Directional Laplacian Centrality for Cyber Situational Awareness

SINAN G. AKSOY, EMILIE PURVINE, and STEPHEN J. YOUNG,

Pacific Northwest National Laboratory

Cyber operations is drowning in diverse, high-volume, multi-source data. To get a full picture of current operations and identify malicious events and actors, analysts must see through data generated by a mix of human activity and benign automated processes. Although many monitoring and alert systems exist, they typically use signature-based detection methods. We introduce a general method rooted in *spectral graph theory* to discover patterns and anomalies without *a priori* knowledge of signatures. We derive and propose a new graph-theoretic centrality measure based on the derivative of the graph Laplacian matrix in the direction of a vertex. To build intuition about our measure, we show how it identifies the most central vertices in standard network datasets and compare to other graph centrality measures. Finally, we focus our attention on studying its effectiveness in identifying important IP addresses in network flow data. Using both real and synthetic network flow data, we conduct several experiments to test our measure's sensitivity to two types of injected attack profiles and show that vertices participating in injected attack profiles exhibit noticeable changes in our centrality measures, even when the injected anomalies are relatively small, and in the presence of simulated network dynamics.

CCS Concepts: • **Security and privacy** → **Intrusion/anomaly detection and malware mitigation**; • **Mathematics of computing** → **Spectra of graphs**;

Additional Key Words and Phrases: Cyber network flow, graph centrality, spectral graph theory

ACM Reference format:

Sinan G. Aksoy, Emilie Purvine, and Stephen J. Young. 2021. Directional Laplacian Centrality for Cyber Situational Awareness. *Digit. Threat.: Res. Pract.* 2, 4, Article 28 (October 2021), 28 pages.
<https://doi.org/10.1145/3450286>

1 INTRODUCTION

Cyber situational awareness comprises many tasks from understanding the threat landscape and network vulnerabilities to real-time network monitoring [27]. Endsley describes general situational awareness of dynamic systems in terms of three levels of awareness: perception, comprehension, and projection (or prediction) [22]. To achieve any of perception, comprehension, or prediction, analysts must navigate a variety of data sources and a deluge of data within each source to get an understanding of current operational posture. Fusing information derived from these data sources in a way that is interpretable to a human analyst provides a holistic situational awareness. In this article, we focus on a specific aspect of perception: network anomaly attribution.

One of the difficult problems of perception in cyber systems is the ability to see through real-time network traffic data produced by benign automated processes and normal human activity to pick out abnormal, potentially malicious, events, and actors. Network traffic data is often messy and massive, requiring assistance from

Authors' addresses: S. G. Aksoy and E. Purvine, 1100 Dexter Ave N., Suite 500, Seattle, WA 98109; emails: {sinan.aksoy, emilie.purvine}@pnnl.gov; S. J. Young, 902 Battelle Blvd, PO Box 999, Richland, WA 99352; email: stephen.young@pnnl.gov.

This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

© 2021 Copyright held by the owner/author(s).

2576-5337/2021/10-ART28 \$15.00

<https://doi.org/10.1145/3450286>

Fig. 1. Example spectral timeline with scan-like anomaly around minute 30.

algorithms or analytics to provide the human analyst with tips and cues for further investigation. Many of the existing monitoring and alert systems use signature-based detection techniques and are highly tailored to specific data types. To detect zero-day attacks and other previously unseen tactics, techniques, and procedures, a signature-agnostic anomaly detection strategy must be utilized. Broadly speaking, two driving challenges are: (1) identifying the presence of anomalous or adversarial behavior within the normal background variation of network data; and (2) identifying which actors or agents within a system are participating in the anomalous behavior. In this regard, a *mathematically grounded* approach is helpful to explore data and discover patterns and anomalies without prior knowledge of behaviors of interest. While there are approaches that attempt to simultaneously address both challenges [39, 41], in this work, we focus on developing a framework for addressing anomaly attribution using a method based on the *spectrum* of the network's graph structure. In addition to these two driving challenges there is also the question of determining maliciousness of an anomaly. The work we propose here is not meant to answer this question but rather to identify the actors or agents participating in anomalous behavior. Subsequent packet inspection and other investigation activities would be necessary to determine the maliciousness of the anomaly.

As a proxy for network structure, and to provide grounding for this work in commonly available network data, we will focus on network flow, a summary of computer to computer communications across the network. The method we propose is applicable to any data with records of the form (*source, destination, time, metadata*), not just within the cyber domain, although interpretation of the method's output will vary by data type. Our approach uses a graph to model a set of network flow records, typically from a small time window. In a spectral approach the structural properties of network communication, as reflected in a graph derived from flow, are analyzed using eigenvalues of associated matrices; these eigenvalues are known as the *graph spectrum*. The temporal sequence of flow data is then analyzed using a sequence of small time windows, each on the order of a few minutes. The spectrum of each small time window is computed separately providing a numerical graph spectrum vector for each window. When stitched together over time this creates a timeline or heartbeat of the network flow data that can be inspected by an analyst to look for patterns and deviations from those patterns. See Figure 1 for an example visualization of synthetic network flow with a scan-like anomaly around minute 30. While this tool has shown promise in helping to identify “low and slow” scans (that is, scans that are not high volume all at once, but rather distributed across a period of time) as well as frequent recurrent behavior in operational data, the process requires a human in the loop to inspect the timeline and identify the actor responsible for the spectral anomaly, which can be error-prone.

This article proposes a tool that we call the Directional Laplacian Centrality to identify “important” **Internet Protocol (IP)** addresses within each time window and focus analyst attention from the broad temporal awareness that the spectrum itself may provide to more specific attribution of anomalies. The underlying philosophy is that an IP is important if it has a significant (measurable) contribution to the spectrum of the associated graph.

Anomalies can then be characterized as times in which a previously unimportant IP becomes significantly more important. After defining the Directional Laplacian Centrality, we will test the method within real and synthetic data, analyzing how the importance of IP addresses change in response to planted attack profiles.

We organize our work as follows: In Section 2, we provide the necessary preliminaries to introduce network flow data and spectral graph theory, justify the advantages of taking a spectral approach, and provide a literature review of other relevant importance measures. In Section 3, we derive and define our proposed importance measure, Directional Laplacian Centrality. In Section 4, we provide some intuition through examples of Directional Laplacian Centrality applied to well-known graph datasets. In Section 5, we describe our specific test data, experimental setup, and report our findings.

2 PRELIMINARIES

A graph $G = (V, E)$ is a set V of elements, called *vertices*, and a set E of unordered pairs of vertices, called *edges*. If $\{u, v\} \in E$, then we say u and v are *adjacent* and write $u \sim v$. We call $\{v \in V : u \sim v\}$ the *neighborhood* of u , and $d_u = |\{v \in V : u \sim v\}|$ the *degree* of u . A set of vertices $S \subseteq V$ is *connected* if for any $u, v \in S$, there exists a sequence of adjacent vertices $u, \dots, v$. A maximal connected subset (one that is not strictly contained within a larger subset) is called a *connected component* and its *size* is the number of vertices in the component, $|S|$. For other basic graph theory terminology, readers are referred to Reference [6].

In this work, for ease of exposition and simplicity, we focus on unweighted, undirected graphs. In most applications, however, the edges of the graph may have natural weight or directionality information. The generalizations of the techniques developed in Section 3 to weighted and directed cases (as appropriate) is a straightforward exercise.

2.1 Network Flow Data

Network flow summarizes data exchanged between pairs of IP addresses in a network. A single network flow is an aggregation of multiple packets that occur within a small time window and have the same source IP, destination IP, source port, destination port, and protocol. There is significant subtlety and engineering that goes into packet capture and flow aggregation (for example, even identifying which IP is source and which is destination is not trivial), as seen in one of the seminal papers [19] and the more recent Reference [26]. However, for this article it is enough to think of a single flow as a record that contains a source and destination IP, and a time stamp. Other metadata are present in a record (e.g., bytes, packets, ports, and protocol) that one could use to weight or filter the data. Initially created for accounting and usage profiling, network flow is a prevalent source of data for cyber situational awareness. Advances in flow monitoring are continuing to enable this usage [28]. A recent paper of Moustafa, Hu, and Slay provides a thorough survey of network anomaly detection approaches using the paradigm of feature identification followed by classification, clustering, machine learning, rule-based, or statistical methods [38]. Our approach fits this paradigm with features derived from a graph interpretation of the network flow data followed by a statistical anomaly detection based on those derived numerical features.

A set of network flow can be modeled as a graph in which vertices represent IP addresses and an edge $\{u, v\}$ indicates that a flow between IPs u and v is contained in the set. Although one could restrict to network flows with specific ports or protocols for a more targeted analysis, in our experiments, we will not filter or label vertices and edges by ports or protocols but will consider all flows. To construct each graph, we restrict to a specific time interval to capture the structure of communications within a given time period. However, network flow gathered from even a few hours typically creates a graph much too large to yield helpful insights, so to represent time and capture structural evolution, we consider a dynamic graph model. A *dynamic graph* is a sequence of graphs $\{G_t\}_{t=1}^N$ where t denotes a specific time, either instantaneous or a short time window, and $G_t = (V_t, E_t)$. Dynamic graphs in which t represents an instantaneous time can be difficult to analyze, since at any given instance the graph may be quite sparse and fragmented; and in the case of network flow, they can evolve very rapidly, as some flows have duration on the order of milliseconds. Instead, we consider small time windows, on the order

Fig. 2. From left to right: a graph, its normalized Laplacian matrix, and normalized Laplacian eigenvalues. Spectral graph theory establishes relationships between eigenvalues and the graph they are derived from, represented in the diagram by the dashed line.

of 30 seconds to 10 minutes, and gather all flow records that overlap the time window into a single graph, G_t . For a dynamic graph created from network flow an edge exists in G_t between two IPs when a flow occurs with one as source and the other as destination within the time window indicated by t . The vertices of G_t are IP addresses present in at least one flow within the time window. Many of the examples in this article will be on a fixed snapshot of network flow data, a single G_t (see Section 5 for details of that graph).

2.2 Spectral Graph Theory

Our methods for identifying potential anomalous actors within networks are rooted in *spectral graph theory*. This area of mathematics studies properties of a graph through the eigenvalues and eigenvectors of matrices associated with the graph, as illustrated in Figure 2. There are a variety of different matrices one could associate with a graph. Perhaps the most basic and well-known, the adjacency matrix A of an n -vertex graph $G = (V, E)$, is an $n \times n$ matrix where

$$A_{ij} = \begin{cases} 1 & \text{if } \{i, j\} \in E \\ 0 & \text{otherwise} \end{cases}.$$

Two other important matrices, which will be the focus of this work, are the combinatorial Laplacian matrix, L , and normalized Laplacian matrix, $\mathcal{L}$:

$$L = D - A,$$

$$\mathcal{L} = D^{-1/2} L D^{-1/2},$$

where A is the adjacency matrix defined above and D is the diagonal matrix with vertex degrees on the diagonal, i.e., $D = \text{diag}(A\mathbf{1})$, where $\mathbf{1}$ is the appropriately sized all-ones vector. The matrices A , L , and $\mathcal{L}$ are all symmetric, and hence have real eigenvalues that we label in increasing order:

$$\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n.$$

When unclear from context, we specify the matrix underlying a particular eigenvalue by writing, for example, $\lambda_i(L)$ denotes that λ_i is the i th eigenvalue of the combinatorial Laplacian.

Unlike A , both L and $\mathcal{L}$ are positive semi-definite, and therefore have non-negative eigenvalues. The eigenvalues, or *spectrum*, of L and $\mathcal{L}$ characterize a number of important properties not captured by the adjacency eigenvalues. Particularly pertinent to our analyses, Laplacian spectra capture various graph connectivity and neighborhood expansion properties. Loosely speaking, such properties quantify the extent to which a graph (or subsets of vertices within a graph) are well-connected. For instance, an elementary fact is that the multiplicity of the eigenvalue 0 of both L and $\mathcal{L}$ equals the number of connected components of the graph. Furthermore, the second eigenvalue λ_2 of a connected graph quantifies the extent to which that graph is well-connected, in several different regards.

In the case of the combinatorial Laplacian L , this second eigenvalue is referred to as *algebraic connectivity*. This quantity is related to the graph's vertex connectivity, the minimum number of vertices that must be deleted to disconnect a graph. Furthermore, the corresponding eigenvector, called Fiedler's vector, is frequently used

to partition a graph into well-connected groups; see Reference [23] for more. In the case of the normalized Laplacian, the second eigenvalue captures neighborhood expansion properties, which measure how many edges leave a set of vertices, relative to the “volume” of that set. One way of quantifying such expansion is via Cheeger’s constant. More precisely, for a vertex subset $X \subseteq V$, letting $e(X, \bar{X})$ denote the number of edges between X and its complement $\bar{X}$ and defining the volume of a set as $\text{vol}(X) = \sum_{i \in X} d_i$, the *Cheeger ratio* of X is

$$\Phi(X) = \frac{e(X, \bar{X})}{\min\{\text{vol}(X), \text{vol}(\bar{X})\}}.$$

The *Cheeger constant* of the graph $\Phi(G)$ is the minimum Cheeger ratio over all vertex subsets. Via Cheeger’s inequality [8], $\lambda_2(\mathcal{L})$ serves as an approximation of the Cheeger’s constant,

$$2\Phi(G) \geq \lambda_2(\mathcal{L}) \geq \frac{\Phi(G)^2}{2}.$$

In this sense, the second eigenvalue controls graph neighborhood expansion properties. Further, just as the aforementioned Fiedler vector may be used to partition a graph, the normalized Laplacian eigenvector corresponding to $\lambda_2(\mathcal{L})$ can be used to partition the graph as well.

While the aforementioned connectivity properties were key to the early development of spectral graph theory, researchers have shown a wide variety of graph structural properties are encoded in eigenvalues. For example, vertex degrees [20], average shortest path length [37], diameter [17], chromatic number [52], independence number [24], number of spanning trees [5], network flows¹ and routing [3], and random walk mixing time and associated parameters [2], can all be bounded, controlled, or characterized in terms of eigenvalues. This body of research attests to eigenvalues of the aforementioned matrices serving as a far-ranging tool for capturing graph properties.

2.3 Importance Measures

One approach towards identifying agents participating in anomalous behavior within a graph is to understand their role within the graph structure relative to other agents and how that role changes over time. In this context, graph-theoretic importance measures provide different ways of quantifying an element’s role in the system. Tracking the importance over time, through a dynamic graph sequence, can identify anomalous changes in behavior or role. Here, we briefly review examples of importance measures to elucidate the breadth of properties they can capture, and to place our proposed importance measure within the literature.

Broadly speaking, one class of graph importance measures are based on the shortest-path structure in the graph. For example, *closeness centrality* ranks vertices based on their average shortest path length to other vertices in the graph: Vertices in close proximity to the rest of the graph via shorter paths are ranked more highly than vertices on the periphery. *Betweenness centrality* measures the importance of a vertex based on the frequency of its occurrence *en route* between other pairs of vertices. Accordingly, vertices that belong to many short paths linking other pairs of vertices are ranked highly. Rather than restrict attention to *shortest* paths, another related class of importance measures are based more generally on *walks*, which may be non-minimal in length. Since a pair of vertices may be linked by infinitely many walks of arbitrary length, such measures tend to rely on asymptotic expressions, such as the limiting distribution of a random walk. For instance, the well-known PageRank algorithm assigns scores to vertices based on the stationary distribution of a modified random walk. Katz centrality considers the total number of walks from a given vertex to the rest of the graph, penalizing longer walks with an attenuation factor.

¹The term “network flow” here is different than network flow data used in this article. Here, it refers to quantities flowing through a weighted and directed graph.

Fig. 3. Different centrality scores on the same graph. Vertex size is proportional to a normalized centrality score, and color is proportional to score percentile, with yellow for higher percentiles and red for lower percentiles.

Figure 3 illustrates these different centrality scores on a single G_t from a dynamic graph of network flow. The specific source and snapshot of the data is given in Section 5, where we describe our experiments in detail but the structure is as outlined in Section 2.1. In Figure 3 vertices are sized proportionally² to their centrality score and colored according to that score's percentile in the graph. Comparing them, we see how seemingly subtle differences in centrality definitions can yield drastically different scores. For betweenness centrality, we observe a heavily skewed score distribution. This is partly due to the prevalence of pendant (i.e., degree 1 vertices), which cannot be on a shortest path between two other vertices, and hence receive a betweenness score of 0. In contrast, the closeness centrality scores are much more uniform across the graph. The high-degree or “hub” vertices are not well-distinguished from the pendant vertices because their distances, on average, to the rest of the graph are comparable. Turning our attention to the two walk-based measures, PageRank and Katz centrality,³ we see that both better distinguish hub from pendant vertices. PageRank assigns relatively larger scores to hubs than Katz centrality. Taking a broader view, this example illustrates the importance of understanding the foundations of any proposed centrality metric to interpret its behavior.

In Section 3, we will propose and justify our own centrality score based in spectral graph theory. While the aforementioned centrality measures are implicitly related to eigenvalues, via the relationship of eigenvalues to random walks and walk-counting noted above, the measures we will consider are *explicitly* derived from eigenvalues and eigenvectors. In this vein, our work will build off prior research by Qi [43], who defined a spectral importance measure they call *Laplacian centrality*. There, Qi defines vertex importance based on its contribution to the entire spectrum, as measured by the *Laplacian energy*:

$$\mathcal{E}(G) = \sum_{i=1}^n \lambda_i^2.$$

More precisely, the Laplacian centrality of a vertex v is the percent change in Laplacian energy from deleting v , i.e.,

$$\frac{\mathcal{E}(G) - \mathcal{E}(G \setminus v)}{\mathcal{E}(G)},$$

where $G \setminus v$ denotes the graph formed by deleting v and all of its edges from G . We note Laplacian centrality is always non-negative [43]. We will return to and visualize Laplacian centrality alongside our proposed notion of centrality at the end of Section 3.

²To make the vertex sizes comparable across centrality measures, the scores within each centrality measure are divided by the sum of all the scores for that centrality measure.

³Here, PageRank is computed with damping parameter $\alpha = 0.85$, and Katz centrality is computed with attenuation parameter $\alpha = \frac{1}{2\lambda_{\max}}$ where $\lambda_{\max}$ is the adjacency matrix spectral radius of G .

3 DIRECTIONAL CENTRALITY MEASURES

As we have noted in Section 2.2, functions of the various spectra of the graph can encode a variety combinatorial properties of the underlying graph. This motivates using spectral information to measure the importance of particular vertices within that structure. Perhaps the most natural way to evaluate the importance of a particular vertex to a spectral function is the aforementioned approach taken by Qi et al: remove the vertex from the graph, recalculate the spectrum and the associated function, and record the change.

However, in some cases deleting a vertex from a graph may induce dramatic changes to the graph structure that disproportionately affect the spectrum. For example, recall from Section 2.2 that the Cheeger constant, a measure of neighborhood expansion, is closely tied to the second eigenvalue of the normalized Laplacian. Consequently, the deletion of *any* vertex that disconnects a graph—including, for example, vertices whose removal only isolates a single vertex from the rest of the graph—converts the Cheeger constant and second eigenvalue to 0. Other graph parameters we surveyed in Section 2.2 may also be changed significantly by the deletion of a vertex that, in turn, induces jump-shifts in the spectrum. For this reason, assessing importance based on outright vertex deletion outcomes may be an insufficiently nuanced approach. Instead, we propose measuring importance using a function of the spectrum based on an *infinitesimal* change in the graph structure. More formally, we consider the derivative of an eigenvalue in the direction of a vertex. In the remainder of this section, we develop this notion and use it to define a more targeted notion of Laplacian centrality that we call *Directional Laplacian Centrality*. Following this technical derivation, we will build overall intuition for the behavior of Directed Laplacian Centrality by applying it to a variety of well-studied networks in Section 4. We will then perform experiments using real and synthetic network flow data in Section 5 to show how the Directional Laplacian Centrality changes in the presence of two types of planted anomalies and indicate how it could be used by an analyst to facilitate situational awareness.

3.1 Eigenvalue Directional Derivative

Before we develop the notion of derivative of an eigenvalue in the direction of a vertex, we first recall some basic facts about the dependence of an eigenpair on the entries of a matrix, (see for instance Reference [35]).

THEOREM 1. *Let X_0 be an real-symmetric matrix in $\mathbb{R}^{n \times n}$ and let (λ_0, v_0) be an associated eigenpair such that $\|v_0\| = 1$. If λ_0 is a simple eigenvalue, then there is a neighborhood $N(X_0)$ and functions $\lambda: N(X_0) \rightarrow \mathbb{R}$ and $v: N(X_0) \rightarrow \mathbb{R}^n$ such that for all $X \in N(X_0)$*

- (1) $\lambda(X_0) = \lambda_0$,
- (2) $v(X_0) = v_0$,
- (3) $\|v(X)\| = 1$, and
- (4) $Xv(X) = \lambda(X)v(X)$.

Furthermore, λ and v are infinitely differentiable on $N(X_0)$.

In particular, this implies that if (λ, v) is a simple eigenpair associated with a real-symmetric matrix $A \in \mathbb{R}^{n \times n}$, then for any other matrix $B \in \mathbb{R}^{n \times n}$, the matrix function $M(t) = A + tB$ defines, via the implicit function theorem, a pair of functions $(\lambda(t), v(t))$ that are infinitely differentiable in a neighborhood, N , of 0 and such that $\|v(t)\| = 1$ and $M(t)v(t) = \lambda(t)v(t)$ for all $t \in N$. If we think of A as the adjacency matrix of a graph and B as the collection of edges incident to a vertex x , then this allows us to naturally define a notion of a directional derivative in the direction of the vertex x for any simple eigenpair. Specifically, recalling that given a parameterized matrix $M(t)$ and the associated eigenpair functions $(\lambda(t), v(t))$, we have that

$$\frac{d\lambda}{dt}(t_0) = v(t_0)^T \frac{dM}{dt}(t_0) v(t_0),$$

and thus, the directional derivative in the direction of x of (λ, v) for the adjacency matrix is given by $v_x \sum_{y \sim x} v_y$. Using this framework, we will define the directional derivative for eigenpairs of both the combinatorial and normalized Laplacian.

LEMMA 2. Let $G = (V, E)$ be a simple graph and let x be an arbitrary vertex. If (λ, v) is an eigenpair of the combinatorial Laplacian L of G , then the derivative of λ in the direction x is given by

$$\sum_{y \sim x} (v_x - v_y)^2.$$

If instead, (λ, v) is an eigenpair of the normalized Laplacian of G , then the derivative of λ in the direction x is given by

$$(1 - \lambda) \sum_{y \sim x} \left(\frac{v_x}{\sqrt{d_x}} - \frac{v_y}{\sqrt{d_y}} \right)^2 - \lambda \sum_{y \sim x} \frac{2v_x v_y}{\sqrt{d_x d_y}}.$$

PROOF. Similarly as above, we define the adjacency matrix in the direction of x to be

$$A_{ij}(t) = \begin{cases} 1 + t & \{i, j\} \in E, x \in \{i, j\} \\ 1 & \{i, j\} \in E, x \notin \{i, j\} \\ 0 & \text{otherwise} \end{cases}.$$

The matrix $D(t)$ is then the diagonal matrix of t -dependent degrees, i.e., the matrix that has $A(t)\mathbb{1}$ on the diagonal. In particular, $D_{xx}(t) = (1 + t)d_x$, $D_{yy}(t) = d_y + t$ for $y \sim x$, and $D_{zz}(t) = d_z$ for $z \not\sim x$. Thus, we have that

$$L'_{ij}(t) = (D'(t) - A'(t))_{ij} = \begin{cases} d_x & i = j = x \\ 1 & i = j, \{i, x\} \in E \\ -1 & \{i, j\} \in E, x \in \{i, j\} \\ 0 & \text{otherwise} \end{cases},$$

and we see that the derivative of λ in the direction x is given by

$$v(0)^T L'(0) v(0) = d_x v_x^2 + \sum_{y \sim x} v_y^2 - 2v_x v_y = \sum_{y \sim x} (v_x - v_y)^2,$$

as desired.

For the normalized Laplacian the situation is slightly more complicated by the need to use the non-commutative product rule, however, recalling that $\mathcal{L}(t) = I - D(t)^{-1/2} A(t) D(t)^{-1/2}$, we have that

$$\frac{d}{dt} \mathcal{L}(t) = \frac{1}{2} D(t)^{-3/2} D'(t) A(t) D(t)^{-1/2} - D(t)^{-1/2} A'(t) D(t)^{-1/2} + \frac{1}{2} D(t)^{-1/2} A(t) D(t)^{-3/2} D'(t).$$

Hence, using that $D(t)^{-3/2} D'(t) = D'(t) D(t)^{-1} D(t)^{-1/2}$ and that $D(0)^{-1/2} A(0) D(0)^{-1/2} v = (1 - \lambda)v$,

$$\begin{aligned} v^T \mathcal{L}'(0) v &= \frac{1}{2} v^T D'(0) D^{-1}(0) (1 - \lambda) v - v^T D(0)^{-1/2} A'(0) D(0)^{-1/2} + (1 - \lambda) v^T D(0)^{-1} D'(0) v \\ &= (1 - \lambda) v^T D'(0) D^{-1}(0) v - v^T D(0)^{-1/2} A'(0) D(0)^{-1/2} \end{aligned}$$

$$\begin{aligned}
&= (1 - \lambda) \left(v_x^2 + \sum_{y \sim x} \frac{v_y^2}{d_y} \right) - \sum_{y \sim x} 2 \frac{v_x v_y}{\sqrt{d_x d_y}} \\
&= (1 - \lambda) \left(v_x^2 - 2 \sum_{y \sim x} \frac{v_x v_y}{\sqrt{d_x d_y}} + \sum_{y \sim x} \frac{v_y^2}{d_y} \right) - \lambda \sum_{y \sim x} 2 \frac{v_x v_y}{\sqrt{d_x d_y}} \\
&= (1 - \lambda) \sum_{y \sim x} \left(\frac{v_x}{\sqrt{d_x}} - \frac{v_y}{\sqrt{d_y}} \right)^2 - \lambda \sum_{y \sim x} \frac{2v_x v_y}{\sqrt{d_x d_y}},
\end{aligned}$$

as desired. $\square$

It is easy to see that for both the combinatorial and normalized Laplacian the derivative with respect to the eigenspace corresponding to eigenvalue 0 ($\mathbb{1}$ for the combinatorial Laplacian and $\langle \sqrt{d_i} \rangle_i$ for the normalized Laplacian) in the direction of any vertex is zero, as expected. Unfortunately, it is clear that if λ is a non-simple eigenvalue (and hence corresponds to a eigenspace of dimension at least 2), then the definition of the derivative depends on the choice of eigenvector associated with λ . However, the following result shows that if we instead define the derivative in terms of the entire eigenspace associated with λ , then the derivative is independent of the particular decomposition of the eigenspace.

LEMMA 3. *Let S be a k -dimensional subspace of $\mathbb{R}^n$ and let $x, y \in \mathbb{R}^n$. There is some constant C such that $\sum_i x^T v_i y^T v_i = C$ for any choice of orthonormal basis of S .*

PROOF. Fix an arbitrary orthonormal basis of S , $\{v_1, \dots, v_k\}$ and let P be an orthonormal matrix in $\mathbb{R}^{k \times k}$. Define $w_i = VPe_i$ $V = [v_1 \dots v_k]$. Now consider

$$\begin{aligned}
\sum_i x^T w_i y^T w_i &= \sum_i x^T (VPe_i) y^T (VPe_i) \\
&= \sum_i x^T VPe_i e_i^T P^T V^T y \\
&= x^T VP \left(\sum_i e_i e_i^T \right) P^T V^T y \\
&= x^T VPP^T V^T y \\
&= x^T VV^T y.
\end{aligned}$$

As this value is independent of the choice of P , and in particular P can be chosen to be the identity, we have the existence of the desired constant. $\square$

Thus, combining the observations of Lemma 2 and 3, we have the following well-founded definition:

Definition (Eigenspace Directional Derivative). Let $G = (V, E)$ be a graph and let λ be an eigenvalue of the combinatorial Laplacian with $\{v^{(1)}, \dots, v^{(k)}\}$ an orthonormal basis of the associated eigenspace. The derivative of λ in the direction $x \in V$ is given by

$$\frac{1}{k} \sum_{i=1}^k \sum_{y \sim x} \left(v_x^{(i)} - v_y^{(i)} \right)^2.$$

If instead λ and $\{v^{(1)}, \dots, v^{(k)}\}$ parameterize an eigenspace of the normalized Laplacian, then the derivative of λ in the direction of $x \in V$ is

$$\frac{1}{k} \sum_{i=1}^k \left((1 - \lambda) \sum_{y \sim x} \left(\frac{v_x^{(i)}}{\sqrt{d_x}} - \frac{v_y^{(i)}}{\sqrt{d_y}} \right)^2 - \lambda \sum_{y \sim x} \frac{2v_x^{(i)}v_y^{(i)}}{\sqrt{d_x d_y}} \right).$$

We note that it is a straightforward exercise to generalize this definition to the case of weighted graphs or to the derivative in the direction of any subset of edges of the graph.

3.2 Directional Laplacian Centrality

We also note that, while the formalism and precise definition may be new, the essential idea of the eigenspace directional derivative is solidly rooted in prior work. For example, the aforementioned work of Qi et al. [43], where they define Laplacian Centrality as $\frac{\mathcal{E}(G) - \mathcal{E}(G \setminus v)}{\mathcal{E}(G)}$, where $\mathcal{E}(G)$ is the sum of squares of eigenvalues of the combinatorial Laplacian, can be understood as an attempt to capture the idea of a directional derivative. Specifically, letting A be the adjacency matrix of G and letting B_x be the adjacency matrix of the edges incident to x , define $A(t) = A + tB_x$. Similarly as above, $D(t) = A(t)\mathbb{1}$ and $L(t) = D(t) - A(t)$. With this notation it is easy to see that the Laplacian Centrality of x can be viewed as a re-scaling of the approximation

$$\left. \frac{d}{dt} \text{trace}(L^2(t)) \right|_{t=0} \approx \frac{\text{trace}(L^2(0)) - \text{trace}(L^2(-1))}{0 - (-1)}.$$

Now, given a complete orthonormal decomposition $\{(\lambda_i, v^{(i)})\}$ for the combinatorial Laplacian, we can evaluate this derivative exactly as

$$\begin{aligned} \left. \frac{d}{dt} \text{trace}(L^2(t)) \right|_{t=0} &= \left. \frac{d}{dt} \left(\sum_i \lambda_i(t)^2 \right) \right|_{t=0} \\ &= \sum_i 2\lambda_i(0)\lambda_i'(0) \\ &= \sum_i 2\lambda_i \sum_{y \sim x} (v_x^{(i)} - v_y^{(i)})^2. \end{aligned}$$

It is interesting to note that, by standard rearrangements of the quadratic form, $\lambda_i = \sum_{\{x,y\} \in E} (v_x^{(i)} - v_y^{(i)})^2$ and thus $\sum_{y \sim x} (v_x^{(i)} - v_y^{(i)})^2$ can be thought of as the portion of λ_i “incident” with x .

These ideas can be further expanded to understand the relative importance of any combinatorial substructure to functions of the spectrum. Namely, if f is a function of $\lambda = (\lambda_1, \dots, \lambda_k)$, the eigenspaces of a matrix associated with a graph $G = (V, E)$, and E' is the set of edges associated with a combinatorial substructure of G , then relative importance of E' to f can be described by

$$\frac{d}{dE'} f(\lambda) = \nabla f(\lambda)^T \frac{d\lambda}{dE'}.$$

Returning to the Laplacian Centrality definition of Qi et al. [43], it is easy to see that it can be phrased within this framework by defining f as the sum over all eigenspaces $\{(\lambda_i, V_i)\}$ of $\dim(V_i)\lambda_i^2$.

Having established this context, we introduce our proposed importance measure Directional Laplacian Centrality.

Definition (S-Directional (Normalized) Laplacian Centrality). Let $G = (V, E)$ be a graph on n vertices with combinatorial and normalized Laplacians L and $\mathcal{L}$, respectively. Let $\{(\lambda_i, v^{(i)})\}$ be an orthonormal decomposition

Fig. 4. (a): Directional Laplacian (top row) and Normalized Laplacian Centrality (bottom row) on the same graph, for the 5 largest (right side) and 5 smallest (left side) nontrivial eigenvalues. (b): Laplacian Centrality (LC) and Normalized Laplacian Centrality (nLC). Vertex size is proportional to score, and color is proportional to score percentile, with yellow for higher percentiles and red for lower percentiles.

of the eigenstructure of L (respectively, $\mathcal{L}$) such that $\lambda_i \leq \lambda_{i+1}$. For any set $S \subseteq \{1, \dots, n\}$, the S -Directional Laplacian Centrality (respectively, S -Directional Normalized Laplacian Centrality) of a vertex $x \in V$ is

$$S\text{-DLC}(x) = \sum_{s \in S} \sum_{y \sim x} \left(v_x^{(s)} - v_y^{(s)} \right)^2,$$

$$S\text{-nDLC}(x) = \sum_{s \in S} \left((1 - \lambda_s) \sum_{y \sim x} \left(\frac{v_x^{(s)}}{\sqrt{d_x}} - \frac{v_y^{(s)}}{\sqrt{d_y}} \right)^2 - \lambda_s \sum_{y \sim x} \frac{2v_x^{(s)}v_y^{(s)}}{\sqrt{d_x d_y}} \right),$$

respectively.

As oftentimes we are interested in the extremal eigenvalues,⁴ we use the notation k -DLC and $\bar{k}$ -DLC to denote the sets $\{t + 1, \dots, t + k\}$ and $\{n - k + 1, \dots, n\}$ where t is the dimension of the nontrivial null space of L . In this way, eigenvectors in the null space are excluded, because the directional derivative for such eigenvectors is always trivially zero. As an example, we compute the Directional Laplacian Centrality and Normalized Laplacian Centrality for the first 5 and last 5 nontrivial eigenvalues of the same graph considered in Figure 3. We present a visualization of the centrality values in Figure 4, alongside Qi's Laplacian Centrality and Normalized Laplacian Centrality.⁵

⁴This preference is motivated by theoretical observations such as the Wigner Semi-circle Law [48–51] and the related extensions [45], as well as the spectral analysis of the random d -regular graph [36, 46] and the Chung-Lu random graph [14–16], which show that behavior of the “bulk” of the spectrum of many random matrices (including those associated with random graphs) has a limiting behavior. This indicates that if we can think of the process that generates a graph in a “low-rank” manner, then the “bulk” of the spectrum is attributable to stochasticity rather than intrinsic properties of the generative process.

⁵Although only defined for the combinatorial Laplacian by Qi in Reference [43], the Normalized Laplacian Centrality can be defined in the same way, as the percent change in the normalized Laplacian energy.

Comparing the visualizations within this figure, the differences underscore the impact the choice of eigenvalues and matrix may have on the scores. These differences in centrality score align with our understanding that the maximal and minimal eigenvalues of a graph capture very different properties. Comparing Figure 4(a) to the previously depicted PageRank, Katz, Betweenness and Closeness centrality visualizations in Figure 3, there are clear qualitative differences. Like those measures, DLC may ascribe importance to “hub” vertices with high degree, particularly as seen in the $\bar{5}$ -DLC. But in certain cases, such as the 5-nDLC, low-degree vertices on long radial paths leading to hub vertices are also highlighted. This may be helpful when looking for lateral movement anomalies in a network, as long paths can indicate such behavior. In contrast, the distance and walk-based measures give minimal scores to these vertices, as do Qi’s Laplacian Centrality and normalized Laplacian Centrality scores visualized in Figure 4(b). The ability to measure structural importance beyond hubs is particularly valuable and intriguing, as this suggests DLC may reflect structural nuances otherwise lost in popular measures. In the next section, we conduct experiments aimed at illustrating the behavior of DLC and nDLC on a variety of standard network science datasets.

4 APPLICATIONS OF DIRECTIONAL LAPLACIAN CENTRALITY

Before addressing the applicability of Directional Laplacian Centrality to cyber-situational awareness in Section 5, we develop an intuition for the behavior of the DLC and nDLC by applying it to several well-understood graph datasets. Below, we consider the Karate Club, Les Misérables, and Network Science collaboration datasets. Then, in Section 4.2, we apply DLC and nDLC to the dynamic graph from the Enron email dataset.

4.1 Social Interaction Data

We consider several prominent datasets frequently utilized to test graph centrality methods. Each of these datasets features one or more vertices that, in their respective context, can be naturally interpreted as the most important or central to that graph. Figure 5 presents the $\bar{5}$ -DLC results for these datasets, which we describe below:

- **Zachary’s Karate Club.** A graph of social interactions amongst members of a karate club [53]. Due to a conflict between an instructor Mr. Hi, and an administrator, John, the club fractured into two groups. The two highest-scoring $\bar{5}$ -DLC values are achieved by the leaders of these groups, Mr. Hi and John.
- **Les Misérables.** A co-occurrence graph from Reference [33] derived from Victor Hugo’s novel. Vertices represent characters that are linked if they co-occur in a scene. The $\bar{5}$ -DLC assigns the highest centrality value to the protagonist, Jean Valjean, while other major characters such as Javert and Gavroche are similarly ranked highly.
- **Network Science.** A collaboration graph from 2006 [40] containing 379 researchers in network science. Two authors are linked if they co-authored a paper together. As noted in Reference [40], the top five authors with the highest $\bar{5}$ -DLC scores are group leaders or senior researchers of groups working in network science.

These observations affirm that DLC may serve as a reasonable centrality measure identifying known important entities in graphs derived from different contexts. While encouraging, such results are not surprising; indeed other centrality measures, such as closeness, betweenness, PageRank, and Katz, would also likely rank some of the same vertices highly.

To more comprehensively compare our method to existing centrality methods, next, we compare not only the top scoring vertices, but how the entire rankings compare to those given by these other well-known methods. Furthermore, we apply our measures under a variety of different parameter settings, considering both DLC and nDLC, and using the top k smallest or largest eigenvalues, for all k . For each such case, we quantify the similarity in ranking using a well-known nonparametric statistical measure, Spearman’s rank correlation coefficient.

Fig. 5. The Directional Laplacian Centrality scores for the 5 largest eigenvalues of the Karate Club, Les Misérables, and Network Science datasets. In the visualisation (top), vertex size is proportional to score, and color is proportional to score percentile, with yellow for higher percentiles and red for lower percentiles. Beneath each visualization, the plots present the top 15 scores.

Spearman's ρ ranges from -1 (if the ordinal rankings are reverses of each other) to 1 (if the ordinal rankings are identical). Figure 6 presents the results for the Network Science collaboration graph, the largest of the three datasets.

Examining the correlation coefficients for closeness and betweenness centrality, neither DLC nor nDLC exhibit strong rank correlations with these measures, regardless of whether the top or bottom k eigenvalues are utilized. However, PageRank and Katz centrality show very strong correlations with $\bar{k}$ -DLC, which tend to increase as more of the top eigenvalues are considered. In contrast, the $\bar{k}$ -nDLC rankings consistently show moderate anti-correlations across all four centrality measures. Across all four centrality measures, we observe a rapid convergence between the k and $\bar{k}$ versions of DLC and nDLC as k approaches the number of vertices. While the agreement of these measures is inevitable (as the set of eigenvalues are eventually the same) it is notable this convergence happens over an abbreviated range of choices for k (especially for nDLC).⁶ We suspect this is a general phenomenon reflective of the tight correlation of small eigenvalues of L and $\mathcal{L}$ with the structure of G . However, as we will see in the next section, there is still significant information that can be recovered by considering both k and $\bar{k}$ variants. Overall, the range of differences observed here underscores the flexibility afforded by the parameter settings and choice of matrix in applying our measures, and confirms DLC and nDLC rankings may be distinct from those given by other prominent centrality measures.

⁶Although it is clear that when the entire spectrum is being considered the exact same ordering will be recovered, we note that equal values of the Spearman's ρ (such as between k -DLC and $\bar{k}$ -DLC for PageRank) does not imply that the orderings induced by those measures are the same.

Fig. 6. Correlations between rankings of authors in the Network Science dataset derived from our Directional Laplacian and Normalized Laplacian centrality under different parameter settings, compared with rankings derived from four other well-known centrality measures.

4.2 Enron Email Data

The Enron email dataset is a well-studied dataset [4, 7, 21, 32, 34, 42, 44] resulting from the public release of emails of approximately 150 Enron employees gathered by a **Federal Energy Regulation Commission (FERC)** investigation into the manipulation of energy market prices in the western region (in particular, California) by Enron. Subsequently, these emails were deemed a public record and were made available to the public in various cleaned formats.⁷ We restrict our attention to the approximately 13,000 emails in the dataset with a time stamp between September 12th, 2001, and December 2nd, 2001 (the day before Enron laid off approximately 4,000 employees). For each seven-day window within this time frame, we build the contact graph where two employees are connected by an edge if there is any email communication between them in that window. From the graph associated with the window, we calculate the percentile rank according to 5-DLC, 5-nDLC, $\bar{5}$ -DLC, and $\bar{5}$ -nDLC for each employee in the giant component. The results of these calculations (sorted by the sum of all percentile scores over the entire time frame) are displayed in Figure 7. While an in-depth investigation of these results is beyond the scope of this work, we wish to point out three high-level observations resulting from this analysis.

First, we note that for 5-DLC, 5-nDLC, and $\bar{5}$ -DLC, the most important individual over the period of interest is the **Vice President of Trading in charge of the market segment corresponding to gas in the western region (VPT-GW)**. As FERC's investigation was focused on the manipulation of energy prices in California, it is unsurprising that the VPT-GW would be a key person of interest and this is reflected in the structure of the

⁷The analysis in this work is based on the version available at <https://s3.amazonaws.com/enron-shetty-adibi/enron-mysqldump.tar.gz>.

Fig. 7. Percentiles of 5-DLC, 5-nDLC, 5-DLC, and 5-nDLC of the Enron email day from September 12th to December 2nd (the day before significant layoffs). Each email graph is formed from the emails for the previous seven days. Each plot is sorted in order of increasing sum the percentiles over all graphs. Important/interesting employees in the dataset are identified by job title, Chief Executive Officer (CEO), Chief Operating Officer (COO), President and Executive Officer (P/EO), Vice President of Trading – Gas West (VPT-GW). The employees labelled with “**” are associated with legal oversight of Enron’s activities (i.e., those with job titles such as General Counsel, Employment Lawyer, Legal Specialist, etc.). The employee identified as Trader is one of several traders in the dataset.

Fig. 8. Flow record density of trailing one-minute window for the LANL dataset. Successful red-team authentications are denoted by vertical gray lines.

email communications gathered by FERC. It is worth noting that we are aware of no other method applied to the Enron dataset that identifies VPT-GW as the most important individual.

The second observation is that the employees associated with legal issues (denoted by a “*” in Figure 7 and including roles such as general council, employment lawyer, legal specialist, etc.) have visually higher overall importance in $\bar{5}$ -DLC and $\bar{5}$ -nDLC than in 5-DLC and 5-nDLC with an average ranking of the summed percentile importance increasing from 42 to 32 for DLC and from 34 to 25 for nDLC. While the precise reason for this shift is not clear, it seems likely that the nature of legal advising within Enron results in a structurally different communication pattern than that associated with trading, to which the bulk of the other employees can be affiliated.

Finally, we note the interesting patterns associated with a single trader in 5-DLC (denoted “Trader” in Figure 7). While there are several other employees who have a similar pattern of only occasionally communicating with other employees in the dataset, this trader is unique in the level of importance during those time periods. In particular, this trader’s average 5-DLC percentile is over 86.1% for the week-long periods when they are present in the giant component. While the reason for this behavior is unclear, this trader appears in evidence brought by **Snohomish County Public Utility District (SPUD)** in their lawsuit against Enron. According to contemporaneous reporting on the lawsuit, in the transcripts provided by SPUD this trader arranged for the shut-down of power plants to manipulate the California energy market through widespread blackouts/brownouts.

5 CYBER-SITUATIONAL AWARENESS EXPERIMENTAL SETUP AND RESULTS

To evaluate the effectiveness of our proposed DLC and nDLC, we analyze their sensitivity to planted anomalies in a graph of network flow. Test datasets with real network data and ground truth labeled anomalies can be hard to come by. There are many causes for this difficulty: For instance, network owners may not want to share technical data regarding their network for fear that this information will allow adversaries to identify critical resources in a network. Additionally, even if network owners are willing to release technical information regarding their network, in many cases even labelling the ground truth is difficult. However, recently **Los Alamos National Laboratory (LANL)** has shared three large de-identified datasets⁸ from their internal network capture [29–31, 47]. One set in particular, *Comprehensive, Multi-Source Cyber-Security Events (CMCE)*, contains authentication logs, process logs, network flow, domain name server lookups, and labeled red team authentication events over the course of 58 consecutive days. Figure 8 depicts the number of flow records in trailing one-minute windows from the LANL CMCE dataset [30] as well as timing of the successful red team authentications. It is worth

⁸<https://csr.lanl.gov/data/>.

Fig. 9. Graph visualization of network flow over one minute from CMCE data (left) and the degree distribution of that graph (right).

mentioning that the red team events do not represent all of the actions of the red team, but rather contain only the successful authentications from the red team. This dataset will be the source of our baseline graph for two of three experiments. While we recognize that no two networks are alike, we note that in our experience with diverse operational datasets there are many similarities within network traffic data across networks. The resulting graphs tend to have roughly power law degree distribution with few hubs (high-degree nodes) and many leaves or short paths. The “dandelion,” a high-degree node connected to many degree 1 nodes, is a common motif found in these graphs. We believe that the LANL data is a good surrogate for a generic network, as we find these structures to be prevalent in this dataset. But, we also recognize the need for future work to consider data from additional networks to feed our baseline.

To have a baseline graph that is not likely to be influenced by red team activities, we chose an arbitrary 1-minute time window (timestamps 1065740 to 1065800) from the CMCE network flow that corresponds to a typically sized graph and is well-separated from red team events (nearly three days after the last successful red team authentications and 10 minutes before the next red team authentication). To avoid minor technical issues with the spectra of disconnected graphs, the experiments will only use the giant component of this flow graph, which contains 94% of the vertices. This largest component has 2,005 vertices, 2,450 edges, and has a maximum degree of 605. Figure 9(a) shows our chosen graph, which was also used in Figures 3 and 4 to illustrate differences in centrality scores. The degree distribution of this graph is roughly power-law with exponent approximately equal to 2.18 (see Figure 9(b)). Into this specific snapshot from the LANL data, we will inject two types of anomalies, represented as extremal subgraphs depicted in Figure 10, the star and clique. The specific method of injection varies by experiment, and so those details are provided in Sections 5.1, 5.2, and 5.3. These two subgraphs represent building blocks of adversarial behavior: network reconnaissance or data gathering (clique) and lateral movement (star). By studying the injection of these two patterns separately, we can better understand how their presence will perturb importance values.

Star anomaly: In the context of network flow, the injection of a star represents the scenario in which a particular IP address (represented by the central, or *root*, vertex r) exchanges data with some additional vertices (the leaves) while the rest of the graph remains exactly the same. Such a scenario might be malicious or benign, depending on the number of additional IPs contacted and the nature of the data exchange. Regardless, this structural change should be reflected as an increase in the importance of r , since the increase in number of IPs with which v exchanges data is the sole change in the graph.

Fig. 10. A star graph (left) and a clique (right).

Clique anomaly: The injection of a clique represents the scenario in which a particular set of IP addresses all exchange data with each other while the rest of the graph remains exactly the same. Such a scenario could be benign, representing a broadcasting operation in a distributed computing scenario, or could represent malicious behavior, such as a prelude to data exfiltration or command-and-control. In either case, the importance of the involved vertices should increase as the inter-set communication represents the sole change in the graph.

Using these anomalies, we conduct three injection experiments, described in the next three subsections, with increasing levels of complexity and stringency. The first two experiments will be performed on our chosen snapshot of LANL network flow, while the third uses a synthetic time-evolving graph. In Section 5.1, we consider injection of both patterns among low importance vertices in the baseline graph mentioned above, and in Section 5.2, we randomly inject the star anomaly. This simulates what could happen if the only temporal change in a graph is the addition of one of these subgraphs. In Section 5.3, we inject the star anomaly within normal temporal evolution. It is worth mentioning that conceptually simpler measures (such as degree change or local density) may also be able to detect the root of the star-anomaly or the vertices of clique-anomaly when compared to a sufficiently stable background. However, in the dynamic environment, such as that discussed in Section 5.3, these measures will have significant difficulty in distinguishing the injection of anomalous behavior amid the natural variation of the underlying graph.

5.1 Anomaly Injection among Low-Importance Vertices

In this first experiment, we deterministically choose vertices with low importance scores to be involved in the anomaly. This simple experiment tests whether, all else equal, the change to lower-importance vertices is registered by the centrality metrics. The procedure for injecting both anomalies into G starts by ordering the vertices of G according to importance in G , lowest to highest. Then, to inject a star anomaly of size s into G :

- (i) Select a *root* vertex r to be one of the vertices of minimum importance.
- (ii) Select the next smallest $s - 1$ vertices in the list of importance to be *leaves*.
- (iii) Add, if necessary, an edge between the root vertex and any leaf.

Similarly, the procedure for injecting a clique anomaly of size s into G is:

- (i) Select a set S of s smallest importance vertices.
- (ii) For every pair of vertices, $u, v \in S$, add, if necessary, an edge between u and v .

Figure 11 displays the percentile of the importance values of the vertices involved in the injection as more vertices are added to the anomaly (indicated on the x axis). For example, in Figure 11(a) at $x = 100$ the corresponding y value for the solid line is the percentile of the 5-DLC for the root vertex after 100 leaf vertices have been attached to it. The y value for the dashed line is the percentile of the average 5-DLC over all of the 100 leaf vertices. For both the leaf vertices and the vertices involved in the clique, the centrality scores are averaged before evaluating the centrality percentile. We use the percentile rather than the raw centrality score

Fig. 11. Importance percentile for vertices in injected anomalies as a function of anomaly size.

to put centralities of all graphs on the same scale. Each x value corresponds to a different graph (the baseline graph with a differently sized anomaly injected) and thus the importance values may not be comparable across all x values. For this reason, we will continue to use the percentile in the remaining two experiments.

The most striking result of this experiment is that when the injection size is larger than approximately 7.5% of the graph for $\bar{5}$ -DLC (Figures 11(a) and 11(c)) the *average* importance of the participating vertices is at or above the 99th percentile, representing an increase in importance percentile of more than three orders of magnitude. A similar increase occurs for the 5-nDLC under clique injection when the anomaly size is approximately 12.5% of the graph. Further, the results of these experiments also clearly indicate that both importance measures can be sensitive to even fairly small injections. For instance, when the injection size is approximately 1% of the size of the graph, both importance measures register an order of magnitude increase for the star injection and an approximately two orders of magnitude increase for the clique injection.

In the negative direction, it is clear by examining Figure 11(b) that beyond an initial increase in relative importance, the nDLC measure is relatively insensitive to the size of the injected star graph. However, this is consistent with the known behavior of the normalized Laplacian. Specifically, by applying the theory of quasi-randomness, it can be shown that the spectrum of the normalized Laplacian is sensitive to presence of small cycles in the graph [9–11, 18]. As a star graph contains no cycles, any small cycles that emerge necessarily have a significant contribution from the underlying graph. Further, as the example graph has relatively few edges, it is natural to expect that, even for quite large injections of a star, there will be relatively few short cycles formed in the graph.

5.2 Anomaly Injection among Random Vertices

In this next experiment, we randomly choose vertices to be involved in the anomaly and consider the results over many trials. In this sense, this experiment reflects average case behavior in a static graph, and accounts for the effects of vertex choice in the prior experiment. *A priori*, there is no reason to believe that the importance measures are independent of the precise set of vertices involved in anomalous behavior. In particular, it is possible that the choice of vertex order in the experiments depicted in Figure 11 is a nontrivial factor in the results. Thus, to test the sensitivity of these results, an additional series of experiments injecting a star anomaly is performed. Namely, as before, the LANL flow graph depicted in Figure 9(a) is utilized as a test case in which structural

changes are planted, and the subsequent changes in importance scores are measured over many trials. The methodology for a single trial is as follows:

- (i) Randomly select a vertex v .
- (ii) Randomly choose a vertex subset S with $k\%$ of the graph's vertices.
- (iii) For any $s \in S$ that is not already connected to v , connect v to s .
- (iv) Compute importance score and score percentile of v before and after these additional edge insertions.

By measuring the change in importance score of v for different levels of k , this experiment tests the importance measure in two regards: (1) whether the importance measure can *detect* the aforementioned structural change (as indicated by an increase in the importance score); and (2) whether the importance measure is *sensitive* to the intensity of these changes (as indicated by larger increases in the importance score for larger values of k). To reduce the overall computational load and limit the effects of averaging over a large set of vertices, this experiment is restricted to studying the importance of the root vertices when a star is injected.

The results of the experiment are presented in Table 1. The experiment was run for $k = 0.1, 0.5, 1, 5$, and 10% , which corresponds to star anomalies with 2, 10, 20, 100, and 201 edges incident to the root v . The table presents the mean scores and mean percentile before and after the star is inserted, along with the mean change, over 500 trials. Turning attention first to the $\bar{5}$ -DLC scores, we observe a positive net change. Even for the smallest value of k in which only two additional edges are added, the $\bar{5}$ -DLC scores are responsive, registering an average score percentile increase of 11%. Additionally, we note this change in percentile increases monotonically for each value of k , with the largest value of k registering an increase of 53% placing the root vertex in the 99th percentile. These results suggest $\bar{5}$ -DLC scores can both detect the star anomaly at minuscule levels of intensity and reflect the relative magnitude of the star anomaly.

In contrast, the nDLC scores *decrease* in response to the star anomaly. Furthermore, these scores decrease more as the star anomaly intensity parameter k increases. This suggests nDLC scores can detect, and are sensitive to, the star anomaly, but in the opposite direction of the DLC score values. To explain why this apparent “inverted scale” occurs for nDLC, the aforementioned theory of quasi-random graphs [9–11, 18] is again relevant. By this theory, the prevalence of short cycles greatly impacts the spectrum, and the star injection is likely to increase the number of such cycles incident to the root vertex v . The net effect of this change on the nDLC score is ultimately negative, because the normalized Laplacian matrix shifts and inverts the sign of the spectrum of the matrix $D^{-1/2}AD^{-1/2}$. It is worth noting these observations are in agreement with the results presented in Figure 11(b), in which root vertices exhibited smaller nDLC values than leaf vertices. However, more research is necessary to better interpret nDLC scores and to reconcile these observations with the clique experiment of Figure 11(d) in which the nDLC scores increase. This dichotomy suggests that nDLC scores can have seemingly counterintuitive interpretations, while DLC scores afford simpler interpretations in the contexts considered here.

5.3 Anomaly Injection with Time-evolution

The prior two experiments show that on a static graph the inclusion of an anomaly is reflected in a change in DLC or nDLC. However, for situational awareness, one must compare a current snapshot with an anomaly to a *prior* snapshot without an anomaly. In this way, we wish to see that the DLC or nDLC are also perturbed significantly in the presence of an anomaly injected at a specific timestep, aiding an analyst in discovering the anomalous vertex when comparing importance scores across timesteps. To that end, in our final experiment, we assume a star anomaly occurs concurrently with the natural time-evolution of a network. This tests the robustness of the measures in the presence of natural noise. To simulate the natural time-evolution of a cyber network, we require a temporal graph model. For this purpose, we utilize an extension of a model developed by Hagberg, Lemons, and Misra that was designed to capture temporal dynamics observed in cyber data.

5.3.1 Hagberg-Lemons-Misra Model. Temporal variations in network traffic comprise a mix of fairly steady computer-generated traffic and more variable human-generated traffic. Overall, in the data we have

Table 1. Importance Scores for a Vertex before and after Star Injection According to the Proportion of Other Vertices to which the Chosen Root Vertex Connects

		0.1% (2 edges)		0.5% (10 edges)		1.0% (20 edges)		5.0% (100 edges)		10.0% (201 edges)	
		<i>Score</i>	<i>PCTL</i>	<i>Score</i>	<i>PCTL</i>	<i>Score</i>	<i>PCTL</i>	<i>Score</i>	<i>PCTL</i>	<i>Score</i>	<i>PCTL</i>
5-DLC	Before	2.30	44%	1.08	45%	1.11	45%	0.83	43%	2.04	46%
	After	2.31	55%	1.11	64%	1.19	67%	1.47	74%	205.68	99%
	Change	0.01	11%	0.03	20%	0.08	22%	0.64	31%	203.63	53%
5-nDLC	Before	-4e-5	50%	-7e-5	52%	-9e-7	50%	5e-5	49%	-9e-6	51%
	After	-8e-4	1%	-2e-3	1%	-4e-3	1%	-2e-2	1%	-3e-2	1%
	Change	-7e-4	-49%	-2e-3	-51%	-4e-3	-49%	-2e-2	-48%	-3e-2	-50%

The results are averaged over 500 trials.

observed, beyond some temporal variability (e.g., more traffic during the day than overnight) the structure of communications, as reflected by the degree distribution of graphs for small time windows, is fairly stable. The temporal graph model developed by Hagberg, Lemons, and Misra [25], which we refer to as the *HLM model*, allows the user to specify a degree sequence and builds on the Chung-Lu model [12, 13], which generates graphs consistent with a given expected degree sequence. HLM accounts for temporal variability by “masking” some vertex pairs at each timestep. Those edges in the current timestep that correspond to vertex pairs not in the masking set are carried to the next timestep, while vertex pairs that are in the masking set are included in the next timestep based on independent Bernoulli trials, regardless of whether they correspond to edges in the prior timestep. All timesteps of the HLM evolution are equal in distribution and equal to a Chung-Lu model with the initial degree sequence. Recent work by the authors of the present article, together with Jenne [1], have extended this model to include a temporally varying degree distribution, now the Temporal HLM or *THELma* model, and shown how to rapidly generate THELma instances. To test our DLC and nDLC measures in realistic time-varying data, we will use a THELma instance with parameters measured from the LANL dataset.

While the THELma model can be based on a variety of random graph models, for these experiments, we will consider the variant based on the Chung-Lu random graph model. Formally, this begins with a degree sequence $w = \langle w_v \rangle_{v \in V}$ as an n -dimensional vector, where $n = |V|$ is the total number of vertices present across all graphs. The value w_v is the expected degree of vertex v across all time instances. The THELma model additionally requires a vector that captures temporal variation in graph density, $\tau = \langle \tau_t \rangle_{t=1}^T$. This corresponds to the expected degree of vertex v at timestep t being $\tau_t \cdot w_v$ when there is no temporal correlation. A final necessary parameter, $\alpha \in [0, 1]$ controls the aforementioned masking set and represents the amount of temporal correlation. That is, if $\alpha = 1$, then there is no correlation between successive graphs; and if $\alpha = 0$, then successive graphs are identical. One could consider α to be a vector in $[0, 1]^{\binom{n}{2}}$, with one α value for each edge, but for simplicity, in this experiment, we use a single value for all edges. Given these parameters, the THELma model generates a sequence of graphs $G_1, G_2, \dots, G_T$. The graph at time t , G_{t+1} , is formed from G_t using the α parameter to control for how much of G_t is carried forwards. Specifically, a masking set M_t is generated where each pair of vertices, $\{u, v\}$, is in M_t independently with probability α . Any pair $\{u, v\} \notin M_t$ that is an edge in G_t is carried forward to be an edge in G_{t+1} , and any unmasked pair of vertices that is not an edge in G_t will not be in G_{t+1} . Then, any pair of vertices $\{u, v\} \in M_t$ is present as an edge in G_{t+1} independently with probability⁹ $\frac{\tau_{t+1} w_u w_v}{\rho}$, regardless of whether $\{u, v\}$ was an edge in G_t , where $\rho = \sum_{v \in V} w_v$. Putting this all together, we see that the probability of

⁹To avoid probabilities greater than 1, we actually use $\min\{1, \frac{\tau_{t+1} w_u w_v}{\rho}\}$ as the probability here.

Fig. 12. Parameters for THeLMa model.

an edge in G_{t+1} is given by

$$P(\{u, v\} \in G_{t+1}) = \begin{cases} 1 - \alpha + \alpha \cdot \frac{\tau_{t+1} w_u w_v}{\rho} & \{u, v\} \in G_t \\ \alpha \cdot \frac{\tau_{t+1} w_u w_v}{\rho} & \{u, v\} \notin G_t \end{cases}.$$

Unlike in the HLM model, which does not allow the degree sequence to be time varying, G_t is not quite equal in distribution to $\mathcal{G}(\tau_t \cdot w)$, the Chung-Lu model with expected degree sequence given by $\tau_t \cdot w$. Rather, the edge $\{u, v\}$ is present in G_t with probability $(1 - \alpha)^t \frac{\tau_1 w_u w_v}{\rho} + \sum_{i=1}^t \alpha(1 - \alpha)^{t-i} \frac{\tau_i w_u w_v}{\rho}$.

5.3.2 Experimental Results. To determine the parameters of the THeLMa model into which we will inject the anomaly, we take a series of 15 snapshots of one-minute traffic intervals in a five-minute window surrounding our example LANL graph (each of these snapshots is shifted 20 seconds in time). The w vector of THeLMa is determined by the average degree sequence (rounded up) of the vertices across these snapshots (with vertices that do not appear in a snapshot counting as having degree 0). The resulting weighting vector has 3,987 entries, with a maximum value of 617, and a average value of 1.74.¹⁰ To mimic the natural circadian rhythms expected in a cyber-security system, the temporal weighting sequence, τ , is given by 500 equally spaced evaluations of $(3 - \cos(x))/2$ over the range $[0, 4\pi]$. The resulting degree distribution and temporal sequence are illustrated in Figure 12.

To simulate a relatively slowly changing environment, we set the evolution parameter at $\alpha = 0.05$, which means that every edge has the potential to change state approximately 25 times during the course of the evolution. The result of the THeLMa model with these parameters is a sequence of 500 graphs with significant variation of parameters as shown in Table 2. The evolution of these parameters is shown in Figure 13.

For each timestep in the sequence of graphs, we will compare the generated graph at time t with the same graph after injecting a single star anomaly with 30 leaves. For consistency's sake, we will be injecting the anomaly to the same vertices across all timesteps. Before selecting the vertices that will participate in the injected anomaly, we observe that almost all vertices have some timestep in the THeLMa model where they are in a small component (i.e., not in the giant component). As the block structure of the combinatorial and normalized Laplacians respects

¹⁰It is worth noting that, because of the extreme sparsity and highly skewed nature of the degree sequence, a standard Chung-Lu model with this parameter has maximum expected degree significantly smaller than 617. In fact, the expected maximum degree is approximately 446 and the expected average degree is 1.62. To be reflective of this difference, we will refer to the value of a vertex in the w as its weight, instead of the more typical usage of expected degree.

Table 2. Statistics for Graphs in THelMa Sequence

	Minimum	Maximum	Mean
Giant Component Size	2,169	3,473	2,949.3
Giant Component Edges	2,863	6,282	4,669.7
Maximum Degree	452	896	660.9
Total Edges	3,123	6,327	4,795.1

Fig. 13. Summary statistics for generated THelMa model.

the component structure of the underlying graph, typical eigenspaces are associated with a single component and are identically zero on all other components. The sole exception to this is the case where multiple components have a common eigenvalue, in which case there does exist a basis for the associated eigenspace where all basis vectors are non-zero on all the components with the common eigenvalue. Even for these eigenspaces it is typical, for computational reasons, to choose a basis that respects the underlying component structure. We will respect this convention in our eigenbasis decomposition. As a consequence, when considering $\bar{5}$ -DLC and 5-nDLC, at any particular timestep there are at most five components whose vertices have non-zero importance. In practice, there is typically only a single component whose vertices have non-zero importance, the giant component. As a consequence, the injection of any anomaly may trivially change the importance of the vertices involved by adding them to the giant component. To avoid this issue, we will choose the anomaly from the set of 66 vertices that are in the giant component for every timestep. This set is naturally biased to those vertices of higher weight, with an average weight of 36.1, and so we further refine this set by removing the top 20% of vertices in terms of weight resulting in a set X of candidate vertices for the anomaly with an average weight of 7.1. The root and 30 leaves forming the anomaly, $\mathcal{A}$, are then selected randomly from the remaining 53 vertices.

Now, given the observed graph G_t at some timestep t , we wish to understand how the behavior of $\bar{5}$ -DLC and 5-nDLC differs between G_{t+1} and $G_{t+1} + \mathcal{A}$. To this end, for each vertex v in $\mathcal{A}$ and for each timestep t , we identify a set of vertices $S_v^{(t)}$ such that for all $s \in S_v$ the percentile ranking of the (normalized) DLC in G_t differs from v by at most 2.5% and s has a non-zero importance in G_{t+1} .¹¹ We refer to this set as the *cohort* associated with v . The former condition allows us to aggregate all vertices sufficiently similar to v , while the second condition accounts for the fact that the vertices of $\mathcal{A}$ were chosen in such a way that they have non-zero importance for all timesteps. By considering the percentiles of the (normalized) DLC for vertices in S_v , we can recover a “typical” evolution of the importance of a vertex like v after one timestep.

In Figure 14, we compare the behavior of the percentile DLC score for the anomaly vertices versus the vertices in the associated cohort. Specifically, in Figures 14(a) and 14(b), we show the distribution and the cumulative distribution, respectively, of the gap between the percentile DLC score with an anomaly present and the percentile

¹¹The sets S_v have an average size of 182.4 and 190.4, for $\bar{5}$ -DLC and 5-nDLC, respectively.

(a) Kernel Density Estimate for the distribution of the gap between the scores with and without the injected anomaly.

(b) Kernel Density Estimate for the cumulative density function for the gap between the scores with and without the injected anomaly.

(c) Kernel Density Estimate for the distribution of the gap in between the scores with the injected anomaly and the score in previous time step.

(d) Kernel Density Estimate for the cumulative density function for the gap in between the scores with the injected anomaly and the score in previous time step.

Fig. 14. $\bar{5}$ -DLC.

score without the anomaly. In Figures 14(c) and 14(d), we instead compare the percentile DLC score with an anomaly with the score in the prior timestep. As Figure 14(b) emphasizes, there is relatively little difference between the percentile scores for the cohort vertices with or without the anomaly present. In contrast, from Figure 14(a) the percentile scores for the vertices participating in the anomaly are typically about 2.5% (for the leaves) to 5% (for the root) higher than they would be without the anomaly. In contrast, in Figures 14(c) and 14(d), we see that for both the anomaly vertices and the associated cohorts, the little typical change in percentile between the previous timestep and the anomaly timestep. However, from Figure 14(c), we can see that the vertices in the cohort of the anomaly have a significantly larger probability of having a sharp decrease in the percentile ranking. Taken together, this indicates that there is likely a tendency to of the scores of vertices to naturally revert to the median behavior under the natural evolution; however, an injected anomaly somewhat moderates this tendency in a localized way (i.e., while the vertices participating in the anomaly resist reverting towards the median, this effect does not carry over to similarly situated vertices).

The figures provided in Figure 15 are analogous to those in Figure 14 except for being applied to the normalized DLC percentiles. We see that qualitatively, the results for the normalized DLC percentile are similar to those for

(a) Kernel Density Estimate for the distribution of the gap between the scores with and without the injected anomaly.

(b) Kernel Density Estimate for cumulative density function for the gap between the scores with and without the injected anomaly.

(c) Kernel Density Estimate for the distribution of the gap between the scores with the injected anomaly and the score in previous time step.

(d) Kernel Density Estimate for the cumulative density function of the gap between the scores with the injected anomaly and the score in previous time step.

Fig. 15. $\bar{5}$ -nDLC.

the DLC percentile. Quantitatively, Figure 15 indicates that the variability of for the normalized DLC percentiles is larger than that of the DLC, but the injection of the anomaly provides a larger boost to the percentile scores, particularly for the root of the anomaly.

6 DISCUSSION

This article introduced two new spectral-based vertex centrality measures, DLC and nDLC, and used the setting of network flow data to illustrate their sensitivity and ability to identify anomalies planted in graphs. Our new centralities are related to existing vertex centralities, most closely to the work of Qi et al. [43], but use a more nuanced approach by measuring changes to eigenvalues in the face of infinitesimal changes to a given vertex. This allows for more subtle structural features in a graph to have measurable centrality.

Our experiments in Section 5 show how the DLC and nDLC measures can be used to provide attribution of certain kinds of anomalies to cyber analysts. In the first experiment, we see that the percentiles of both DLC and nDLC, calculated using the top five eigenvalues, show significant increases when a star or clique anomaly is injected onto low importance vertices. This is true even when the injected anomaly is a rather small percentage

of the graph. In the second experiment, rather than considering only low importance vertices, we inject a star anomaly of varying size randomly into the graph. Again, we see that even for small anomalies there is a measurable change in both DLC and nDLC percentiles for the root vertex of the star. These two experiments taken together show that on average, even small star anomalies manifest measurable changes in both centrality measures. This indicates that looking for vertices that have a large increase in centrality from one timestep to the next may narrow the set of IP addresses to investigate for malicious behavior.

Those first two experiments did not take into account temporal variation in the network. Rather, we compared a static, typical network flow graph G to that same graph with an anomaly added. In the third experiment, we added the additional complexity of temporal variation. Using the THeLMa model to create a dynamic graph with properties similar to daily rhythm of network flow data, we injected a small star anomaly onto a set of vertices that are in what one could consider the *core* of the graph, i.e., vertices that are in the giant component for all timesteps. We showed that when comparing the DLC and nDLC of vertices in $G_{t+1} + \mathcal{A}$ to both G_{t+1} and G_t those vertices participating in the anomaly were more perturbed than their “cohort” vertices (those vertices having roughly the same centrality score in G_t). This difference was more pronounced for nDLC than DLC, but still visually evident in both. Given this observation, we hypothesize that DLC and nDLC could be used to detect an anomaly of this type by tracking cohorts over time and identifying when vertices change cohorts.

Situational awareness comprises understanding the current posture of the network and predicting the security states of future time points. This article provides a tool for identifying anomalies in a network to assess the current network state. While our measure does not provide a high-level overview of the network’s security posture, the ability of these measures to point to even small anomalies may aid operators in preventing those from becoming large anomalies.

ACKNOWLEDGMENTS

The authors would like to thank Helen Jenne for helpful discussions.

REFERENCES

- [1] Sinan G. Aksoy, Helen Jenne, Emilie Purvine, and Stephen J. Young. in preparation. Rapid Generation and Parameter Recovery of Correlated Temporal Graphs.
- [2] David Aldous and James Fill. 1995. Reversible Markov chains and random walks on graphs. <https://www.stat.berkeley.edu/~aldous/RWG/book.pdf>
- [3] Noga Alon, Fan R. K. Chung, and Ronald L. Graham. 1994. Routing permutations on graphs via matchings. *SIAM J. Discr. Math.* 7, 3 (1994), 513–530.
- [4] Brett W. Bader, Michael W. Berry, and Murray Browne. 2008. In *Discussion Tracking in Enron Email Using PARAFAC*. Springer London, 147–163.
- [5] Francis T. Boesch and Helmut Prodinger. 1986. Spanning tree formulas and Chebyshev polynomials. *Graphs Combinat.* 2, 1 (1986), 191–200.
- [6] John Adrian Bondy, Uppaluri Siva Ramachandra Murty, et al. 1976. *Graph Theory with Applications*. Vol. 290. Macmillan London.
- [7] Anurat Chapanond, Mukkai S. Krishnamoorthy, and Bülent Yener. 2005. Graph theoretic and spectral analysis of Enron email data. *Computat. Math. Organiz. Theor.* 11, 3 (2005), 265–281.
- [8] Fan Chung. 1997. *Spectral Graph Theory*. Vol. 92. American Mathematical Society.
- [9] Fan Chung and Ronald Graham. 2002. Sparse quasi-random graphs. *Combinatorica* 22, 2 (2002), 217–244.
- [10] Fan Chung and Ron Graham. 2008. Quasi-random graphs with given degree sequences. *Rand. Struct. Algor.* 32, 1 (2008), 1–19.
- [11] F. R. K. Chung, R. L. Graham, and R. M. Wilson. 1988. Quasi-random graphs. *Proc. Nat. Acad. Sci.* 85, 4 (1988), 969–970.
- [12] Fan Chung and Linyuan Lu. 2002. The average distances in random graphs with given expected degrees. *Proc. Nat. Acad. Sci.* 99, 25 (2002), 15879–15882.
- [13] Fan Chung and Linyuan Lu. 2004. The average distances in random graphs with given expected degrees. In *Internet Mathematics*, Vol. 1. 91–113.
- [14] Fan Chung, Linyuan Lu, and Van Vu. 2003. Eigenvalues of random power law graphs. *Ann. Comb.* 7, 1 (2003), 21–33.

- [15] Fan Chung, Linyuan Lu, and Van Vu. 2003. Spectra of random graphs with given expected degrees. *Proc. Nat. Acad. Sci.* 100, 11 (2003), 6313–6318.
- [16] Fan Chung, Linyuan Lu, and Van Vu. 2004. The spectra of random graphs with given expected degrees. *Internet Math.* 1, 3 (2004), 257–275.
- [17] Fan R. K. Chung. 1989. Diameters and eigenvalues. *J. Amer. Math. Soc.* 2, 2 (1989), 187–196.
- [18] Fan R. K. Chung, Ronald L. Graham, and Richard M. Wilson. 1989. Quasi-random graphs. *Combinatorica* 9, 4 (1989), 345–362.
- [19] Kimberly C. Claffy, H.-W. Braun, and George C. Polyzos. 1995. A parameterizable methodology for Internet traffic flow profiling. *IEEE J. Select. Areas Commun.* 13, 8 (1995), 1481–1494.
- [20] Kinkar Ch. Das and R. B. Bapat. 2005. A sharp upper bound on the largest Laplacian eigenvalue of weighted graphs. *Lin. Algeb. Applic.* 409 (2005), 153–165.
- [21] J. Diesner and K. M. Carley. 2005. Exploration of communication networks from the Enron email corpus. In *Proceedings of the Workshop on Link Analysis, Counterterrorism and Security, SIAM International Conference on Data Mining*. 3–14. Retrieved from http://www.andrew.cmu.edu/user/jdiesner/publications/diesner_carley_siam_enron_03_05.pdf.
- [22] Mica R. Endsley. 1995. Toward a theory of situation awareness in dynamic systems. *Hum. Fact.* 37, 1 (1995), 32–64.
- [23] Miroslav Fiedler. 1975. A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory. *Czech. Math. J.* 25, 4 (1975), 619–633.
- [24] Chris D. Godsil and Mike W. Newman. 2008. Eigenvalue bounds for independent sets. *J. Combin. Theor. Series B* 98, 4 (2008), 721–734.
- [25] Aric Hagberg, Nathan Lemons, and Sidhant Misra. 2016. Temporal reachability in dynamic networks. In *Dynamic Networks and Cyber-Security*. World Scientific, 181–208.
- [26] Rick Hofstede, Pavel Čeleda, Brian Trammell, Idilio Drago, Ramin Sadre, Anna Sperotto, and Aiko Pras. 2014. Flow monitoring explained: From packet capture to data analysis with NetFlow and IPFIX. *IEEE Commun. Surv. Tutor.* 16, 4 (2014), 2037–2064.
- [27] Sushil Jajodia, Peng Liu, Vipin Swarup, and Cliff Wang. 2009. *Cyber Situational Awareness*. Springer.
- [28] T. Jirsik and P. Čeleda. 2020. Cyber situation awareness via IP flow monitoring. In *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS 2020)*. 1–6. DOI: <https://doi.org/10.1109/NOMS47738.2020.9110327>
- [29] Alexander D. Kent. 2014. User-computer Authentication Associations in Time. Los Alamos National Laboratory. DOI: <https://doi.org/10.11578/1160076>
- [30] Alexander D. Kent. 2015. Comprehensive, Multi-source Cyber-security Events. Los Alamos National Laboratory. DOI: <https://doi.org/10.17021/1179829>
- [31] Alexander D. Kent. 2015. Cybersecurity data sources for dynamic network research. In *Dynamic Networks in Cybersecurity*. Imperial College Press.
- [32] Bryan Klimt and Yiming Yang. 2004. The Enron corpus: A new dataset for email classification research. In *Machine Learning: ECML, Jean-François Boulicaut, Floriana Esposito, Fosca Giannotti, and Dino Pedreschi (Eds.)*. Springer, Berlin, Heidelberg, 217–226.
- [33] Donald Ervin Knuth. 1993. *The Stanford GraphBase: A Platform for Combinatorial Computing*. ACM Press New York, NY.
- [34] Jure Leskovec, Kevin J. Lang, Anirban Dasgupta, and Michael W. Mahoney. 2010. Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Math.* 6, 1 (2010), 20–123.
- [35] Jan R. Magnus. 1985. On differentiating eigenvalues and eigenvectors. *Economet. Theor.* 1 (1985), 179–191.
- [36] Brendan D. McKay. 1981. The expected eigenvalue distribution of a large regular graph. *Linear Algeb. Appl.* 40 (1981), 203–216. DOI: [https://doi.org/10.1016/0024-3795\(81\)90150-6](https://doi.org/10.1016/0024-3795(81)90150-6)
- [37] Bojan Mohar. 1991. Eigenvalues, diameter, and mean distance in graphs. *Graphs Combinat.* 7, 1 (1991), 53–64.
- [38] Nour Moustafa, Jiankun Hu, and Jill Slay. 2019. A holistic review of network anomaly detection systems: A comprehensive survey. *J. Netw. Comput. Applic.* 128 (2019), 33–55.
- [39] Ido Nevat, Dinil Mon Divakaran, Sai Ganesh Nagarajan, Pengfei Zhang, Le Su, Li Ling Ko, and Vrizlynn L. L. Thing. 2017. Anomaly detection and attribution in networks with temporally correlated traffic. *IEEE/ACM Trans. Netw.* 26, 1 (2017), 131–144.
- [40] Mark E. J. Newman. 2006. Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* 74, 3 (2006), 036104.
- [41] Nikolaos Pitropakis, Emmanouil Panaousis, Alkiviadis Giannakoulas, George Kalpakis, Rodrigo Diaz Rodriguez, and Panayiotis Sari-giannidis. 2018. An enhanced cyber attack attribution framework. In *Proceedings of the International Conference on Trust and Privacy in Digital Business*. Springer, 213–228.
- [42] Carey E. Priebe, John M. Conroy, David J. Marchette, and Youngser Park. 2005. Scan statistics on enron graphs. *Computat. Math. Organiz. Theor.* 11, 3 (2005), 229–247.
- [43] Xingqin Qi, Eddie Fuller, Qin Wu, Yezhou Wu, and Cun-Quan Zhang. 2012. Laplacian centrality: A new centrality measure for weighted networks. *Inf. Sci.* 194 (2012), 240–253.
- [44] Jitesh Shetty and Jafar Adibi. 2005. Discovering important nodes through graph entropy the case of enron email database. In *Proceedings of the 3rd International Workshop on Link Discovery*. 74–81.
- [45] Terence Tao and Van Vu. 2008. Random matrices: The circular law. *Commun. Contemp. Math.* 10, 02 (2008), 261–307.

- [46] Linh V. Tran, Van H. Vu, and Ke Wang. 2013. Sparse random graphs: Eigenvalues and eigenvectors. *Rand. Struct. Algor.* 42, 1 (2013), 110–134. DOI: <https://doi.org/10.1002/rsa.20406>
- [47] Melissa J. M. Turcotte, Alexander D. Kent, and Curtis Hash. 2018. In *Unified Host and Network Data Set*. World Scientific, 1–22. DOI: https://doi.org/10.1142/9781786345646_001
- [48] Eugene P. Wigner. 1955. Characteristic vectors of bordered matrices with infinite dimensions. *Ann Math.* 62, 3 (1955), 548–564. Retrieved from <http://www.jstor.org/stable/1970079>.
- [49] Eugene P. Wigner. 1958. On the distribution of the roots of certain symmetric matrices. *Ann. Math.* 67, 2 (1958), 325–327. Retrieved from <http://www.jstor.org/stable/1970008>.
- [50] Eugene P. Wigner. 1993. Characteristic vectors of bordered matrices with infinite dimensions I. In *The Collected Works of Eugene Paul Wigner*. Springer, 524–540.
- [51] Eugene P. Wigner. 1993. Characteristic vectors of bordered matrices with infinite dimensions II. In *The Collected Works of Eugene Paul Wigner*. Springer, 541–545.
- [52] Herbert S. Wilf. 1967. The eigenvalues of a graph and its chromatic number. *J. London Math. Soc.* 42, 1967 (1967), 330.
- [53] Wayne W. Zachary. 1977. An information flow model for conflict and fission in small groups. *J. Anthropol. Res.* 33, 4 (1977), 452–473.

Received August 2020; revised November 2020; accepted February 2021