EFFECTS OF NETWORK TOPOLOGY ON THE CONDITIONAL DISTRIBUTIONS OF SURROGATED GENERALIZED COHERENCE ESTIMATES

Lauren Crider and Douglas Cochran

School of Mathematical & Statistical Sciences
Arizona State University; Tempe AZ 85287-1804 USA

ABSTRACT

Coherence estimation is an established approach in multiple-channel detection and estimation, providing optimal solutions in many cases. Recent work has considered the use of maximum-entropy matrix completion when elements are missing from the gram matrix from which the coherence statistics are formed. This is desirable in sensor network settings, for example, where direct communication is not available between every pair of nodes in the network. This paper examines the role of network topology in determining the conditional distributions of the statistic obtained by the matrix completion process under both signal-present and signal-absent hypotheses.

1. INTRODUCTION

Coherence estimation is an established approach in multi-channel detection and estimation. Detectors and estimators based on this approach are typically functions of the eigenvalues of a gram matrix formed from segments of time series data collected at the sensors [8, 6, 12, 10, 9, 11, 13]. In some active sensing settings, a signal replica may be used in addition to sensor data in forming the gram matrix [1]. Recent work has also used multiple segments of time series from the same sensor in coherence-based processing [12, 2].

The generalized coherence (GC) estimate [8, 6], which is based on the gram matrix determinant, is the statistic for an optimal test of diagonal covariance versus arbitrary non-diagonal covariance across the channels. It has been used, in particular, for detecting the presence of a common but unknown signal in noisy data collected at a suite of sensors. While this paper focuses on the GC estimate, the approach for handling scenarios involving incompletely connected sensor networks introduced in [7] and analyzed here also applies to other statistics based on the gram matrices.

As discussed more precisely in Section 2, the elements in the gram matrix are inner products (correlations) between normalized segments of time series data collected at each pair of sensors in the network. When two nodes are in direct communication (i.e., share an edge in the network graph), their corresponding inner product can be computed locally. Thus in the case of a completely connected network, only scalar values must be accumulated globally to form the gram matrix. This research is primarily motivated by incompletely connected sensor networks, wherein local computation of some inner products is not possible. The approach set forth in [7] replaces missing inner product data with "surrogate" values obtained via a maximum-entropy method, in essence performing maximum-entropy completion of the gram matrix. The loss in detection performance compared to having the full gram matrix was observed in simulations to be modest when the number of surrogated entries was small.

This paper examines the statistical behavior of the GC estimate formed from a gram matrix obtained by completing a partial gram matrix according to maximum-entropy criterion in the absence of a common signal. The focus is on two questions:

1. Under the signal-absent hypothesis ($\mathcal{H}_0$), does the distribution of the GC estimate depend on *which* elements of the gram matrix have been surrogated, or only on how many?

2. Under the signal-present hypothesis ($\mathcal{H}_1$), does the distribution of the GC estimate depend on which elements of the gram matrix have been surrogated?

Item 1 is of particular importance in setting detection thresholds that correspond to desired false alarm probabilities at the fusion center. If the signal-absent distribution does not depend explicitly on the network topology, it will only be necessary for the fusion center to store one set of detection thresholds for each *number* of possible surrogated values rather than one set for each possible network topology, of which there are a huge quantity for even a ten-node network.

Section 2 introduces the problem formulation and establishes the GC detector. Section 3 describes the approach taken for applying coherence-based detectors in examines in incomplete networks, where surrogate values for channels that are not in direct communication are provided by maximum-entropy completion of the gram matrix. Simulation results showing the performance of this method in specific scenarios are presented in Section 4 and an invariance property that is the main result of this paper is presented in Section 5. A short discussion closes the paper.

2. GENERALIZED COHERENCE DETECTION

Consider a set of M complex N-vectors, $x_1, \cdots, x_M$ that represent time series data from each of the M spatially distributed sensor channels. It is assumed, as is normally true in practice, that $N \gg M$. The GC estimate for $M \geq 2$ sensor channels is defined as

$$\hat{\gamma}^2(x_1, \ldots, x_M) = 1 - \frac{\det G(x_1, \ldots, x_M)}{||x_1||^2 \cdots ||x_M||^2}$$

where

$$G(x_1, \ldots, x_M) = \begin{bmatrix} \langle x_1, x_1 \rangle & \cdots & \langle x_1, x_M \rangle \\ \vdots & \ddots & \vdots \\ \langle x_M, x_1 \rangle & \cdots & \langle x_M, x_M \rangle \end{bmatrix}$$

is the $M \times M$ positive semi-definite hermitian matrix whose entries are comprised of the inner products $\langle x_i, x_j \rangle = x_j^\dagger x_i$, where $\dagger$ denotes conjugate transpose. Denoting X the $N \times M$ matrix whose m^{th} column is x_m, the gram matrix can be written as $G = X^\dagger X$. In typical multiple-channel detection applications, the value of $\hat{\gamma}^2$ is compared to a threshold to decide between signal-absent ($\mathcal{H}_0$) and signal-present ($\mathcal{H}_1$) hypotheses, defined as follows

$$\mathcal{H}_0 : x_k = \nu_k, \qquad k = 1, \ldots, M$$

$$\mathcal{H}_1 : x_k = \alpha_k \mathbf{s} + \nu_k, \qquad k = 1, \ldots, M$$

where $s \in \mathbb{C}^N$ is the common, but unknown, deterministic signal and the random vectors ν_k are independent, complex, zero-mean, white gaussian noise, i.e. $\nu_k \sim \mathcal{CN}[\mathbf{0}, \sigma^2 \mathbb{I}_{N \times N}]$ with $\mathbf{0}$ the zero vector in $\mathbb{C}^N$ and $\mathbb{I}$ the $N \times N$ identity matrix. Each α_k represents an unknown complex gain.

Consider the four-node sensor network ($M = 4$) depicted in Figure 1(a). In this diagram, each pair of elements connected by an edge in the network graph are assumed to be in direct communication, and consequently able to form the inner product of their data locally and transmit only a scalar value to the fusion center for processing. So in the network configuration of Figure 1(a) the fusion center receives all necessary information to implement the GC detector. However, when the network graph is not complete, as in the network configuration of Figure 1(b), inner products of data vectors corresponding to each edge in the network graph is not sufficient to enable the fusion center to compute the GC estimate.

The following section describes an approach for completing the gram matrix and performing an approximate GC estimate when the network is incompletely connected.

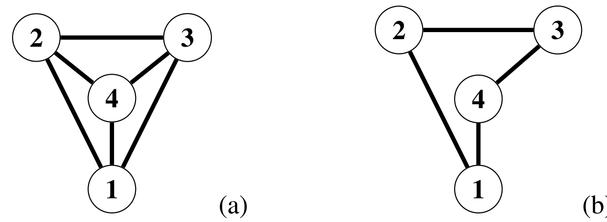

Fig. 1. (a) A complete sensor network with four nodes. (b) An incomplete four-node network.

3. MAXIMUM-ENTROPY SURROGATION IN GC-BASED DETECTION

The problem of completing a partial gram matrix corresponding to sensor network configurations, as in the case depicted in Figure 1(b), has been addressed in recent work. A method for forming an approximate GC estimate was recently introduced in [7]. The approach outlined in this paper uses the principle of maximum entropy to introduce maximum-entropy "surrogate" values for the missing inner product entries of the gram matrix.

Briefly, the method of maximum-entropy assumes that there is a complex random variable x_m modeling the time series data collected at each sensor. The ability to communicate between sensors k and j linked in the network graph gives an estimation of the covariance $\text{cov}(x_k, x_j)$. Thus for a complete sensor network, it is possible to estimate the full $M \times M$ covariance matrix C of the variables $x_1, \ldots, x_M$. Under the suitable assumption that these variables are independent with mean zero, the optimal estimator $\hat{C}$ is proportional to the gram matrix $G(x_1, \ldots, x_M)$. The method of maximum-entropy holds that missing values in C should be surrogated in such a way as to introduce no new assumptions about the nature of the random variables or of the network itself. The joint distribution of the random variables $x_1, \ldots, x_M$ that best describes the current knowledge (i.e. the estimated covariances between all pairs of directly connected sensors) is the maximum-entropy distribution constrained by the available data. The problem of finding the maximum-entropy completion of a covariance matrix has been studied in prior literature, where it is noted that the covariance matrix of this maximum-entropy distribution will have the property that its inverse will have zeros in positions corresponding to the missing covariance values [14]. This method gives a direct means for computing the necessary surrogate values for a partial gram matrix in small network regimes. It is noted that computing k surrogate values via the method of maximum-entropy becomes formidable for $M > 5$. Fortunately, such maximum-entropy covariance matrix completion problems can be efficiently solved by convex programming techniques [4].

4. DETECTION PERFORMANCE IN INCOMPLETELY CONNECTED NETWORKS

Consider the four-sensor incomplete network ($M = 4$) with two surrogated links, as pictured in Figure 2. This is the smallest setting for which there are two distinct network topologies.

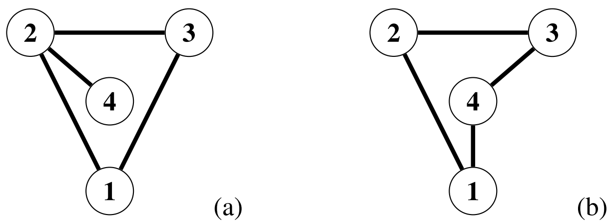

Fig. 2. The two distinct topologies for a four-node network with four edges.

ROC curves obtained from a 1,000-trial Monte Carlo simulation with equal SNR on all sensor channels are shown in Figure 3 below. The scenario is detection of a white complex gaussian signal vector in white complex gaussian noise on $M = 4$ channels. The vectors are of length $N = 128$. The top set of curves is the detection performance with SNR -7.5 dB, the center has SNR -9 dB, and the bottom set of curves is the detection performance with SNR -12 dB. Within each set, the solid curve is the detection performance assuming all data are available and aggregated at the fusion center, while the dashed lines are for the incomplete networks with maximum-entropy completion of the gram matrix.

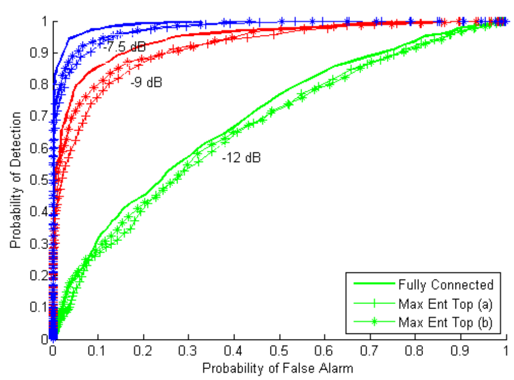

Fig. 3. ROC curves for detection in complete and surrogated four-node networks at three different SNRs.

The loss in detection performance between the fully connected network and the incompletely connected networks in the two topologies compared is modest, in agreement with [3, 7]. More importantly for this paper, examination of the detection performance for the two incomplete topologies (each with two surrogated values) suggests that performance may be identical.

Further empirical tests for various network sizes and topologies show that the different networks of M sensors with equal numbers of surrogate values give indistinguishable detection performance. Figure 4 illustrates ROC curves obtained from a 1,000-trial Monte Carlo simulation with equal SNR on all sensor channels. The scenario is detection of a white complex gaussian signal vector in while complex gaussian noise on $M = 20$ channels. The vectors are of length $N = 128$. Each curve corresponds to a distinct randomly generated topology with $k = 5$ surrogate values. We remark that the number of distinct network topologies for a 20-node sensor network with 5 surrogations is on the order of 2×10^{34}.

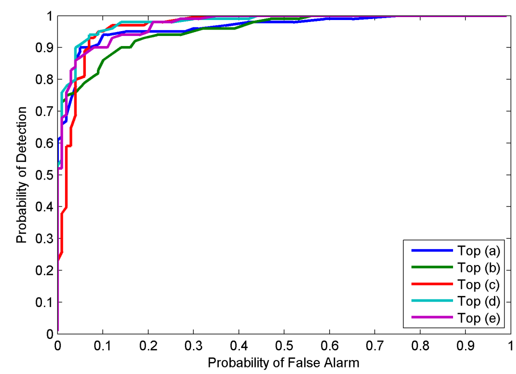

Fig. 4. ROC curves for detection in complete and surrogated twenty-node networks at three different SNRs. Maximum-entropy matrix completion in this case was accomplished numerically via convex optimization.

5. NULL DISTRIBUTION OF SURROGATED COHERENCE

In the context of multiple-channel signal detection, the value of knowing the conditional distribution under the signal-absent hypothesis is its utility in determining detection thresholds that correspond to desired false alarm probabilities. Past work [5, 6] has studied the conditional distributions under the signal-present and signal-absent hypotheses. In fact, the statistical behavior of the GC estimate is well understood in the absence of a common signal, thereby allowing such thresholds to be readily established. Recording the distribution of the surrogated $\hat{\gamma}^2$ value under signal-absent conditions with 1,000 independent trials provides empirical cumulative distribution functions, shown in Figure 5(a), for the two distinct four-sensor network topologies depicted in Figure 2. By standard goodness of fit tests, these distributions are indistinguishable.

The invariance to network topology observed in the surrogated null distribution is not expected to hold under the signal-present hypothesis, particularly when channels carry signal of differing strengths. Without investigating any detail, this conclusion is verified in one example as follows. Reconsider the four node sensor network in Figure 2(a). Since sensors 1 and 3 do not receive any information from sensor 4, if 4 carries a stronger signal than do sensors 1, 2, and 3, detection performance degradation is expected. Figure 5(b) illustrates sample distribution functions under such a model.

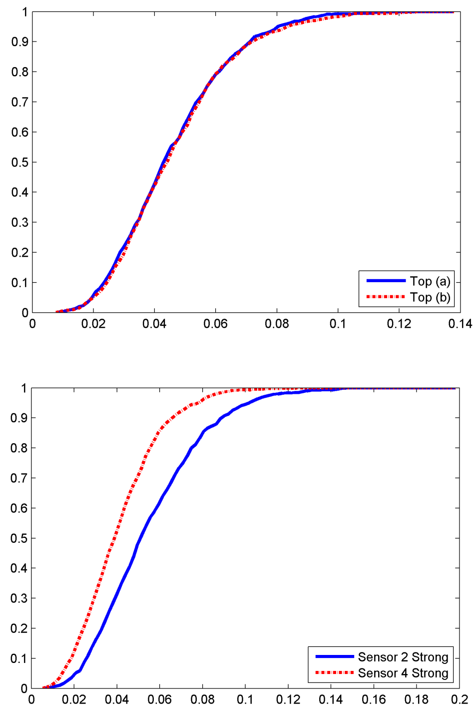

(a)

(b)

Fig. 5. (a) Empirical cumulative distribution functions obtained under the null hypothesis for the two distinct network topologies with four nodes and four edges. (b) Empirical CDFs obtained with various SNRs on the four nodes.

6. DISCUSSION

In the simulation results presented above the conditional distribution, under the signal-absent hypothesis, of the GC estimate formed from a gram matrix obtained by maximum-entropy completion of a partial gram matrix is shown to depend only on the number of elements missing in the original incomplete gram matrix. In particular, under $\mathcal{H}_0$, it does not matter which links are surrogated, only how many. This result enables the fusion center to store only one set of detection thresholds for each number of possible network topologies under a fixed number of surrogations rather than one set for each possible network topology. These empirical results are sufficiently compelling to motivate a comprehensive, theoretical study. At the time of this publication, however, an analytical proof of the desired invariance has not yet been forthcoming.

7. ACKNOWLEDGMENTS

This work was supported in part by the U.S. Army Research Office under MURI Grant No. W911NF-11-1-0391 and by the U.S. Air Force Office of Scientific Research under Grant No. FA9550-12-1-0225.

8. REFERENCES

[1] Kaitlyn Beaudet and Douglas Cochran. Multiple-channel detection in active sensing. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 3910–3914, May 2013.

[2] Kaitlyn Beaudet and Douglas Cochran. Estimation of subspace occupancy. In *Conference Record of the Asilomar COnference on Signals, Systems, and Computers*, November 2014.

[3] Kaitlyn Beaudet, Lauren Crider, and Douglas Cochran. Detection in networked radar. In *Proceedings of the SPIE Defense, Security, and Sensing Conference*, May 2013.

[4] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[5] A. Clausen and D. Cochran. Asymptotic analysis of the generalized coherence estimate. *IEEE Transactions on Signal Processing*, 49(1):45–53, 2001.

[6] D. Cochran, H. Gish, and D. Sinno. A geometric approach to multiple-channel signal detection. *IEEE Transactions on Signal Processing*, 43(9):2049–2057, 1995.

[7] D. Cochran, S. D. Howard, B. Moran, and H. A. Schmitt. Maximum-entropy surrogation in network signal detection. In *Proceedings of the IEEE Statistical Signal Processing Workshop*, August 2012.

[8] H. Gish and D. Cochran. Generalized coherence. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 2745–2748, April 1988.

[9] D. Ramírez, J. Iscar, J. Vía, I. Santamaría, and L. L. Scharf. The locally most powerful invariant test for detecting a rank-P gaussian signal in white noise. In *Proceedings of the IEEE Sensor Array and Multichannel Signal Processing Workshop*, June 2012.

[10] D. Ramírez, G. Vazquez-Vilar, R. López-Valcarce, J. Vía, and I. Santamaría. Detection of rank-P signals in cognitive radio networks with uncalibrated multiple antennas. *IEEE Transactions on Signal Processing*, 59(8):3764–3775, 2011.

[11] D. Ramírez, J. Vía, and I. Santamaría. The locally most powerful test for multiantenna spectrum sensing with uncalibrated receivers. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, March 2012.

[12] D. Ramírez, J. Vía, I. Santamaría, and L. L. Scharf. Detection of spatially correlated Gaussian time series. *IEEE Transactions on Signal Processing*, 58(10):5006–5015, 2010.

[13] S. Sirianunpiboon, S. D. Howard, and D. Cochran. Multiple-channel detection of signals having known rank. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 6536–6540, May 2013.

[14] Lieven Vandenberghe, Stephen Boyd, and Shao-Po Wu. Determinant maximization with linear matrix inequality constraints. *SIAM Journal on Matrix Analysis and Applications*, 19(2):499–533, 1998.