015.04.024

Many physical-world and cyber-network synchronization processes operate in complex environments, which can effect local topological changes. These local changes can induce network-wide propagative impacts that modify the network's core function, and hence there is a motivation for quickly detecting topological changes. Here, we study whether local changes in a synchronization network's graph can be detected using only passive monitoring at a few network locations.

This study contributes to a growing research effort on inferring complex network topologies (and/or topological changes), either from partial knowledge of the topology or by leveraging a dynamics imposed on the network. One approach taken in the literature is to impose a dedicated algorithm for identifying particular topological changes, e.g. cuts or local damage [9–11]. Entirely structure-based approaches have also been widely studied, such as data mining techniques to reconstruct a full network topology from partial knowledge [12]. A third approach, which is more closely connected to our focus here, is to use passive measurements of the network's dynamics to infer modal or topological characteristics [13–18]. These techniques fundamentally leverage the network's native dynamical responses for inference. While our study is aligned with this perspective, we are particularly concerned with detecting local structural changes (specifically link failures) using measurements taken only at a very few network locations, which requires us to exploit remote signatures. In particular, this study examines whether a monitor can infer link failures in a discrete-time network synchronization process from local noisy temporal response data. This type of problem arises e.g. in infrastructure-control and security applications, where stakeholders may only be able to probe the network at a few points, and depend on the propagative impact of topological changes for detection.

The approach taken here is to phrase the problem of inferring link failures as a *Maximum A Posteriori Probability* (MAP) detection problem, in which the monitor seeks to distinguish between alternative failure hypotheses using its noisy local response data, with or without knowledge of the network's dynamical state. The key contribution of the work is to characterize the performance of the MAP detector in terms of the graph topology of the synchronization process, and the locations of the failure and measurement relative to the graph. Our analysis shows that propagative responses caused by local structural changes allow the monitor to distinguish failures from remote measurement signatures, under broad network-connectivity conditions. In the case where the observer has knowledge of the state, perfect detection is possible asymptotically as long as the network is strongly connected. In the unknown-state case, perfect inference is typically not possible, but the measurement data permits effective inference as long as the measurement noise is limited.

Our analysis also shows that detectability, and the time required for detection, fundamentally depends on how the network spectrum changes due to the failure of links. These dependences are related to, and informed by, the wide literature connecting a network's spectrum with its graph structure (e.g. Ref. [19]). Of particular note, the detection task can be viewed as an extension of the inverse problem phrased as *Can one hear the shape of a network?*, which asks if the graph eigenvalue uniquely identifies the underlying topology [20]. Here we take the further step of using sparse local measurements of a network's dynamics to identify a change in the topology, which is seen to be possible precisely if the spectrum changes. The spectral analysis of detectability also admits an interesting interpretation from the perspective of Ref. [21], which shows that networks with the same Laplacian eigenvalues can have substantially different synchronization behavior.

The rest of the article is organized as follows. In Section 2 we introduce the link-failure detection problem, the MAP detection methodology, and metrics for quantifying detection performance. We differentiate our analysis on the MAP scheme into two cases based on the monitor's knowledge of the initial state. In Section 3 the performance of the MAP scheme is analyzed when the initial state is exactly known. Section 4 focuses on the case where the initial state is unknown and modeled a random variable. Finally, illustrative examples are given in Section 5. Our main focus throughout is to develop graph-theoretic insights into the performance of the MAP detector.

2. Problem description

In this section, we describe the network-synchronization model, pose the problem of detecting (inferring) link failures, and introduce the MAP detection approach.

2.1. Nominal networked synchronization process

Nominally, we consider a synchronization dynamics defined on a weighted, directed graph (or digraph) $\Gamma = (\mathcal{V}, E : W)$. Here, $\mathcal{V}$ is a set containing the n vertices in the graph, which we label $1, 2, \ldots, n$. E is a set of directed edges or arcs, each of which is an ordered pair of distinct vertices (say (i, j), for the edge from vertex i to vertex j). Each edge (i, j) in the graph has associated with it a positive weight w_{ij}, as specified in the weight set W. We concisely represent an arc (directed edge) and its associated weight as $(i, j : w_{ij})$. We refer to i and j as the starting and trailing end of the directed edge (i, j).

A nominal synchronization process is defined on Γ. Specifically, a process with n components is considered, where component i corresponds to vertex i in Γ. Edges in Γ correspond to links in the network synchronization process. Each network component k has associated with it a *synchronization state* $x_k[t]$, which evolves along a discrete time axis $t \in Z^+$. At each time step, component k's state evolves based on a weighted average of differences from other components' states, where the weights are specified by the directed edges into vertex k in Γ. Specifically, $x_k[t]$ is nominally governed by:

$$x_k[t+1] = x_k[t] + \sum_{i \in \mathbb{N}_k} w_{ik}(x_i[t] - x_k[t]) \tag{1}$$