

Quickest Hub Discovery in Correlation Graphs

Taposh Banerjee

School of Engineering and Applied Sciences
Harvard University
Cambridge, MA, USA
Email: tbanerjee@seas.harvard.edu

Alfred O. Hero III

Department of EECS
University of Michigan
Ann Arbor, MI, USA
Email: hero@umich.edu

Abstract—A sequential test is proposed for detection and isolation of hubs in a correlation graph. Hubs in a correlation graph of a random vector are variables (nodes) that have a strong correlation edge. It is assumed that the random vectors are high-dimensional and are multivariate Gaussian distributed. The test employs a family of novel local and global summary statistics generated from small samples of the random vectors. Delay and false alarm analysis of the test is obtained and numerical results are provided to show that the test is consistent in identifying hubs, as the false alarm rate goes to zero.

I. INTRODUCTION

In this paper we consider the problem of quickest detection of a change in correlation between variables in a sequence of high-dimensional vectors. We are also interested in isolating or identifying the variables whose magnitude correlation with other variables has changed the most. We refer to such variables as *hubs*. A precise definition will be given below.

The problem of correlation detection and estimation is an important one in science and engineering. The problem is particularly difficult when the data involved is high-dimensional; see [1], [2], and [3] for a survey of the literature. We are interested in the sequential version of the problem where the variables are initially uncorrelated or independent, and due to some event, the correlation between the variables abruptly changes. This has applications in problems of fault detection, anomaly detection, detection of changes in time-series data or financial data, etc. For simplicity, we restrict our attention to multivariate Gaussian random vectors. However, the results in this paper are application to the broader class of elliptically contoured distributions [4].

If the pre-change parameter are known (means and variances of the uncorrelated variables), the classical approach to sequential detection of a change in the covariance matrix of a sequence of Gaussian random vectors is to use the generalized likelihood ratio (GLR) based Cumulative Sum (CUSUM) method of Lorden [5]. In this algorithm one sequentially estimates the covariance of the data and applies classical quickest change detection (QCD) test using the estimated covariance. But, this algorithm works only for low dimensional data. One approach to handle high-dimensional data, is to make assumptions about the structure of the covariance of the matrix, like sparsity, and then estimate the high-dimensional covariance matrix. Although this approach is appealing, it is

generally hard to design sequential tests based on such an approach. When the pre-change parameters are not known, as is the case in this paper, the problem is even harder. For low dimensional data, one can use a double-GLR approach [6], but optimality of such tests is still under investigation.

In this paper we take a random matrix based approach to solving this problem. We collect the sequence of high-dimensional vectors in batches of successive samples from which we form a random matrix. We then define summary statistics for the random matrix and obtain approximations for the asymptotic distribution for the summary statistics as the dimension of the vector goes to infinity. The asymptotic distribution we obtain belongs to a one-parameter exponential family, and a change in correlation of the random vectors induces a change in the parameter of the random matrix. We use this fact to detect the change in correlation. This works differs from our effort in [4] in that we also define a family of local summary statistics, one for each variable, and obtain their approximate asymptotic distributions. We use these local summary statistics to isolate or identify variables that have seen the most amount of change in correlation with other variables. This fault isolation process is what we call hub discovery. We obtain asymptotic expressions for delay and false alarm for our proposed procedure, and show via numerical results that the test is consistent in hub discovery.

II. PROBLEM DESCRIPTION

We have a sequence of samples of independent and identically distributed (i.i.d.) Gaussian p -variate random vectors $\{\mathbf{X}(k)\}$ with densities

$$f_{\mathbf{X}(k)}(\mathbf{x}) = \frac{1}{\sqrt{2\pi}^p |\Sigma_k|^{p/2}} e^{(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_k)^T \Sigma_k^{-1}(\mathbf{x}-\boldsymbol{\mu}_k))}. \quad (1)$$

For some time parameter γ the samples are assumed to have common covariance parameter $\Sigma_0 = D_p$ for $k < \gamma$, where D_p is a diagonal matrix, and common covariance parameter $\Sigma \neq D_p$ and for $k \geq \gamma$. γ is called the change point and the pre-change and post-change distributions of $\mathbf{X}(k)$ are denoted by $f_{\mathbf{X}}^0$ and $f_{\mathbf{X}}^1$, respectively. No assumptions are made about the mean parameter $\boldsymbol{\mu}_k$, and can take different values for different

k . More specifically, the change-point model is described by:

$$\begin{aligned} \mathbf{X}(k) &\sim f_{\mathbf{X}}^0(\mathbf{x}) = \frac{1}{\sqrt{2\pi}^p |\Sigma_0|^{p/2}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_k)^T \Sigma_0^{-1}(\mathbf{x}-\boldsymbol{\mu}_k)}, \\ & \qquad \qquad \qquad k < \gamma, \\ &\sim f_{\mathbf{X}}^1(\mathbf{x}) = \frac{1}{\sqrt{2\pi}^p |\Sigma|^{p/2}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_k)^T \Sigma^{-1}(\mathbf{x}-\boldsymbol{\mu}_k)}, \\ & \qquad \qquad \qquad k \geq \gamma \end{aligned} \quad (2)$$

Thus, in the above change point model, the variables are *uncorrelated before change*, and *correlated after change*. The objective is to detect this change in correlation as quickly as possible, while avoiding false alarms, and also correctly isolating the variables that have experienced the most change in their correlation level (to be made precise below). The mean parameters $\{\boldsymbol{\mu}_k\}$, the diagonal entries in Σ_0 and the post-change covariance matrix Σ , are *unknown*.

A decision-maker sequentially acquires the samples $\{\mathbf{X}(k)\}$. At each time point k the decision-maker decides to either stop sampling, declaring that the change has occurred, i.e., $k \geq \gamma$, or to continue sampling. The decision to stop at time k is only a function of $(\mathbf{X}(1), \dots, \mathbf{X}(k))$. Thus, the time at which the decision-maker decides to stop sampling, say τ , is a stopping time for the vector sequence $\{\mathbf{X}(k)\}$. At the stopping time τ , the decision maker incurs a delay of $\tau - \gamma$, if $\tau \geq \gamma$. If $\tau < \gamma$, we say that we have an event of false alarm. At time τ , using the available data $(\mathbf{X}(1), \dots, \mathbf{X}(\tau))$, the decision maker also has to identify or isolate variables that have undergone the most amount of change in correlation. We refer to such variables as hubs. We now formalize the notion of a hub.

For a non-diagonal covariance matrix Σ with correlation coefficients $\{\rho_{ki}\}$ define

$$V_k(\Sigma) = \max_{i \neq k} |\rho_{ki}|, \quad \text{for } k \in \{1, \dots, p\}, \quad (3)$$

to be the maximum magnitude correlation coefficient for the k th variable. Hubs are defined as

$$\mathbf{H} = \{k : V_k = \max_{1 \leq j \leq p} V_j\}. \quad (4)$$

Define the correlation graph $\mathcal{G}(\Sigma)$ associated with the post-change correlation matrix Σ as an undirected graph with p vertices, each representing a variable in the vector \mathbf{X} . An edge is present between vertices i and j if the magnitude of the correlation coefficient ρ_{ij} between the i^{th} and j^{th} components of the random vector \mathbf{X} is nonzero, i.e., if $|\rho_{ij}| > 0$, $i \neq j$. The correlation graph $\mathcal{G}(\Sigma)$ can be treated as a weighted graph with weight on edge between node i and j equal to $|\rho_{ij}|$, provided the latter is nonzero. Then, hubs are the nodes with highest maximum weights.

There are also other interesting ways to define hubs. For example, one can select a threshold ρ and define hubs as those variables for which $V_k > \rho$. Another way to define hubs is to redefine correlation graphs with edges present only of $|\rho_{ij}| \geq \rho$, and then define hubs as the nodes with the highest

degree. Results on hub discovery for these other type of hubs will be reported elsewhere. In this paper we only discuss hub discovery for hubs defined in (4).

At time τ , let

$$\mathcal{D}(\mathbf{X}(1), \dots, \mathbf{X}(\tau)) \in 2^{\{1, \dots, p\}} \quad (5)$$

be the decision function that selects a subset of the p variables as the hubs. The decision-maker's objective is to choose a test, i.e. $\{\tau, \mathcal{D}\}$, to detect this change in correlation of the random vectors, as quickly as possible, subject to a constraint on the rate of false alarms, and on the rate of false isolation.

A problem formulation that captures the above performance trade-offs is the following. We seek $\{\tau, \mathcal{D}\}$ to solve

$$\begin{aligned} \min_{\{\tau, \mathcal{D}\}} \quad & \sup_{\gamma \geq 1} \mathbb{E}_\gamma[\tau - \gamma | \tau \geq \gamma] \\ \text{subj. to} \quad & \mathbb{E}_\infty[\tau] \geq \beta, \\ & \text{and } \sup_{\gamma} \mathbb{P}_\gamma[\mathcal{D} \neq \mathbf{H}] \leq \zeta, \end{aligned} \quad (6)$$

where \mathbb{E}_γ is the expectation with respect to the probability measure under which the change occurs at γ , \mathbb{E}_∞ is the expectation when the change never occurs, and β and ζ are user-specified constraints on the mean time to false alarm and the probability of false isolation, respectively.

This problem formulation is also studied in [7] in the context of detecting and isolating a single affected stream out of a finite number of independent streams of observations, when the exact pre- and post-change distributions are known. Even with independent streams, and known distributions, the author in [7] obtained a solution for only a relaxed version of the above problem, by replacing the false isolation probability metric by $\mathbb{P}_1[\mathcal{D} \neq \mathbf{H}]$.

In our case there are three major challenges towards developing a test or extending/applying the results from [7]:

- 1) The means $\{\boldsymbol{\mu}_k\}$, covariance matrices Σ_0 and Σ are unknown
- 2) The high dimension p that is much larger than the total number of samples we will possibly accumulate.
- 3) The streams (one for each dimension) in our observation models are not independent: In fact, we are trying to detect a change in the level of dependence between the streams.

In the sequential analysis literature, for such intractable problems, an alternative approach often taken is to propose a sub-optimal test, and provide a detailed performance analysis of the proposed test. The performance analysis is often used to design the test (to choose thresholds). The effectiveness of the proposed test is then verified by simulations. See for example [8], [9], and [10].

In [4], we proposed a test for the non-parametric family of vector elliptically contoured densities, of which the Gaussian case considered here is a special case, with the objective of minimizing detection delay while avoiding false alarms. In our test in [4], we used a *global* summary statistic to detect the change in correlation. However, the global nature of our summary statistics in [4] does not allow for classification of

where the changes occurred in the correlation matrix, i.e., correlation hub discovery. In this paper we propose a family of p local summary statistics, one for each variable, and obtain asymptotic densities for these local statistics. We use these and the global statistic from [4] to perform sequential classification of the hubs in addition to detection of the existence of hubs.

III. A RANDOM MATRIX SOLUTION APPROACH

We propose a random matrix based solution to our detection problem. The decision maker analyzes the p -dimensional vector sequence $\{\mathbf{X}(k)\}$ in batches of size n , where $p \gg n$. This leads to a $n \times p$ random matrix sequence $\{\mathbb{X}(m)\}_{m \geq 1}$. Specifically,

$$\begin{aligned} \mathbb{X}(1) &= [\mathbf{X}(1), \dots, \mathbf{X}(n)]^T \\ \mathbb{X}(2) &= [\mathbf{X}(n+1), \dots, \mathbf{X}(2n)]^T, \text{ etc.} \end{aligned} \quad (7)$$

Thus, for each m , each of the n rows of the random matrix $\mathbb{X}(m)$ is an i.i.d. samples of the multivariate Gaussian p -variate random vector $\mathbf{X} = [X_1, \dots, X_p]^T$ (1).

A change in the covariance of the random vector sequence $\{\mathbf{X}(k)\}$ as per the description in (2) will change the covariance of the random matrix sequence $\{\mathbb{X}(m)\}$. We propose to detect the change by using a stopping rule for the sequence $\{\mathbb{X}(m)\}$. Note that:

- 1) A stopping rule for the sequence $\{\mathbb{X}(m)\}$ is also a stopping rule for the sequence $\{\mathbf{X}(k)\}$, but it takes values in multiples of n (with respect to index k). The precise choice of n will be discussed below.
- 2) The law of the random matrix \mathbb{X} at the change point may consist of random vectors \mathbf{X} from both pre- and post-change distributions. We ignore this issue in the rest of the paper and assume that all the vectors in the matrix either have density $f_{\mathbf{X}}^0$ or have density $f_{\mathbf{X}}^1$ (2). However, see [11] and [12], where such transient behavior is addressed.

In random matrix theory (RMT), asymptotic distributions are obtained for various functions of the random matrix, e.g., a semicircle law for the empirical eigen value distribution of a Wigner matrix, etc [13]. Such results are in general valid when both the dimensions of the matrix are taken to infinity. For our random matrix \mathbb{X} , however, the batch size n is fixed, and the dimension p is large. This is called the *purely high-dimensional regime* [14]. Thus, traditional RMT theorems cannot be applied in this regime. In this paper, we obtain approximate asymptotic distribution for a proposed summary statistic (in fact, a family of them) for the random matrix \mathbb{X} in this regime. We then use this approximate distribution(s) to detect a change in the law of the random matrix. We also propose techniques to isolate the hubs.

IV. LOCAL SUMMARY STATISTICS FOR THE RANDOM DATA MATRIX

In this section we define a family of local summary statistics to detect hubs (4). The summary statistics are a function of the sample correlation matrix obtained from the random data matrix \mathbb{X} . We then obtain their asymptotic distribution

in the purely high dimensional regime of $p \rightarrow \infty$, n fixed. The asymptotic distributions are members of a one-parameter exponential family to be specified below. In Section VI, we will use these summary statistics and their asymptotic distributions for sequential hub discovery.

For random data matrix \mathbb{X} we write

$$\mathbb{X} = [\mathbf{X}_1, \dots, \mathbf{X}_p] = [\mathbf{X}_{(1)}^T, \dots, \mathbf{X}_{(n)}^T]^T,$$

where $\mathbf{X}_i = [X_{1i}, \dots, X_{ni}]^T$ is the i^{th} column and $\mathbf{X}_{(i)} = [X_{i1}, \dots, X_{ip}]$ is the i^{th} row. Define the sample covariance matrix as

$$\mathbf{S} = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{X}_{(i)} - \bar{\mathbf{X}})^T (\mathbf{X}_{(i)} - \bar{\mathbf{X}}),$$

where $\bar{\mathbf{X}}$ is the sample mean of the n rows of \mathbb{X} . Also define the sample correlation matrix as

$$\mathbf{R} = \mathbf{D}_{\mathbf{S}}^{-1/2} \mathbf{S} \mathbf{D}_{\mathbf{S}}^{-1/2},$$

where $\mathbf{D}_{\mathbf{A}}$ denotes the matrix obtained by zeroing out all but the diagonal elements of the matrix \mathbf{A} . Note that, under our assumption that the covariance matrix Σ of the rows of \mathbb{X} is positive definite, $\mathbf{D}_{\mathbf{S}}$ is invertible with probability one. Thus \mathbf{R}_{ij} , the element in the i^{th} row and the j^{th} column of the matrix \mathbf{R} , is the sample correlation coefficient between the i^{th} and j^{th} columns of \mathbb{X} .

Define the summary statistics (see also (3))

$$V_k(\mathbb{X}) := V_k(\mathbf{R}) := \max_{i \neq k} |\mathbf{R}_{ki}|, \quad k \in \{1, \dots, p\}. \quad (8)$$

This defines a set of p local statistics $\{V_k\}_{k \leq p}$, one for each random variable in the random vector. We will use these statistics for detection of correlation strength hubs.

Note that the summary statistics are the empirical estimates of the variables $\{V_k\}$ defined for the population correlation matrix Σ in (3). If we could collect enough samples such that $n > p$, then we can rely on these estimates to check if they are close to zero, because the population version of these quantities are identically zero when the variables are uncorrelated. But, since $p \gg n$, we may have to collect a large number of samples before we can rely on these estimates. Such a delay may not be acceptable in practice, or those many samples may not even be available. In fact, for large p , $V_k(\mathbb{X})$ can be arbitrarily close to 1; see [15] and Fig. 2 in this paper.

However, as we show below, one can obtain asymptotic law of these estimates as $p \rightarrow \infty$. The precise statement is provided in the theorems below. A change in correlation between the variables may be reflected in the asymptotic law of these estimates. This phenomena can be utilized for detecting a change in correlation. The rest of this section is devoted to obtaining the asymptotic laws. The algorithms are discussed in Section VI.

We say that a matrix is row sparse of degree j if there are no more than j nonzero entries in any row. We say that a matrix is block sparse of degree j if the matrix can be reduced to block diagonal form having a single $j \times j$ block, via row-column permutations. For a threshold parameter $\rho \in [0, 1]$

define the correlation graph $\mathcal{G}_\rho(\Sigma)$ associated with the post-change correlation matrix Σ as an undirected graph with p vertices, each representing a variable in the vector \mathbf{X} . An edge is present between vertices i and j if the magnitude of the correlation coefficient ρ_{ij} between the i^{th} and j^{th} components of the random vector \mathbf{X} is greater than ρ , i.e., if $|\rho_{ij}| \geq \rho$, $i \neq j$. We define d_k to be the degree of vertex k in the graph $\mathcal{G}_\rho(\Sigma)$:

$$d_k(\Sigma, \rho) := \#\{i \neq k : |\rho_{ki}| \geq \rho\}, \quad (9)$$

$$k \in \{1, \dots, p\}, \rho \in [0, 1],$$

where $\#A$ is the size of set A . Finally, throughout the index k will be associated with the k th variable.

Theorem 4.1: Let Σ be the covariance matrix of \mathbf{X} or the rows of \mathbb{X} , Σ could be diagonal or non-diagonal depending on whether it is from the pre- or post-change data, respectively. Let Σ be row sparse of degree $j = o(p)$. Also let $p \rightarrow \infty$ and $\rho = \rho_p \rightarrow 1$ such that $p(1 - \rho^2)^{n/2} \rightarrow 0$. Then,

$$P(V_k(\mathbb{X}) \leq \rho) \rightarrow e^{-\Lambda_k}, \quad (10)$$

where

$$\Lambda_k = \lim_{p \rightarrow \infty, \rho \rightarrow 1} \Lambda_{k,\rho},$$

with

$$\Lambda_{k,\rho} = (p-1)P_0(\rho)J_k,$$

$$P_0(\rho) = I_{1-\rho^2}((n-2)/2, 1/2),$$

with $I_x(a, b)$ being the regularized incomplete beta function with parameter a and b . The parameter J_k is a positive real number that is a function of the asymptotic joint density of \mathbf{X} . Further, if the covariance matrix Σ of the p -variate vector \mathbf{X} is block sparse of degree j , then

$$J_k = 1 + O((j/p)^2).$$

In particular, if Σ is diagonal then

$$J_k = 1.$$

Proof: The events $\{V_k(\mathbb{X}) \geq \rho\}$ and $\{d_k > 0\}$ are equivalent. Hence

$$P(V_k(\mathbb{X}) \geq \rho) = P(d_k(\mathbb{X}, \rho) > 0). \quad (11)$$

The result now follows from Theorem 4.2 below. \blacksquare

Theorem 4.2 ([16]): Let Σ be row sparse of degree $j = o(p)$. Also let $p \rightarrow \infty$ and $\rho = \rho_p \rightarrow 1$ such that $p(1 - \rho^2)^{n/2} \rightarrow 0$. Then $d_k(\mathbb{X}, \rho)$ converges in distribution to a Poisson random variable with rate parameter Λ_k .

Based on the asymptotic distribution obtained in Theorem 4.1, the large p distribution of V_k defined in (8) can be approximated by

$$P(V_k(\mathbb{X}) \leq \rho) \approx \exp(-\Lambda_{k,\rho}), \quad (12)$$

$$= \exp(-(p-1)P_0(\rho)J_k) \rho \in [0, 1].$$

where $\Lambda_{k,\rho}$ is as defined in Theorem 4.1. Using [16, Thm. 3.1] it can easily be shown that the approximation error associated

with (14) decays to zero at least as fast as $p(1 - \rho^2)^{n/2}\Delta$ where Δ is a dependency coefficient associated with the set of U-scores..

The distribution (12) is differentiable everywhere except at $\rho = 0$ since $P(V_k(\mathbb{X}) = 0) > 0$. For $\rho > 0$ and large p , V_k has density

$$f_{V_k}(\rho) \approx -(p-1)P_0'(\rho)J_k \exp(-(p-1)P_0(\rho)J_k), \rho \in (0, 1]. \quad (13)$$

Note that f_V in (13) is the density of the Lebesgue continuous component of the distribution (12) and that it integrates to $1 - O(e^{-p^2})$ over $\rho \in (0, 1]$.

For each k , the density f_{V_k} is a member of a one-parameter exponential family with J_k as the unknown parameter. Then, since $P_0(\rho)' = -\frac{2(1-\rho^2)^{(n-4)/2}}{B((n-2)/2, 1/2)}$, where $B(a, b)$ is the beta function, we have for $\rho \in [0, 1]$, the exponential family form of the density f_V with parameter J_k :

$$f_{V_k}(\rho; J_k)$$

$$= \frac{2(p-1)J_k(1-\rho^2)^{\frac{n-4}{2}}}{B((n-2)/2, 1/2)} \exp(-(p-1)J_k P_0(\rho)). \quad (14)$$

Note that the density f_{V_k} corresponds to the k th variable, and the whole family $\{f_{V_k}\}$ depends on the index k only through the parameter J_k , which can be different for different k . We thus rewrite the family in (14) (with complete specification for easy accessibility) as

$$f_V(y; J) = J C(p, n) (1 - y^2)^{\frac{n-4}{2}} e^{-(p-1)J P_0(y)}, \quad (15)$$

where

$$C(p, n) = \frac{2(p-1)}{B((n-2)/2, 1/2)},$$

$$P_0(y) = I_{1-y^2}((n-2)/2, 1/2),$$

$$I_x(a, b) = \frac{B(x; a, b)}{B(a, b)}, \quad (16)$$

$$B(x, a, b) = \int_0^x t^{a-1}(1-t)^{b-1} dt,$$

$$B(a, b) = B(1, a, b).$$

In Fig. 1 we have plotted the density f_V for various values of J for $n = 10$, and $p = 100$. We note that for the chosen values of n and p , the density is concentrated close to 1, consistent with large values of ρ arising in the purely high dimensional regime assumed in Theorem 4.1 [14]. Note that *the variable V tends to take on higher values as the parameter J increases*. This is the fundamental idea used for hub discovery in this paper.

In Fig. 2 we have plotted the normalized histogram of the samples $V_1(\mathbb{X})$ for multivariate Gaussian independent random vectors and compare it with the approximation $f_V(\cdot; 1)$. Note that the approximation is quite accurate. Also, note that for this case the population $V_1 = 0$, but the samples are concentrated in the interval (0.55, 0.95). This is due to the fact that $p \gg n$.

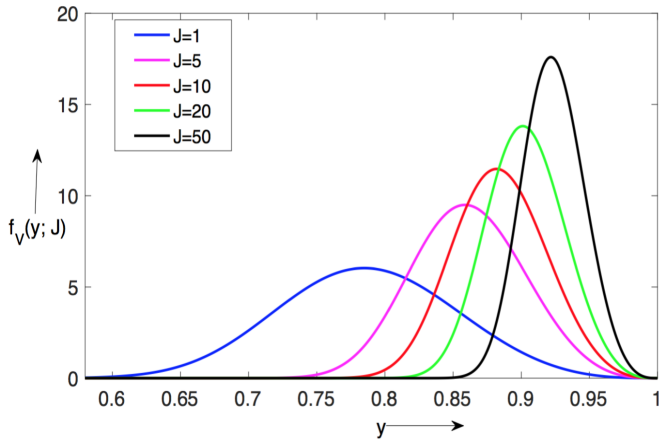


Fig. 1. Plot of density f_V in (15) for various values of the parameter J for $n = 10, p = 100$.

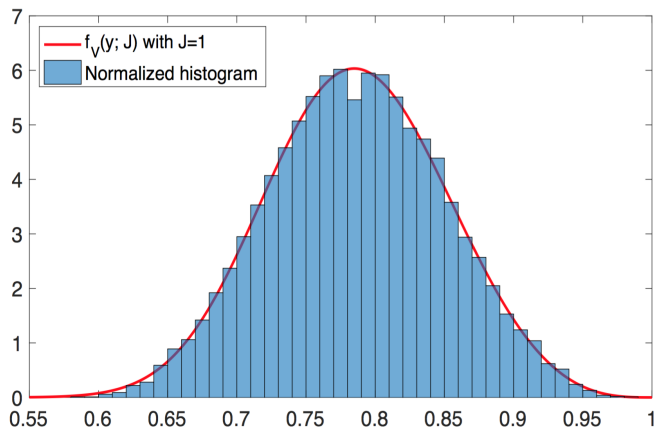


Fig. 2. Comparison of density f_V in (15) for $J = 1$ for multivariate Gaussian independent random variables with $n = 10, p = 100$ with the normalized histogram obtained from 10,000 samples. Note that the population covariance is diagonal, yet the density is concentrated close to 1.

V. GLOBAL SUMMARY STATISTICS

In [4] we proposed a global summary statistic and obtained its asymptotic distribution. Specifically, we showed that the statistic

$$U(\mathbb{X}) = \max_k V_k(\mathbb{X}), \quad (17)$$

in the limit $p \rightarrow \infty$ (taken in a particular sense) and under sparsity assumptions similar to that in Theorem 4.1, has approximate density in an exponential family given by

$$g(u; \theta) = \frac{D\theta}{2} (1-u^2)^{\frac{n-4}{2}} \exp\left(-\frac{D}{2}\theta T(u)\right), \quad u \in (0, 1], \quad (18)$$

where

$$D = D_{p,n} = 2p(p-1)B((n-2)/2, 1/2), \quad (19)$$

$$T(u) = \int_u^1 (1-s^2)^{\frac{n-4}{2}} ds.$$

The parameter θ is a positive parameter and has similar interpretation and nature as that of the parameter J in (15). Specifically, if the rows of \mathbb{X} have diagonal covariance Σ then $\theta = 1$, otherwise $\theta \neq 1$. Compared to $V_k(\mathbb{X})$, $U(\mathbb{X})$ has a global view of the covariance matrix. If any of the V_k s undergo a change, this change will be reflected through a change in the distribution of U . However, it is not possible to isolate hubs based only on the statistic U . Sequential tests for detecting a change in correlation using the statistic U is studied in detail in [4]. The sequential test from [4] will play a fundamental role in what follows.

VI. MULTIVARIATE QUICKEST HUB DISCOVERY

We now discuss how to use the results from the previous sections to detect and isolate correlation hubs.

Recall that we have a sequence of p -variate Gaussian distributed vectors $\{\mathbf{X}(k)\}$ with a change point model given in (2) and $p \gg 1$. Thus, the vectors are uncorrelated before change and correlated with covariance matrix Σ after change. We collect these vectors in batches of size n and obtain a sequence of $n \times p$ random matrices $\{\mathbb{X}(m)\}$; see (7). For quick detection, we should choose a batch size n . Thus, we also have $p \gg n$.

Assuming that the post-change covariance matrix Σ is row sparse, a change in the law of \mathbf{X} changes the law of the random matrix \mathbb{X} , which in turn affects the law of summary statistics $U(\mathbb{X})$ and $\{V_k(\mathbb{X})\}_{k=1}^p$. Since, the distributions of these statistic each belong to a one-parameter exponential family (15) and (18), a change in law within this family is a simple change in the value of the parameters $\{J_k\}$ and θ . We propose to detect the change in correlation in \mathbf{X} through a change in the parameters $\{J_k\}$ and θ

From (15) and (18), it follows that we can choose $n > 4$, because the densities are well defined for these values of n . Note also that here we assume that the change in correlation in \mathbf{X} affects the parameters $\{J_k\}$ and θ , otherwise our test is insensitive to the change in correlation.

A. Correlation Change Detection Using Global Statistic

Assuming that the post-change covariance matrix Σ is row sparse, we map the random matrix sequence $\{\mathbb{X}(m)\}$ to the summary statistic sequence $\{U(\mathbb{X}(m))\}$. The random variables $\{U(\mathbb{X}(m))\}$ have approximate density given by (18) which is in an exponential family with parameter θ which is 1 before the change and some value $\theta \neq 1$ after change. For simplicity, we refer to $U(\mathbb{X}(m))$ simply by $U(m)$. The QCD problem on the density $f_{\mathbf{X}}$, depicted in (2), is reduced to the following QCD problem

$$U(m) \sim g(\cdot; 1), \quad m < \gamma$$

$$\sim g(\cdot; \theta), \quad \theta \neq 1, m \geq \gamma. \quad (20)$$

Consider the following generalized likelihood ratio (GLR)

based QCD tests (see [5], [4]) defined by the stopping time:

$$\tau_u = \min \left\{ m : \max_{1 \leq \ell \leq m} \sup_{\theta: |\theta-1| \geq \epsilon_u} \sum_{i=\ell}^m \log \frac{g(U(i); \theta)}{g(U(i); 1)} > A_u \right\}, \quad (21)$$

where we assign $\tau_u = \infty$ if the right hand side of (23) is the empty set. Here $A_u > 0$ and $\epsilon_u > 0$ are user-defined parameters, and ϵ_u represents the minimum magnitude of change, away from $\theta = 1$, that the user wishes to detect.

B. Correlation Change Detection Using Local Statistics

We can also detect a change in correlation using the local sample correlation statistics $\{V_k(\mathbb{X})\}_{k=1}^p$ defined in (8). For simplicity, $V_k(\mathbb{X}(m))$ is denoted simply by $V_k(m)$. For large p and row sparse covariance, the random variables $\{V_k(m)\}$ have approximate density given by (15) which is in an exponential family with parameter J_k . This parameter is 1 before the change and takes on some value $J_k \neq 1$ after the change. A natural locally pooled sequential change detection test is to implement p local sequential change detection tests in parallel, one for each variable $\{V_k(\mathbb{X})\}_{k=1}^p$ (equivalently, each variable in \mathbf{X}). An alarm is raised when *any* of the local tests raises an alarm.

Mathematically, the QCD problem on the density $f_{\mathbf{X}}$, depicted in (2), is reduced to the following family of QCD problems (one for each k):

$$\begin{aligned} V_k(m) &\sim f_V(\cdot; 1), & m < \gamma \\ &\sim f_V(\cdot; J_k), & J_k \neq 1, m \geq \gamma. \end{aligned} \quad (22)$$

Consider the following GLR-based QCD tests defined by the family of local stopping times [5]:

$$\tau_v^{(k)} = \min \left\{ m : \max_{1 \leq \ell \leq m} \sup_{J_k: |J_k-1| \geq \epsilon_v} \sum_{i=\ell}^m \log \frac{f_V(V_k(i); J_k)}{f_V(V_k(i); 1)} > A_v \right\}, \quad (23)$$

where, again we assign $\tau_v^{(k)} = \infty$ if the right hand side of (23) is the empty set, and $A_v > 0$ and $\epsilon_v > 0$ are user-defined parameters. This stopping time can be used to detect a change in the parameter J_k for the k th variable. Define the overall local stopping time to be

$$\tau_v = \min_k \tau_v^{(k)} = \min \left\{ m : \max_{1 \leq k \leq p} G_k(m) > A_v \right\}, \quad (24)$$

where

$$G_k(m) := \max_{1 \leq \ell \leq m} \sup_{J_k: |J_k-1| \geq \epsilon_v} \sum_{i=\ell}^m \log \frac{f_V(V_k(i); J_k)}{f_V(V_k(i); 1)} \quad (25)$$

is the GLR statistic. The test τ_v can be used to detect if the parameter of any of the p variables is affected by the change. ϵ_v represents the minimum magnitude of change, away from $J_k = 1$, that the user wishes to detect.

C. Joint Detection and Hub Discovery

For joint detection and hub discovery, we combine the local and global tests and propose to use the following maximum of the stopping times:

$$\tau_{\text{HB}} = \max\{\tau_v, \tau_u\}. \quad (26)$$

Let $D_k(\mathbb{X}(1), \dots, \mathbb{X}(\tau_{\text{HB}}))$ be the binary decision variable with value 1 if variable k is declared a hub. We use the following rule for hub discovery: fix a positive integer q and set

$$D_k(\mathbb{X}(1), \dots, \mathbb{X}(\tau_{\text{HB}})) = \mathbb{I}\{G_k(\tau_{\text{HB}}) \text{ is top } q \text{ statistic}\}. \quad (27)$$

To appreciate the above rule, recall that a correlation strength hub is defined as the set of those variables that achieve the maximum in (17); see (4). Thus, there are at least two variables that are correlation strength hubs in the correlation graph. These hub variables in general have value of parameter J_k that is higher than the parameter values for other variables. However, there may be many variables that have experienced a change in correlation, and their V_k value, and hence the value of parameter J_k , might be close to the values attained by the hubs. To illustrate this idea, we have plotted in Fig. 3 estimated $\{J_k\}$ values from samples of a multivariate Gaussian vector with a row sparse covariance matrix (see Section VII on how the covariance matrix is generated). Here the hubs are the two variables with the highest values of J_k . But, as seen in the figure, there are many other variables that have experienced significant changes in correlation, and hence also have high values of J_k . In Fig. 4, we have plotted the values of the GLR statistic $\{G_k\}$ at the time of stopping τ_{HB} . As seen in the figure, in addition to the G_k values of hubs, the values corresponding to other variables are also significant. Thus, a simple strategy like choosing the variable with the highest G_k value as hub, will not work well.

However, what is clear from Fig. 3 and Fig. 4 is that high G_k values at the time of stopping correspond to high J_k values of the summary statistics, and these values together have significantly higher magnitude than the values for the rest of the variables. This motivates the pick-the-top approach. In fact, in practice, *one can plot the G_k values this way and identify all the variables that have experienced a change in correlation.*

Picking the variables corresponding to the top q values where $q > 2$ ensures that the hubs are discovered with high probability, provided a slightly modified definition of hub discovery is used: the variables in the hub H are considered discovered if

$$H \subset \{k : D_k = 1\}. \quad (28)$$

Indeed, this fact is used in Section VII for numerical computations, where it is shown that the proposed rule's probability of false hub discovery goes to zero, as the false alarm rate goes to zero; see Fig. 5.

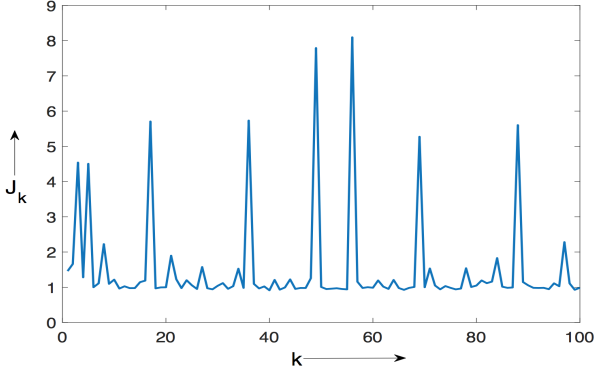


Fig. 3. Estimated J_k values from a $p \times p$ row-sparse covariance matrix with $p = 100$. Here k is the index of the variables in \mathbf{X} . Note that the values of parameter J_k for non-hub variables are also significant. Hubs are the two variables with the highest values of the parameter.

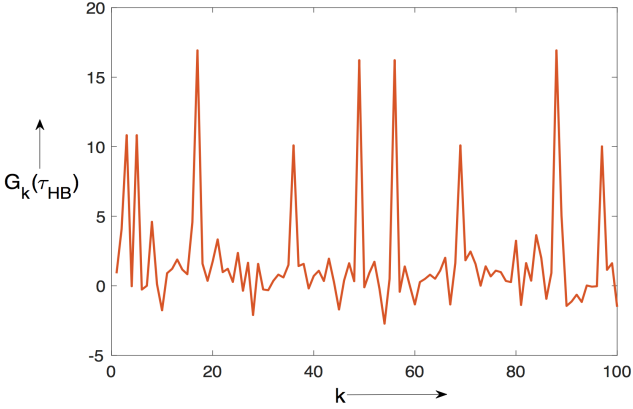


Fig. 4. Values of GLR statistic G_k at the time of stopping τ_{HB} corresponding to the model used in Fig. 3. Note that the G_k values for the non-hub variables are comparable to those of the hub variables. In practice, one can just plot these G_k values to capture all the variables that are affected by a change in correlation. A close mathematical equivalent to this idea is the pick-the-top q approach (27).

D. Delay and False Alarm Analysis

The approximate delay and false alarm performance of the stopping rule τ_{HB} is given in the theorem below.

Theorem 6.1: Fix any $\epsilon_u > 0$ and $\epsilon_v > 0$. If the densities (15) and (18) are the true distributions of the samples then we have the following:

- 1) For the stopping rule τ_{HB} , the supremum in (6) is achieved at $\gamma = 1$, i.e.,

$$\sup_{\gamma \geq 1} \mathbf{E}_\gamma[\tau_{\text{HB}} - \gamma | \tau_{\text{HB}} \geq \gamma] = \mathbf{E}_1[\tau_{\text{HB}} - 1].$$

- 2) Setting $A_u = A_v = \log \beta$ ensures that as $\beta \rightarrow \infty$,

$$\mathbf{E}_\infty[\tau_{\text{HB}}] \geq \beta(1 + o(1)).$$

- 3) For each possible true post-change parameters J and θ , with $|J - 1| \geq \epsilon_v$, and $|\theta - 1| \geq \epsilon_u$, as $\beta \rightarrow \infty$

$$\mathbf{E}_1[\tau_{\text{HB}}] \leq \left(\frac{\log \beta}{I(J)} + \frac{\log \beta}{I(\theta)} \right) (1 + o(1)), \quad (29)$$

where $I(J)$ is the Kullback-Leibler divergence between the densities $f_V(\cdot; J)$ and $f_V(\cdot; 1)$, and $I(\theta)$ is the Kullback-Leibler divergence between the densities $g(\cdot; \theta)$ and $g(\cdot; 1)$.

Proof: The proof follows from standard QCD arguments in the literature (see [17]) and because

$$\tau_{\text{HB}} \geq \tau_U,$$

and

$$\tau_{\text{HB}} \leq \tau_U + \tau_V. \quad \blacksquare$$

VII. NUMERICAL RESULTS

In Fig. 5 we report the hub discovery performance of the proposed algorithm. We set $\Sigma_0 = I_p$, the identity matrix. We assume that the post-change covariance matrix Σ is a row-sparse matrix of degree j , obtained as follows. A $p \times p$ sample from the Wishart distribution is generated and some of the entries are forced to be zero in such a way that no row has more than j non-zero elements. Specifically, we retain the top left $j \times j$ block of the matrix, and for each row k , $j + 1 \leq k \leq (p + j)/2$, all but the diagonal and the $(p + j + 1 - k)$ th element is forced to zero. Each time an entry (k, i) is set to zero, the entry (i, k) is also set to zero, to maintain symmetry. Finally, a scaled diagonal matrix is also added to Σ to restore its positive definiteness. We set $n = 10$, $p = 100$, and $j = 5$.

To implement τ_V we have chosen $\epsilon_v = 1$, and we use the the maximum likelihood estimator for J_k which, as a function of m samples $(V_k(1), \dots, V_k(m))$ from $f_V(\cdot, J)$, is given by

$$\hat{J}(V_k(1), \dots, V_k(m)) = \frac{1}{(p-1) \frac{1}{m} \sum_{i=1}^m P_0(V_k(i))}. \quad (30)$$

Specifically,

$$\begin{aligned} \arg \max_{J: J \geq 2} \log \sum_{i=\ell}^m \frac{f_V(V_k(i); J)}{f_V(V_k(i); 1)} \\ = \max\{2, \hat{J}(V_k(\ell), \dots, V_k(m))\}. \end{aligned} \quad (31)$$

Similar techniques were used to implement τ_U ; see [4].

We consider three scenarios, each with two hubs:

1. $H = \{19, 86\}$ with $J_{19} = 2.56$ and $J_{86} = 2.51$.
2. $H = \{44, 61\}$ with $J_{44} = 5.88$ and $J_{61} = 5.85$.
3. $H = \{12, 93\}$ with $J_{12} = 17$ and $J_{93} = 16$.

The parameter J_k s are estimated using 1000 samples. For each case, the total number of false isolations is calculated using 1000 sample paths by assuming $\gamma = 1$. The parameter q was chosen to be 10. The definition (28) was used as a criterion calculating false isolation or false hub discovery. The results for all the three cases is plotted in Fig. 5 as a function of the threshold $A = A_v = A_u$. As shown in Fig. 5, the number of false discoveries goes to zero as the threshold is increased. Thus, empirical evidence suggests that the test is consistent.

In Fig. 6 we plot the delay ($\mathbf{E}_1[\tau_{\text{HB}}]$) vs the log of mean time to false alarm ($\log \mathbf{E}_\infty[\tau_{\text{HB}}]$) for various cases specified in

IX. ACKNOWLEDGMENTS

This work was partially supported by the Consortium for Verification Technology under Department of Energy National Nuclear Security Administration award number DOE-NA0002534.

REFERENCES

- [1] M. Pourahmadi, *High-dimensional covariance estimation: with high-dimensional data*. John Wiley & Sons, 2013.
- [2] J. Fan, Y. Liao, and H. Liu, "An overview of the estimation of large covariance and precision matrices," *The Econometrics Journal*, vol. 19, no. 1, pp. C1–C32, 2016.
- [3] T. T. Cai, Z. Ren, H. H. Zhou, *et al.*, "Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation," *Electronic Journal of Statistics*, vol. 10, no. 1, pp. 1–59, 2016.
- [4] T. Banerjee, H. Firouzi, and A. O. Hero, "Non-parametric quickest change detection for large scale random matrices," in *IEEE International Symposium on Information Theory (ISIT)*, pp. 146–150, June 2015. <http://arxiv.org/abs/1508.04720>.
- [5] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, pp. 1897–1908, Dec. 1971.
- [6] T. L. Lai, "Sequential changepoint detection in quality control and dynamical systems," *J. Roy. Statist. Soc. Suppl.*, vol. 57, no. 4, pp. pp. 613–658, 1995.
- [7] A. G. Tartakovsky, "Multidecision quickest change-point detection: Previous achievements and open problems," *Sequential Analysis*, vol. 27, pp. 201–231, Apr. 2008.
- [8] D. Siegmund and E. S. Venkatraman, "Using the generalized likelihood ratio statistic for sequential detection of a change-point," *Ann. Statist.*, vol. 23, pp. 255–271, Feb. 1995.
- [9] F. Desobry, M. Davy, and C. Doncarli, "An online kernel change detection algorithm," *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 2961–2974, 2005.
- [10] H. Chen, "Sequential change-point detection based on nearest neighbors," *arXiv preprint arXiv:1604.03611*, 2016.
- [11] G. V. Moustakides and V. V. Veeravalli, "Sequentially detecting transitory changes," in *Information Theory (ISIT), 2016 IEEE International Symposium on*, pp. 26–30, 2016.
- [12] T. Banerjee, Y. C. Chen, A. D. Dominguez-Garcia, and V. V. Veeravalli, "Power system line outage detection and identification – a quickest change detection approach," in *IEEE Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2014. submitted.
- [13] G. W. Anderson, A. Guionnet, and O. Zeitouni, *An introduction to random matrices*, vol. 118. Cambridge university press, 2010.
- [14] A. O. Hero and B. Rajaratnam, "Foundational principles for large-scale inference: Illustrations through correlation mining," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 93–110, 2016.
- [15] A. Hero and B. Rajaratnam, "Hub discovery in partial correlation graphs," *IEEE Trans. Inf. Theory*, vol. 58, no. 9, pp. 6064–6078, 2012.
- [16] H. Firouzi and A. O. Hero, "Local hub screening in sparse correlation graphs," in *SPIE Optical Engineering+ Applications*, pp. 88581H–88581H, International Society for Optics and Photonics, 2013.
- [17] V. V. Veeravalli and T. Banerjee, *Quickest Change Detection*. Elsevier: E-reference Signal Processing, 2013. <http://arxiv.org/abs/1210.5552>.

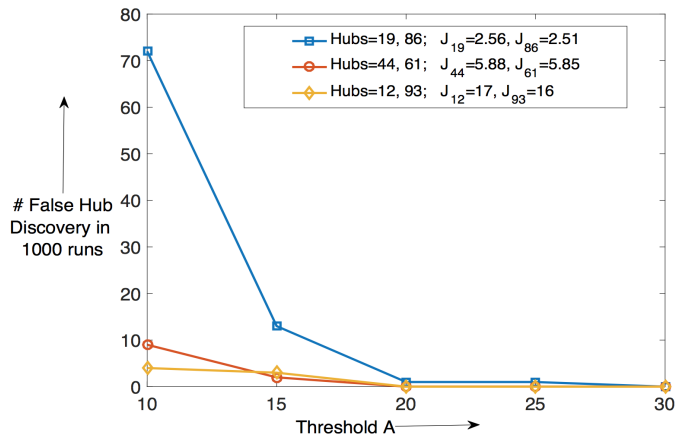


Fig. 5. False Isolation performance for three scenarios in (32)

the figure. The values in the figure are obtained by choosing different values of the threshold A and estimating the delay by choosing the change point $\gamma = 1$ and simulating the test for 500 sample paths. The mean time to false alarm values are estimated by simulating the test for 1500 sample paths.

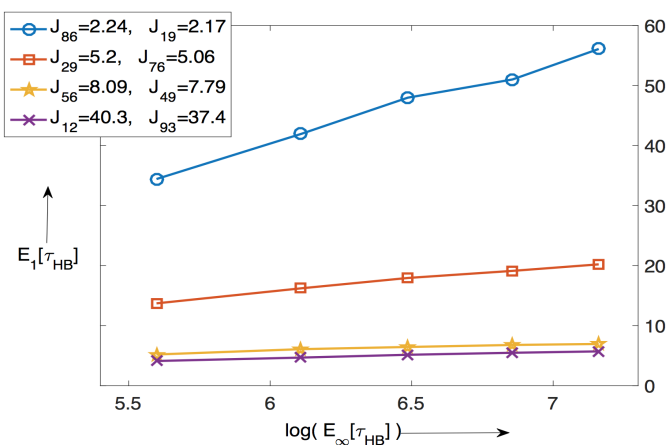


Fig. 6. MFA-Delay performance for the stopping rule τ_{HB} for various choices of parameter $\{J_k\}$. The subscripts of J are the hub indices.

VIII. CONCLUSIONS AND FUTURE WORK

We have introduced the quickest hub detection and localization problem for high dimensional correlation graphs. Local and global summary statistics were proposed based on purely high dimensional asymptotic theory for the presence of hub nodes, defined as vertices having either high degree or large correlations. The asymptotic densities of these summary statistics were shown to be in the one dimensional exponential family of densities. These densities were used to define a hybrid local and global quickest detection test for hubs and the asymptotic performance of this test was evaluated in terms of mean time to detect a change and mean time to false alarm.