Using Scan-Statistical Correlations for Network Change Analysis

Adriel Cheng and Peter Dickinson

Command, Control, Communications and Intelligence Division
Defence Science and Technology Organisation, Department of Defence, Australia
{adriel.cheng,peter.dickinson}@dsto.defence.gov.au

Abstract. Network change detection is a common prerequisite for identifying anomalous behaviours in computer, telecommunication, enterprise and social networks. Data mining of such networks often focus on the most significant change only. However, inspecting large deviations in isolation can lead to other important and associated network behaviours to be overlooked. This paper proposes that changes within the network graph be examined in conjunction with one another, by employing correlation analysis to supplement network-wide change information. Amongst other use-cases for mining network graph data, the analysis examines if multiple regions of the network graph exhibit similar degrees of change, or is it considered anomalous for a local network change to occur independently. Building upon Scan-Statistics network change detection, we extend the change detection technique to correlate localised network changes. Our correlation inspired techniques have been deployed for use on various networks internally. Using real-world datasets, we demonstrate the benefits of our correlation change analysis.

Keywords: Mining graph data, statistical methods for data mining, anomaly detection.

1 Introduction

Detecting changes in computer, telecommunication, enterprise or social networks is often the first step towards identifying anomalous activity or suspicious participants within such networks. In recent times, it has become increasingly common for a changing network to be sampled at various intervals and represented naturally as a time-series of graphs [1,2]. In order to uncover network anomalies within these graphs, the challenge in network change detection lies not only with the type of network graph changes to observe and how to measure such variations, but also the subsequent change analysis that is to be conducted.

Scan-statistics [3] is a change detection technique that employs statistical methods to measure variations in network graph vertices and surrounding vertex neighbourhood regions. The technique uncovers large localised deviations in behaviours exhibited by subgraph regions of the network. Traditionally, the subsequent change analysis focuses solely on local regions with largest deviations.

J. Li et al. (Eds.): PAKDD 2013 Workshops, LNAI 7867, pp. 1–13, 2013.

Despite usefulness in tracking such network change to potentially anomalous subgraphs, simply distinguishing which subgraph vertices contributed most to the network deviation is insufficient.

In many instances, the cause of significant network changes may not be restricted to a single vertex or subgraph region only, but multiple vertices or subgraphs may also experience similar degrees of deviations. Such vertex deviations could be interrelated, acting collectively with other vertices as the primary cause of the overall network change. Concentrating solely on the most significantly changed vertex or subgraph, other localised change behaviours would be hidden from examination. The dominant change centric analysis may in fact hinder evaluation of the actual change scenario experienced by the network.

To examine the network more conclusively, rather than inspect the most deviated vertex or subgraph, scan-statistic change analysis should characterise the types of localised changes and their relationships with one another across the entire network.

With this in mind, we extend scan-statistics with correlation based computations and change analysis. Using our approach, correlations between the edge-connectivity changes experienced by each pair of network graph vertices (or subgraphs) are examined. Correlation measurements are also aggregated to describe the correlation of each vertex (or subgraph) change with all other graph variations, and to assess the overall correlation of changes experienced by the network as a whole.

The goal in supporting scan-statistical change detections with our correlations-based analysis is to seek-out and characterise any relational patterns in the localised change behaviours exhibited by vertices and subgraphs. For instance, if a significant network change is detected and attributed to a particular vertex, do any other vertices in the network show similar deviation in behaviours? If so, how many vertices are considered similar? Do the majority of the vertices experience similar changes, or are these localised changes independent and not related to other regions of the network.

Accounting for correlations between localised vertex or subgraph variations provides further context into the possible scenarios triggering such network deviations. For example, if localised changes in the majority of vertices are highly correlated with one another, this could imply a scenario whereby a network-wide initialisation or re-configuration took place. In a social network, such high correlations of increased edge-linkages may correspond to some common holiday festive event, whereby individuals send/receive greetings to everyone on the network collectively within the same time period. Or if the communication links (and traffic) of a monitored network of terrorist suspects intensifies as a group, this could signal an impending attack.

On the other hand, a localised vertex or subgraph change which is uncorrelated to other members of the graph may indicate a command-control network scenario, In this case, any excessive change in network edge-connectivity would be largely localised to the single command vertex. Another example could involve the failure of a domain name system (DNS) server or a server under a denial-of-service attack. In this scenario, re-routing of traffic from the failed server to an alternative server would take place. The activity changes at these two server vertices would be highly localised and not correlated to the remainder of the network.

To the best of our knowledge, examining scan-statistical correlations of network graphs in support of further change analysis has not been previously explored. Hence, the contributions of this paper are two-fold. First, to extend scan-statistics network change detection with correlations analysis at multiple levels of the network graphs. And second, to facilitate visualisation of vertex clusters and reveal interrelated groups of vertices whose collective behaviour requires further investigation.

The remainder of this paper is as follows. Related work is discussed next. Section 3 gives a brief overview of scan-statistics. Sections 4 to 6 describe the correlation extensions and correlation inspired change analysis. This is followed by experiments demonstrating the practicality of our methods before the paper concludes in Section 8.

2 Related Work

The correlations based change analysis bears closest resemblance to the anomaly event detection work of Akoglu and Faloutsos [4]. Both our technique and that of Akoglu and Faloutsos employ 'Pearson ρ' correlation matrix manipulations. The aggregation methods to compute vertex and graph level correlation values are also similar. However, only correlations between vertices in/out degrees are considered by Akoglu and Faloutsos, whereas our method can be adapted to examine other vertex-induced k hop subgraph correlations as well – e.g. diameter, number of triangles, centrality, or other traffic distribution subgraph metrics [2,9]. The other key difference between our methods lies with their intended application usages.

Whilst their method exposes significant graph-wide deviations, employing correlation solely for change detection suffers from some shortcomings. Besides detecting change from the majority of network nodes, we are also interested in other types of network changes, such as anomalous deviations in behaviours from a few (or single) dominant vertices. In this sense, our approach is not to deploy correlations for network change detection directly, but to aid existing change detection methods and extend subsequent change analysis.

In another related paper from Akoglu and Dalvi [5], anomaly and change detection using similar correlation methods from [4] is described. However, their technique is formalised and designated for detecting 'Eigenbehavior' based changes only. In comparison, our methods are general in nature and not restricted to any particular type of network change or correlation outcome.

Another relevant paper from Ide and Kashima [6] is their Eigenspace inspired anomaly detection work for web-based computer systems. Both our approach and [6] follow similar procedural steps. But whilst our correlation method involves a graph adjacency matrix populated and aggregated with simplistic correlation computations, the technique in [6] employs graph dependency matrix values directly and consists of complex Eigenvector manipulations.

Other areas of research related to our work arise from Ahmed and Clark [7], and Fukuda et. al. [8]. These papers describe change detections and correlations that share similar philosophy with our methods. However, their underlying change detection and correlation methods, along with the type of network data differ from our approach.

The remaining schemes akin to our correlation methodology are captured by the MetricForensics tool [9]. Our technique span multiple levels of the network graphs, in contrast, MetricForensics applies correlation analysis exclusively at a global level.

3 Scan-Statistics

This section summaries the scan-statistics method. For a full treatment of scan-statistics, we refer the reader to [3]. Scan-statistics is a change detection technique that applies statistical analysis to sub-regions of a network graph. Statistical analysis is performed on graph vertices and vertex-induced subgraphs in order to measure local changes across a time-series of graphs. Whenever the network undergoes significant global change, scan-statistics detects and identifies the network vertices (or subgraphs) which exhibited greatest deviation from prior network behaviours.

In scan-statistics, local graph elements are denoted by their vertex induced k-hop subgraph regions. For every k-hop subgraph region, a *vertex-standardized* locality statistic is measured for that region. In order to monitor changes experienced by these subgraph regions, their locality statistics are measured for every graph throughout the time-series of network graphs. The vertex-standardized statistic $\tilde{\Psi}$ is :

$$\tilde{\Psi}_{k,t}(v) = \frac{\Psi_{k,t}(v) - \hat{\mu}_{k,t,\tau}(v)}{\max(\hat{\sigma}_{k,t,\tau}(v), 1)} \tag{1}$$

where k is the number of hops (edges) from vertex v to create the induced subgraph, v is the vertex from which the subgraph is induced from, t is the time denoting the time-series graph, τ is the number (*window*) of previous graphs in the time-series to evaluate against current graph at t, Ψ is the local statistic that provides some measurement of behavioural change exhibited by v, and μ and σ are the mean and variance of Ψ.

The vertex-standardized locality statistic equation (1) above is interpreted as follows. For the network graph at time t, and for each k-hop vertex v induced subgraph, equation (1) measures the local subgraph change statistic Ψ in terms of the number of standard deviations from prior variations.

With the aid of equation (1), scan-statistics detects any subgraph regions whose chosen behavioural characteristics Ψ deviated significantly from its recent history. By applying equation (1) iteratively to every vertex induced subgraph, scan-statistics uncovers local regions within the network that exhibit the greatest deviations from their expected behaviours.

With scan-statistics change detection, typically the subsequent change analysis focuses on individual vertices (or subgraphs) that exhibited greatest deviation from their expected prior behaviours only. Our scan-statistical correlation method bridges this gap by examining all regions of change within the network and their relationships with one another using correlation analysis.

4 Scan-Statistical Correlations

In order to examine correlations between local network changes uncovered by scan-statistics, we use Pearson's ρ correlation. We examine and quantify possible relationships in local behavioural changes between every pair of vertex induced k-hop subgraphs in the network. For every pair of vertices v_1 and v_2 induced subgraphs, we extend scan-statistics with correlation computations using Pearson's ρ equation :

$$\rho_{k,t,\tau}(v_1,v_2)=\frac{\sum_{t'=t-\tau}^{t-1}(\Psi_{k,t'}(v_1)-\hat{\mu}_{k,t,\tau}(v_1))(\Psi_{k,t'}(v_2)-\hat{\mu}_{k,t,\tau}(v_2))}{\sqrt{\sum_{t'=t-\tau}^{t-1}(\Psi_{k,t'}(v_1)-\hat{\mu}_{k,t,\tau}(v_1))^2}\sqrt{\sum_{t'=t-\tau}^{t-1}(\Psi_{k,t'}(v_2)-\hat{\mu}_{k,t,\tau}(v_2))^2}} \qquad (2)$$

where k, v, t, τ, Ψ, and $\hat{\mu}$ are defined the same as for (1).

The scan-statistical correlation scheme is outlined in Fig. 1. For every network graph in the time-series, correlations between local vertex (or subgraph) changes are computed according to corresponding vertex behaviours from the recent historical window τ of time-series graphs. The raw correlations data are then populated into an $n\times n$ matrix of n vertices from the network graph. This matrix provides a simplistic assessment of positive, low, or possibly opposite correlations in change behaviours.

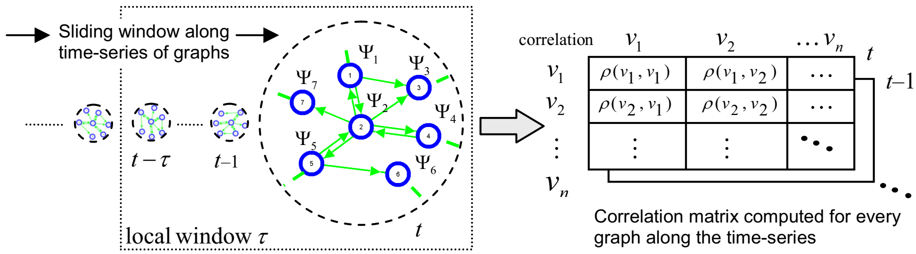

Fig. 1. Correlation is computed for each time-series graph and populated into a matrix

5 Multi-level Correlations Analysis

5.1 Aggregation of Correlation Data

To facilitate analysis of correlations amongst behavioural changes at higher network graph levels, the raw correlation data (i.e. correlation matrix in Fig. 1) is aggregated into other representative results. A number of aggregation schemes were examined. However, compared to basic aggregation methods, spectral, Perron-Frobenius, Eigenvector, and other matrix-based methods did not present any additional benefits and took longer processing times. Hence, for the remainder of this paper, we restrict our discussions to aggregation schemes employing straightforward averaging.

The aggregation of correlation data is described in Fig. 2. In the first step, for each vertex, the vertex's correlation with every other vertex is aggregated together. The outcome is to provide an overall correlation measure for every vertex against the majority of other vertices throughout the network. From the perspective of an individual vertex, this aggregated correlation value indicates if the behavioural change experienced by that vertex is also exhibited by the majority or only a small number of other vertices (discussed further in Section 5.3).

In the second step, using individually aggregated vertex correlation values from Step 1, an overall correlation measure is acquired for the network graph. The network graph correlation indicates if the change experienced by the network is part of a broader graph-wide change, or if the network deviation is due to few local regions.

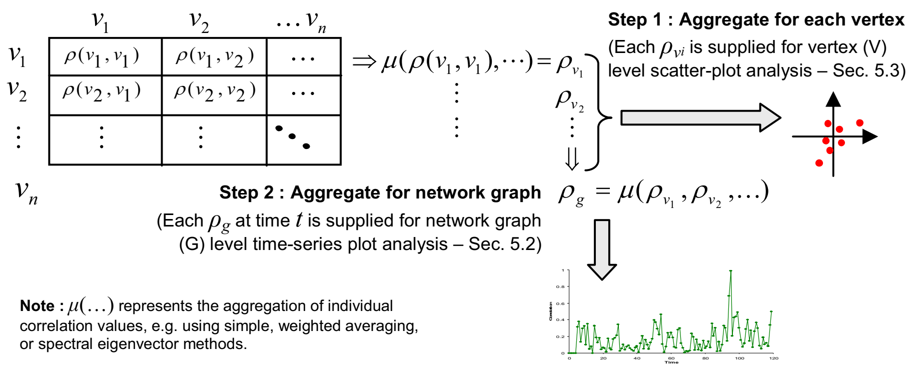

Fig. 2. Correlation data is aggregated to provide results for each individual vertex and the graph

5.2 Global Network Graph Correlation (G)

Using the aggregated correlations values ρ_g, the change in *global* graph correlation levels across the time-series of network graphs can be examined. This enables network analysts to monitor for patterns between the network deviations uncovered from a change detection time-series plot against corresponding graph-wide correlation from the correlation time-series plot. For instance, are significant network deviations due to widespread changes throughout the graph at multiple change-points? In the case of high correlation, this indicates a large majority of network subgraph regions exhibit similar degree of change in network behaviour (as demonstrated in Section 7).

5.3 Vertex Level Correlation (V)

Beneath the global graph level, the aggregated correlation of every vertex to all other vertices is analysed. The aggregated correlation ρ_{vi} is acquired via Step 1 in Fig. 2. For every vertex that undergoes significant change, the aggregated vertex correlation indicates if the change experienced by that vertex is exhibited by the majority or only a few vertices. A high correlation indicates that the change by the vertex was also experienced by other vertices as well. On the other hand, a low correlation signifies the change was likely restricted to that vertex only.

From a change-analysis perspective, besides simply identifying which vertices contributed most to the network change (as per conventional scan-statistics), using vertex level correlations, further insight regarding how these individual deviations relate to the wider network can be deduced. Correlation inspired visualisation techniques may then be employed to observe these changes.

Scatter-Plot Visualisation

To effectively analyse vertex level correlations, a scatter-plot visualisation scheme is employed. The scatter-plots reveal how an individual vertex, groups of vertices, and the overall network vary from one time interval to the next. The two types of scatter-plots are : (i) a scatter-plot of scan-statistic locality deviation value $\tilde{\Psi}_{(v)}$ of every vertex, and (ii) the scatter-plot of aggregated correlation ρ_v of every vertex.

Fig. 3 summarises our scatter-plot concept. To examine how certain vertices of interest vary over time, scatter-plots are created for network graphs that undergo significant network change. For each of these network graph change-points, the deviation (or correlation) results of every vertex from the previous and current graphs (i.e. graphs before and at the change-point) are plotted on the scatter-plot.

On the scatter-plot, every vertex is plotted as a xy-coordinate point. The y-axis represents the vertex deviation statistic or correlation value held by the vertex in the previous graphs, and the x-axis value corresponds to the vertex deviation or correlation of the current graph under examination.

By examining where individual or clusters of vertices appear on the scatter-plots, the vertices that experienced most significant changes are easily identified, and the types of changes can be inferred immediately. Dynamically changing scatter-plots, whereby a scatter-plot is displayed for consecutive time-series graphs also reveal how specific vertices of interest or clustered changes transpire over time.

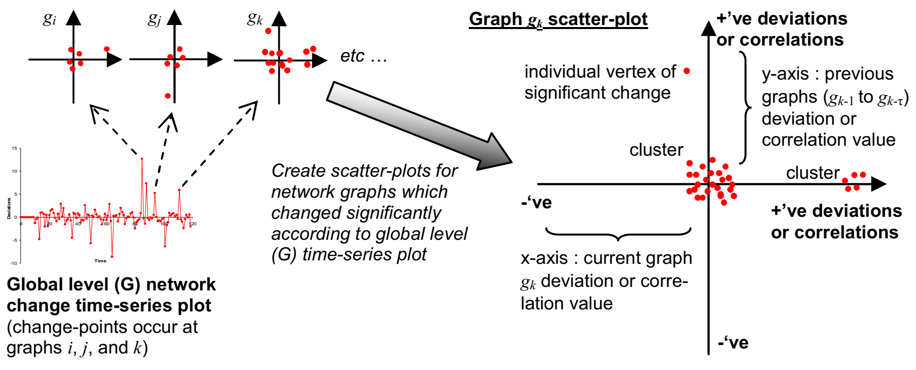

Fig. 3. Concept of vertex (V) level scatter-plots created for significantly deviated graphs

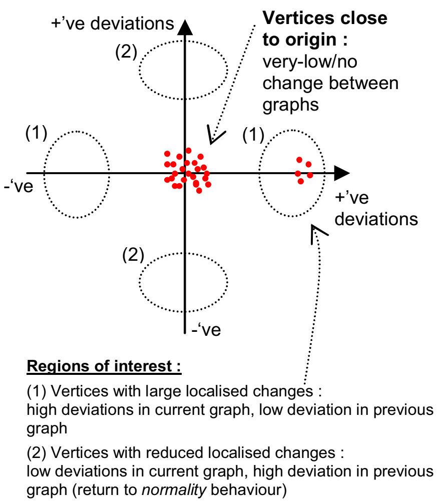

Fig. 4. Vertex (V) level deviation scatter-plot

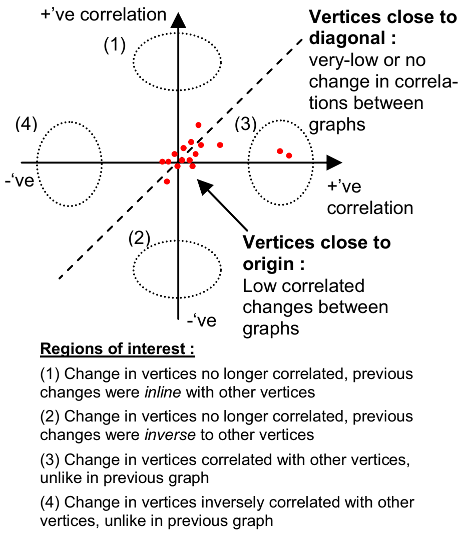

Fig. 5. Vertex (V) level correlation scatter-plot

For instance, the above concept for the scan-statistic deviation scatter-plot is shown in Fig. 4. This scatter-plot reveals the change in edge-connectivity (k=0 hop) of vertices; in particular, the extent of edge-connectivity changes and clustering of vertices with similar connectivity variations.

Various *regions of interest* are identified on this scatter-plot and described in Fig. 4. We observe which regions vertices fall into and how these vertices shift across the scatter-plots throughout the time-series of graphs. This allows the connectivity deviations of every vertex (and clusters) relative to the wider network to be examined.

For the correlations scatter-plot (Fig. 5), various similar regions of interests are also identified. This scatter-plot reveal if specific vertices acted alone in exhibiting localised changes, or if their behavioural deviations were part of a collective network-wide change. For example, do single, multiple, clustered vertices exhibit similar or independent deviations in edge-connectivity from prior network behaviours?

5.4 Vertex-to-Vertex Correlation (V×V)

At the lowest level, the raw correlation data matrix in Fig. 1 allows for the examination of individual vertex-to-vertex correlations in change behaviours. A possible use-case of such vertex-to-vertex correlations monitoring is to detect highly similar or duplication of behaviours from multiple vertices. For example, the discovery of a vertex in the network attempting to falsely mimic the characteristics and assume the identity of another legitimate network vertex. Over a sufficient period of time, if the correlations between two vertices are suspiciously maintained as highly correlated, then this may indicate the presence of illegal vertex imposters.

6 A Multi-level Network Change Analysis Scheme

In this section, we bring together the different levels of correlations data and analysis above to outline a scheme for examining network changes uncovered by scan-statistics. The approach is depicted in Fig. 6 as a flow diagram. It involves examining various forms of network correlation data, to gather evidence for establishing possible scenarios and the context in which network changes were triggered.

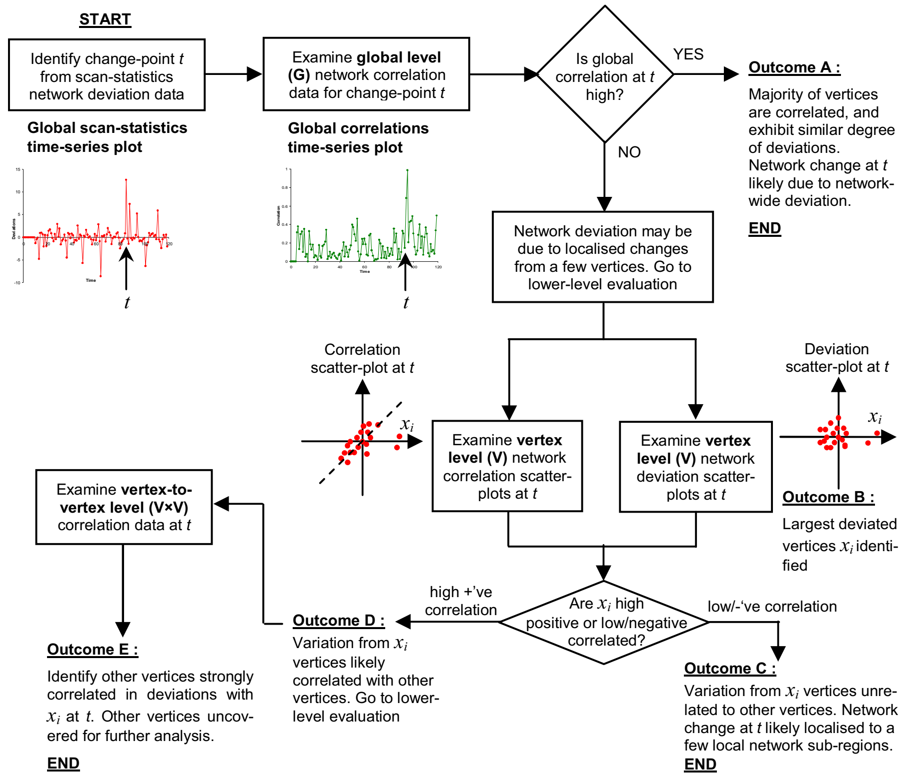

Fig. 6. Multi-level network change analysis process using scan-statistics correlations data

7 Experiments

The scan-statistical correlations techniques have been applied to a variety of network datasets, both internally within and external to our organisation. For demonstration purposes in this paper, we make use of the Enron email dataset [3]. The dataset contains emails from 150 Enron employees including those of senior management between 1999 and 2001, covering events leading up to the collapse of the company. A network graph is extracted from email transactions during each week resulting in a time-series of 120 graphs. In each graph, an edge is established between two vertices if at least one email was sent between them during that weekly period.

Global (G) Level Network Deviations and Correlations

We adopt the multi-level correlations change analysis outlined in Fig. 6. For scan-statistics and correlations evaluations, the *k*=0 hop edge connectivity deviations of every vertex is computed – i.e. the change in number of new emails between individuals and their overall emailing connectivity is our main focus. From preliminary test runs, the window τ of graphs to establish statistical edge-connectivity mean and compute correlations was set at a size of 5.

The global level (G) scan-statistical deviation and correlation time-series plots are shown in Fig. 7 and 8. The x-axis corresponds to network graph change-points at weekly intervals, and the y-axis is the number of deviations or correlation level.

Fig. 7 reveals a number of positive and negative change-spikes when emailing edge-connectivity rose excessively or dropped sharply from prior weeks. These change-spikes indicate when significant increases or decreases in emails were sent/received amongst individuals from one week to another when compared to the recent history of expected emailing behaviours.

The interesting period occurred between weeks 64 and 101. We focus on three change-points, weeks 85, 88, and 94, when new emailing activity escalated significantly. Weeks 85 and 88 corresponds to the periods before and after the resignation of the Enron CEO; who oversaw the price-fixing and illegitimate practices of the company. The week leading up to the investigations by the Securities and Exchange Commission (SEC) into Enron's operations is also highlighted during week 94.

Besides uncovering when significant change events occurred, the scan-statistics global time-series plot does not reveal much. For instance, were these change-points triggered by a single individual vertex, a few or clusters of vertices, or collectively by a majority of the network graph nodes? In Fig. 8, the global (G) correlation time-series plot is the first step in examining how emailing behaviour deviates collectively and individually across all Enron employees.

At weeks 85 and 88, the corresponding global correlation in emailing deviations is low. This suggests the large emailing deviations indicated by Fig. 7 are not widespread. This is not surprising given that the planned stepping down of a CEO concerns highest level executives. As expected, the emailing behaviour of the majority of employees does not change from previous weeks. At week 94, the global correlation is higher, indicating the emailing network change may involve a larger population of employees. To assess this possibility and identify individual vertices that triggered the network change, the emailing behaviour at the vertex level (V) are examined next.

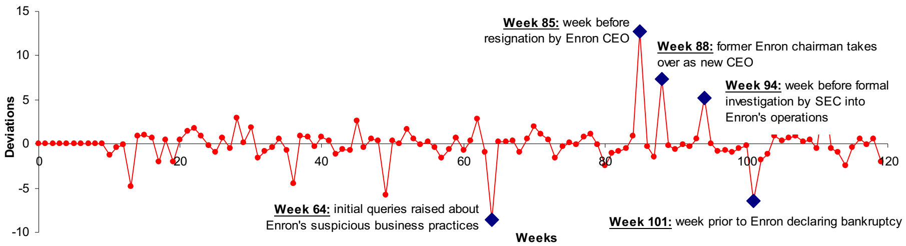

Fig. 7. Global (G) deviation time-series plot : Enron dataset

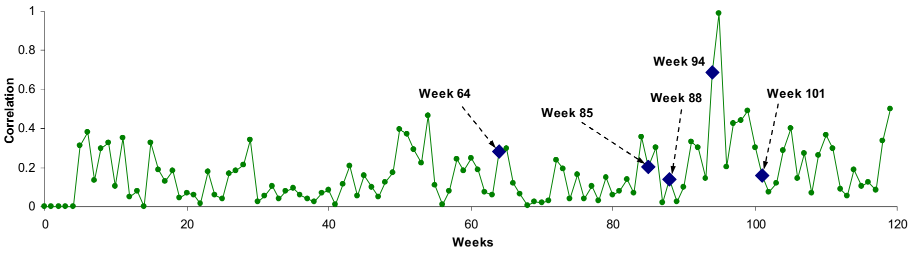

Fig. 8. Global (G) correlation time-series plot : Enron dataset

Vertex (V) Level Network Deviations and Correlations

Fig. 9, 10 and 11 show the scan-statistics deviation and correlation scatter-plots. The deviation units are normalised such that deviations on the scatter-plot are relative to the maximum deviation exhibited from the vertex that deviated the most during the specified week.

For week 85, Fig. 9(a) shows that vertex 118 (and 60) exhibit the largest individual deviation. Their location at the extreme positive x-axis and near-zero y-axis region indicates their emailing activity increased up to 16 times from the prior weeks of low emailing. Relative to these two vertices, the deviation in emailing activity of other vertices remained low, all being clustered around the origin.

Tracing back our vertex labelling, it is no surprise that vertex 118 is the Enron CEO (and vertex 60 is a president of Enron Online Business). Prior to stepping down, it is usual for the CEO to send various emails to large groups of individuals to close off remaining official duties or resolve other matters (e.g. even a company-wide email announcing his resignation). The CEO may also receive many new emails from other employees offering farewell messages. In an organisation such as Enron, even a small proportion of such emailing activity for one individual would cause their emailing edge connectivity to spike up.

However, with vertex 60, suspicious reasoning behind the concurrent spike in emailing with vertex 118 may exist. The Enron online business was a major part of Enron's suspicious practices, hence it would be interesting to establish what anomalous relationship exists between the CEO and the online business president.

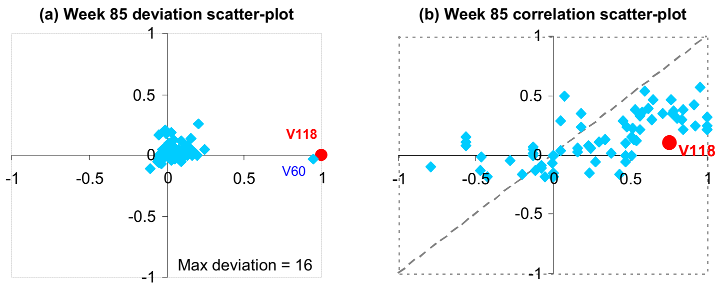

Fig. 9. Vertex (V) level scatter-plots : Enron dataset – Week 85

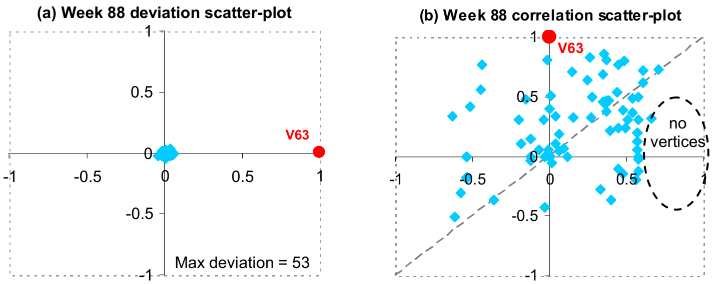

Fig. 10. Vertex (V) level scatter-plots : Enron dataset – Week 88

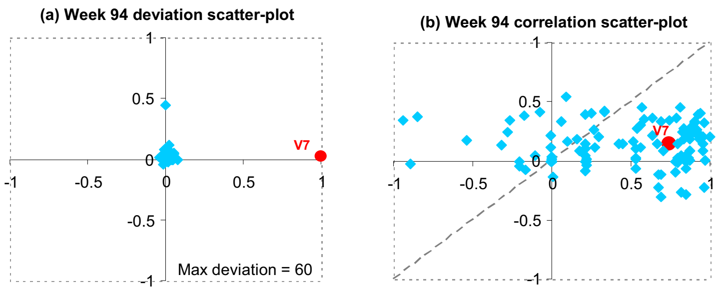

Fig. 11. Vertex (V) level scatter-plots : Enron dataset – Week 94

For week 88, Fig. 10(a) shows the extreme emailing deviation is entirely due to vertex 63, who happens to be the former Enron chairman that took over the new CEO role during that week. For week 94, Fig. 11(a) shows that the network change was dominated by vertex 7, the Enron chief operating officer (COO). These scan-statistics scatter-plots identify vertices that deviated vertices most, as per Outcome B in Fig. 6.

From our deviation scatter-plots, it is clear emailing deviation is dominated by a single (or two) vertex. Deviations of other vertices are much smaller, concentrated at the scatter-plot origin. But it remains beneficial to examine if deviations from the rest of the network are correlated (or disparate) with the dominant single vertex deviation.

From an overall network perspective, the majority of graph vertices on the week 88 correlation scatter-plot (Fig. 10(b)) do not reveal any discernable pattern. However, we make two key observations. First, no vertices are located on the high positive x-axis region; and second, vertex 63 is shown to hold low correlation during week 88 but high correlation previously (i.e. at low x-axis and high positive y-axis). This signifies the extreme deviation of vertex 63 is limited to itself (i.e. Outcome C in Fig. 6).

Previously, vertex 63's emailing was consistent with other vertices, but in week 88, the new Enron CEO was acting alone with its excessive emailing activity. Such behaviour may be justified given that a new CEO would send/receive many more emails as new responsibilities are taken up under his control. Given the new CEO was a former Enron board chairman and may not be involved with day-to-day activities of Enron, the sudden and excessive spike in emailing makes sense. But if such emailing behaviour was detected for other vertices, then this could be deem anomalous.

The correlation scatter-plot in Fig. 11(c) shows the largest deviated vertex 7 (Enron COO) was previously uncorrelated (low y-axis previous graphs correlation), but its extreme deviation is correlated with other vertices during week 94 (high x-axis current graph correlation). In addition, a significant portion of vertices are also located on the low y-axis and higher positive x-axis region. This suggests wider network change occurred at week 94, and vertex 7 was not acting alone.

Whilst other vertices may not have deviated as excessively as vertex 7, they still exhibited the same positive degree of new emailing connectivity as the Enron COO. These vertices may represent employees working closely together to respond to the pending SEC investigations; of which, the Enron COO was likely the coordinator for the company. Intuitively, this accounts for the large spike in new emailing deviation from vertex 7, and associated deviation from other vertices. For week 85, similar correlation analysis as week 94 can be deduced.

8 Conclusions

This paper presented a correlations based network change analysis technique. The correlations between network graph deviations amongst vertices are examined in order to gain greater insight into different scenarios triggering the network changes. We extend scan-statistical change detection with correlations inspired network analytics operating at multiple levels of the network graph. Whilst our technique is beneficial for any networks in general, the Enron dataset was used to demonstrate valuable outcomes using correlations analysis. In the future, our research shall focus on applying higher order statistics to acquire other in-depth network change correlation results.

References

1. Shoubridge, P., Kraetzl, M., Ray, D.: Detection of Abnormal Change in a Time-Series of Graphs. Journal of Interconnection Networks 3(1-2), 85–101 (2002)
2. Bunke, H., Dickinson, P., Kraetzl, M., Wallis, W.: A Graph-Theoretic Approach to Enterprise Network Dynamics, Birkhauser (2007)
3. Priebe, C., Conroy, J., Marchette, D., Park, Y.: Scan Statistics on Enron Graphs. Computational and Mathematical Organization Theory 11(3), 229–247 (2005)
4. Akoglu, L., Faloutsos, C.: Event Detection in Time-series of Mobile Communication Graphs. In: 27th Army Science Conference, pp. 1–8 (2010)
5. Akoglu, L., Dalvi, B.: Structure, Tie Persistence and Event Detection in Large Phone and SMS Networks. In: Mining & Learning with Graphs, pp. 10–17. ACM (2010)
6. Ide, T., Kashima, H.: Eigenspace-based Anomaly Detection in Computer Systems. In: Knowledge Discovery and Data Mining Conference, pp. 440–449. ACM (2005)
7. Ahmed, E., Clark, A.: Characterising Anomalous Events Using Change-Point Correlation on Unsolicited Network Traffic. In: Secure IT Systems, pp. 104–119. Springer (2009)
8. Fukuda, K., Hirotsu, T., Akashi, O., Sugawara, T.: Correlation among Piecewise Unwanted Traffic Time-series. In: IEEE Global Telecommunications Conference, pp. 1–5 (2008)
9. Henderson, K., Eliassi-Rad, T., Faloutsos, C., Akoglu, L., Li, L., Maruhashi, K., Prakash, B., Tong, H.: Metric Forensics: a Multi-level Approach for Mining Volatile Graphs. In: Knowledge Discovery and Data Mining Conference, pp. 163–172. ACM (2010)