

# Seurat: A Pointillist Approach to Anomaly Detection

Yinglian Xie<sup>1</sup>, Hyang-Ah Kim<sup>1</sup>  
David R. O’Hallaron<sup>1,2</sup>, Michael K. Reiter<sup>1,2</sup>, and Hui Zhang<sup>1,2</sup>

<sup>1</sup> Department of Computer Science

<sup>2</sup> Department of Electrical and Computer Engineering  
Carnegie Mellon University

Email: {ylxie, hakim, droh, reiter, hzhang}@cs.cmu.edu

**Abstract.** This paper proposes a new approach to detecting aggregated anomalous events by correlating host file system changes across space and time. Our approach is based on a key observation that many host state transitions of interest have both temporal and spatial locality. The intuition is that abnormal state changes, which may be hard to detect in isolation, become apparent when they are correlated with similar changes on other hosts. In particular, the goal is to detect similar, coincident changes to the patterns of file updates that are shared across multiple hosts. We have implemented this approach in a prototype system called *Seurat* and demonstrated its effectiveness using a combination of real workstation cluster traces gathered over three months, simulated attacks, and a manually launched Linux worm.

**Keywords:** Anomaly detection, Pointillism, Correlation, File updates, Clustering

## 1 Introduction

Correlation is a recognized technique for improving the effectiveness of intrusion detection by combining information from multiple sources. For example, many existing works have proposed correlating different types of logs gathered from distributed measurement points on a network (e.g., [1–3]). By leveraging collective information from different local detection systems, we can detect more attacks while reducing the false positive rates.

In this paper, we propose a new approach to anomaly detection based on the idea of correlating host state transitions such as memory usage or file system updates. The idea is to correlate host state transitions across both space (multiple hosts) and time (the past and the present), detecting similar coincident changes to the patterns of host state updates that are shared across multiple hosts. Examples of such coincident events include administrative updates that modify files that have not been modified before, and worm propagations that cause certain log files, which are modified daily, to cease being updated.

Our approach is based on the key observation that changes in host state in a network system often have both temporal and spatial locality. Both administrative updates and malware propagation exhibit spatial locality, in the sense that similar updates tend to occur across many of the hosts in a network. They also exhibit temporal locality in the

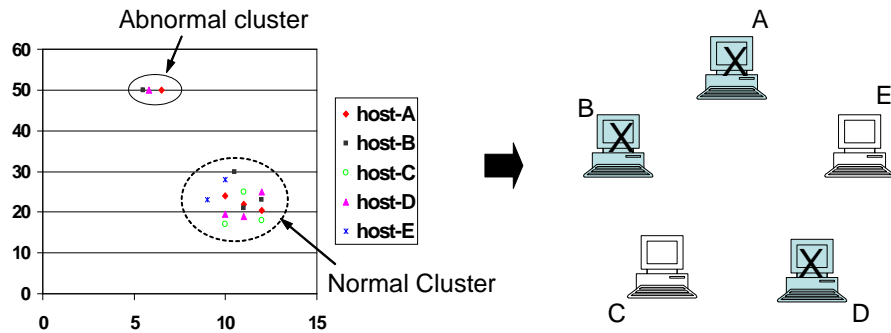
sense that these updates tend to be clustered closely in time. Our goal is to identify atypical such aggregate updates, or the lack of typical ones.

By exploring both the temporal and spatial locality of host state changes in a network system, our approach identifies anomalies *without foreknowledge of normal changes* and *without system-specific knowledge*. Existing approaches focus on the temporal locality of host state transitions, while overlooking the spatial locality among different hosts in a network system. They either define a model of normal host state change patterns through learning, or specify detailed rules about normal changes. The learning based approaches train the system to learn characteristics of normal changes. Since they focus only on the temporal locality of single-host state transitions, any significant deviation from the normal model is suspicious and should raise an alarm, resulting in high false alarms. The rule-based approaches require accurate, specific knowledge of system configurations and daily user activity patterns (e.g., Tripwire [4]) on a specific host. Violation of rules suggests malicious intrusions. Although rule-based anomaly detection raises fewer false alarms, it requires system administrators to manually specify a set of rules for each host. The correlation capability of our approach across both space and time allows us to learn the patterns of normal state changes over time, and to detect those anomalous events correlated among multiple hosts due to malicious intrusions. This obviates the need for specific rules while eliminating the false alarms caused by single host activity pattern shifts.

The correlation is performed by clustering points, each representing an individual host state transition, in a multi-dimensional feature space. Each feature indicates the change of a file attribute or memory state, with all features together describing the host state transitions of an individual machine during a given period (e.g., a day). Over time, the abstraction of point patterns inherently reflects the aggregated host activities. For normal host state changes, the points should follow some regular pattern by roughly falling into several clusters. Abnormal changes, which are hard to detect by monitoring that host alone, will stand out when they are correlated with other normal host state changes. Hence our approach shares some flavor of *pointillism* – a style of painting that applies small dots onto a surface so that from a distance the dots blend together into meaningful patterns. And we name our prototype system *Seurat*, who is the 19th century founder of pointillism [5].

Figure 1 illustrates the pointillist approach to anomaly detection. There are five hosts in the network system. We represent each host state changes daily as a point in a 2-dimensional space. On normal days, the points roughly fall into the dash-circled region. The appearance of a new cluster consisting of three points (indicated by the solid circle) suggests the incidence of anomaly on host A, B, and D, which may all have been compromised by the same attack. Furthermore, if we know that a certain host (e.g., host A) is already compromised (possibly detected by other means such as a network based IDS), then we can correlate the state changes of the compromised hosts with the state changes of all other hosts in the network system to detect more infected hosts (e.g., host B and D).

We have implemented a prototype system, called *Seurat*, that uses file system updates to represent host state changes for anomaly detection. *Seurat* successfully pinpoints the existence of a real Linux worm on a number of hosts in an isolated cluster



**Fig. 1. Pointillist approach to anomaly detection:** Normal points are clustered by the dashed circle. The appearance of a new cluster consisting of three points suggests anomalous events on host A, B, and D.

where we manually launched the attack. Seurat has a low false alarm rate when evaluated by a real deployment. These alarms are caused by either system re-configurations or network wide experiments. The false negative rate and detection latency, evaluated with simulated attacks, are both low for fast propagating attacks. For slowly propagating attacks, there is a tradeoff between false negative rate and detection latency. For each alarm, Seurat identifies the list of hosts involved and the related files, which we expect will be extremely helpful for system administrators to examine the root cause and dismiss alarms.

The rest of the paper is organized as following: Section 2 introduces the threat model of Seurat. Section 3 describes the algorithm we used for correlating host state changes across both space and time. Section 4 evaluates our approach. Section 5 discusses the limitations of Seurat and suggests possible improvements. Finally, Section 6 presents related work before we conclude in Section 7.

## 2 Attack Model

The goal of Seurat is to automatically identify anomalous events by correlating the aggregated state change events of all hosts in a network system. Hence Seurat defines an anomalous event as an unexpected state change close in time across *multiple* hosts in a network system.

We focus on rapid propagating Internet worms, virus, zombies or other malicious attacks that compromise multiple hosts in a network system at a time (e.g., one or two days). We have observed that, once fast, automated attacks are launched, most of the vulnerable hosts get compromised due to the rapid propagation of the attack and scanning preference of the automated attack tools. According to CERT's analysis [6], the level of automation in attack tools continues to increase. The automation speeds vulnerable hosts searching and the attack propagation. Recent Slammer worm hit 90 percent of vulnerable systems in the Internet within 10 minutes [7]. Much worse, the lack of diversity in systems and software run by Internet-attached hosts enables massive and

fast attacks. Computer clusters tend to be configured with the same operating systems and softwares. This results in high level of temporal and spatial locality of host state changes by attacks, where the efficacy of Seurat comes.

Although Seurat is originally designed to detect the system changes due to fast propagating attacks, Seurat can be generalized to detect slowly propagating attacks as well. This can be done by varying the time resolution of reporting and correlating the collective host state changes. We will discuss this issue further in Section 5. However, Seurat’s global correlation can not detect abnormal state changes that are unique to only one single host in the network system.

Seurat represents host state changes using *file system updates*. Pennington et al. [8] found that 83% of the intrusion tools and network worms they surveyed modify one or more system files. These modifications would be noticed by monitoring file system updates. There are many security tools such as Tripwire [4] and AIDE [9] that rely on monitoring file systems integrity violations for intrusion detection.

We denote each file in the network system as its complete path name when correlating the file updates. Different versions of a file represented as a common path name will be regarded as a same file across different hosts. This is because we are mostly interested in system files which tend to have canonical path names exploited by malicious attacks. We treat files, with different path names on different hosts, as user files even when they are identical in content. User files would rarely be updated by an intrusion and the global updates of user files would be infrequent.

For the detection of anomalies caused by attacks, we found this representation of host state changes is effective and useful. However, we may need different approach to represent host state transitions for other applications of Seurat such as file sharing detection, or for the detection of more sophisticated future attacks that alter files at arbitrary locations as they propagate. As an ongoing work, we are investigating the use of file content digests instead of file names.

### 3 Correlation based Anomaly Detection

We define a  $d$ -dimensional feature vector  $\mathbf{H}_{ij} = \langle v_1, v_2, \dots, v_d \rangle$  to represent the file system update attributes for host  $i$  on day  $j$ . Each  $\mathbf{H}_{ij}$  can be plotted as a point in a  $d$ -dimensional feature space. Our pointillist approach is based on correlating the feature vectors by clustering. Over time, for normal file updates, the points follow some regular pattern (e.g., roughly fall into several clusters). From time to time, Seurat compares the newly generated points against points from previous days. The appearance of a new cluster, consisting only of newly generated points, indicates abnormal file updates and Seurat raises an alarm.

For clustering to work most effectively, we need to find the most relevant features (dimensions) in a feature vector given all the file update attributes collected by Seurat. We have investigated two methods to reduce the feature vector dimensions: (1) *wavelet-based selection*, and (2) *principle component analysis* (PCA).

In the rest of this section, we first present how we define the feature vector space and the distances among points. We then describe the methods Seurat uses to reduce

feature vector dimensions. Finally, we discuss how Seurat detects abnormal file updates by clustering.

### 3.1 Feature Vector Space

Seurat uses binary feature vectors to represent host file updates. Each dimension in the feature vector space corresponds to a unique file (indexed by the file name concatenated with its path). We define the *detection window* to be the period that we are interested in finding anomalies. In the current prototype, the detection window is one day. For each vector  $\mathbf{H}_{ij} = \langle v_1, v_2, \dots, v_d \rangle$ , we set  $v_k$  to 1 if host  $i$  has updated (added, modified, or removed) the  $k$ -th file on day  $j$ , otherwise, we set  $v_k$  to 0.

The vectors generated in the detection window will be correlated with vectors generated on multiple previous days. We treat each feature vector as an independent point in a set. The set can include vectors generated by a same host on multiple days, or vectors generated by multiple hosts on a same day. In the rest of the paper, we use  $\mathbf{V} = \langle v_1, v_2, \dots, v_d \rangle$  to denote a feature vector for convenience. Figure 2 shows how we represent the host file updates using feature vectors.


	F <sub>1</sub>	F <sub>2</sub>	F <sub>3</sub>	F <sub>4</sub>	F <sub>5</sub>
$\mathbf{V}_1 = \mathbf{H}_{11} =$	$\langle 1,$	$1,$	$0,$	$1,$	$1 \rangle$
$\mathbf{V}_2 = \mathbf{H}_{21} =$	$\langle 1,$	$1,$	$1,$	$0,$	$0 \rangle$
$\mathbf{V}_3 = \mathbf{H}_{12} =$	$\langle 1,$	$1,$	$0,$	$1,$	$0 \rangle$

**Fig. 2. Representing host file updates as feature vectors:**  $F_1, F_2, F_3, F_4, F_5$  are five different files with unique path names. Accordingly, the feature vector space has 5 dimensions in the example.

The correlation is based on the distances among vectors. Seurat uses a cosine distance metric, which is a standardized similarity measure between vectors [10, 11]. We define the distance  $D(\mathbf{V}_1, \mathbf{V}_2)$  between two vectors  $\mathbf{V}_1$  and  $\mathbf{V}_2$  as their angle  $\theta$  computed by the cosine value:

$$D(\mathbf{V}_1, \mathbf{V}_2) = \theta = \cos^{-1}\left(\frac{\mathbf{V}_1 \cdot \mathbf{V}_2}{|\mathbf{V}_1||\mathbf{V}_2|}\right)$$

For each day  $j$  (the detection window), Seurat correlates the newly generated vectors with vectors generated in a number of previous days  $j-1, j-2, \dots$ . We assume that the same abnormal file update events on day  $j$ , if any, have not occurred on those previous days. We define the *comparison window* of day  $j$  as the days that we look back for comparison, and the *correlation window* of day  $j$  as the inclusive period of day  $j$  and its comparison window. Vectors generated outside the correlation window of day  $j$  are not used to identify abnormal file updates on day  $j$ . Figure 3 illustrates the concepts of detection window, comparison window, and correlation window.



**Fig. 3. Detection window, comparison window, and correlation window.** The detection window is day  $j$ . The comparison window is from day  $j - t$  to day  $j - 1$ . The correlation window is from day  $j - t$  to day  $j$ .

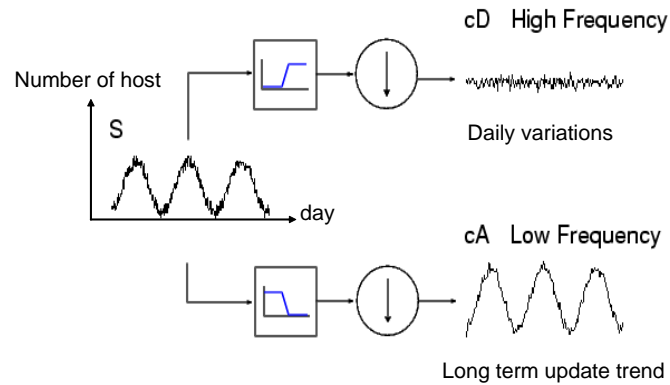
Since each vector generated during the comparison window serves as an example of normal file updates to compare against in the clustering process, we explore the temporal locality of normal update events by choosing an appropriate comparison window for each day. The comparison window size is a configurable parameter of Seurat. It reflects how far we look back into history to implicitly define the model of normal file updates. For example, some files such as `/var/spool/anacron/cron.weekly` on Linux platforms are updated weekly. In order to regard such weekly updates as normal updates, administrators have to choose a comparison window size larger than a week. Similarly, the size of detection window reflects the degree of temporal locality of abnormal update events.

Since Seurat correlates file updates across multiple hosts, we are interested in only those files that have been updated by at least two different hosts. Files that have been updated by only one single host in the network system throughout the correlation window are more likely to be user files. And we do not select them as relevant dimensions to define the feature vector space.

### 3.2 Feature Selection


Most of the file updates are irrelevant to anomalous events even after we filter out the file updates reported by a single host. Those files become noise dimensions when we correlate the vectors (points) to identify abnormal updates, and increase the complexity of the correlation process. We need more selective ways to choose relevant files and reduce feature vector dimensions. Seurat uses wavelet-based selection method and principle component analysis (PCA) for this purpose.

**Wavelet-based Selection** The wavelet-based selection method regards each individual file update status as a discrete time series signal  $S$ . Given a file  $i$ , the value of the signal on day  $n$  denoted by  $S_i(n)$  is defined as the total number of hosts that update file  $i$  on day  $n$  in the network system. Each such signal  $S_i$  can be decomposed into a low frequency signal  $cA_i$  reflecting the long term update trend, and a high frequency signal  $cD_i$  reflecting the daily number of host variations (see Figure 4). If the high frequency signal  $cA_i$  shows a spike on a certain day, we know that a significantly larger number of hosts updated file  $i$  than on a normal day. We then select file  $i$  as a relevant feature dimension in defining the feature vector space.




**Fig. 4. Representing file update status with wavelet transformation:** The original signal is  $S$ , which can be decomposed into a low frequency signal  $cA$  reflecting the long term update trend, and a high frequency signal  $cD$  reflecting the daily variations.

There are several different algorithms proposed to remove long term periodicity using wavelet transforms. Seurat detects the signal spikes using the residual signal of long-term trend. Such technique was used to detect disease outbreaks[12]. To detect anomalies on day  $j$ , the algorithm takes as input the list of files that have been updated by at least two different hosts in the correlation window of day  $j$ , and selects from the input a subset of files that will be used to define the feature vector space.



**Fig. 5. Wavelet-based feature selection**

Figure 5 shows the steps to select features by wavelet-based method. Given a fixed correlation window of day  $j$ , the algorithm starts with constructing a time series signal  $S_i$  for each file  $i$ , and decomposes  $S_i$  into  $cA_i$  and  $cD_i$  using wavelet transformation as described. Then we compute the residual signal value  $R_i(j)$  of day  $j$  by subtracting the trend value  $cA_i(j-1)$  of day  $j-1$  from the original signal value  $S_i(j)$  of day  $j$ . If  $R_i(j)$  exceeds a pre-set threshold  $\alpha$ , then the actual number of hosts who have updated file  $i$  on day  $j$  is significantly larger than the prediction  $cA_i(j-1)$  based on the long term trend. Therefore, Seurat selects file  $i$  as an interesting feature dimension for anomaly detection on day  $j$ . As an example, Figure 6 shows the original signal and the residual signal of a file using a 32-day correlation window in a 22-host teaching cluster.



**Fig. 6. Wavelet transformation of file update status: (a) The original signal of the file update status (b) The residual signal after wavelet transformation**

Note the threshold value  $\alpha$  of each file is a parameter selected based on the statistical distribution of historical residual values.

**PCA based Dimension Reduction** PCA is a statistical method to reduce the data dimensionality without much loss of information [13]. It finds a set of  $d'$  orthogonal vectors, called principle components in data space, that account for the variance of the input data as much as possible. Dimensionality reduction is then achieved by projecting the original  $d$ -dimensional data onto the subspace spanned by these  $d'$  orthogonal vectors. Most of the intrinsic information of the  $d$ -dimensional data is preserved in the  $d'$ -dimensional subspace.

We note that the updates of different files are usually correlated. For example, when a software package is updated on a host, many of the related files will be modified together. Thus we can perform PCA to identify the correlation of file updates.

Given a  $d$ -dimensional feature space  $\mathcal{Z}_2^d$ , and a list of  $m$  feature vectors  $\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_m \in \mathcal{Z}_2^d$ , we perform the following steps using PCA to obtain a new list of feature vectors  $\mathbf{V}'_1, \mathbf{V}'_2, \dots, \mathbf{V}'_m \in \mathcal{Z}_2^{d'}$  ( $d' < d$ ) with reduced number of dimensions:

1. Standardize each feature vector  $\mathbf{V}_k = \langle v_{1k}, v_{2k}, \dots, v_{dk} \rangle$  ( $1 \leq k \leq m$ ) by subtracting each of its elements  $v_{ik}$  by the mean value of the corresponding dimension  $u_i$  ( $1 \leq i \leq d$ ). We use  $\bar{\mathbf{V}}_k = \langle \bar{v}_{1k}, \bar{v}_{2k}, \dots, \bar{v}_{dk} \rangle \in \mathcal{Z}_2^d$  to denote the standardized vector for the original feature vector  $\mathbf{V}_k$ . Then,

$$\bar{v}_{ik} = v_{ik} - u_i \quad (u_i = \frac{\sum_{j=1}^m v_{ij}}{m}, 1 \leq i \leq d)$$

2. Use the standardized feature vectors  $\bar{\mathbf{V}}_1, \bar{\mathbf{V}}_2, \dots, \bar{\mathbf{V}}_m$  as input data to PCA in order to identify a set of principle components that are orthogonal vectors defining a set of transformed dimensions of the original feature space  $\mathcal{Z}_2^d$ . Select the first  $d'$  principle components that count for most of the input data variances (e.g., 90% of data variances) to define a subspace  $\mathcal{Z}_2^{d'}$ .



3. Project each standardized feature vector  $\bar{\mathbf{V}}_k \in \mathcal{Z}_2^d$  onto the PCA selected subspace  $\mathcal{Z}_2^{d'}$  to obtain the corresponding reduced dimension vector  $\mathbf{V}'_k \in \mathcal{Z}_2^{d'}$ .

Note that PCA is complementary to wavelet-based selection. Once we fix the correlation window of a particular day, we first pick a set of files to define the feature vector space by wavelet-based selection. We then perform PCA to further reduce the data dimensionality.

### 3.3 Anomaly Detection by Clustering

For every day, once we obtain a list of transformed feature vectors using feature selection described in the previous subsection, we cluster the vectors based on the distance between every pair of them.

We call the cluster a *new cluster* if it consists of multiple vectors only from the days in the detection window. The appearance of a new cluster indicates possibly abnormal file updates occurred on that day and should raise an alarm.

There are many existing algorithms for clustering, for example, K-means [14, 15] or Single Linkage Hierarchical Clustering [11]. Seurat uses a simple iterative algorithm, which is a common method for K-means initialization, to cluster vectors without prior knowledge of the number of clusters [16]. The algorithm assumes each cluster has a hub. A vector belongs to the cluster whose hub is closest to that vector compared with the distances from other hubs to that vector. The algorithm starts with one cluster whose hub is randomly chosen. Then, it iteratively selects a vector that has the largest distance to its own hub as a new hub, and re-clusters all the vectors based on their distances to all the selected hubs. This process continues until there is no vector whose distance to its hub is larger than the half of the average hub-hub distance.

We choose this simple iterative algorithm because it runs much faster, and works equally well as the Single Linkage Hierarchical algorithm in our experiments. The reason that even the simple clustering algorithm works well is that the ratio of inter-cluster distance to intra-cluster distance significantly increases after our feature selection. Since the algorithm currently used is sensitive to outliers, we plan to explore and test other more robust algorithms for the best clustering results.

Once we detect a new cluster and so generate an alarm, we examine further to identify the involved hosts and the root cause files. The suspicious hosts are just the ones whose file updates correspond to the feature vectors in the new cluster. To determine which files possibly cause the alarm, we only focus on the files picked by the wavelet-based selection to define the feature vector space. For each of those files, if it is updated by all the hosts in the new cluster on the day under detection, but has not been updated by any host during the corresponding comparison window, Seurat outputs this file as a candidate file.

Based on the suspicious hosts and the candidate files for explaining root causes, system administrators can decide whether the updates are known administrative updates that should be suppressed, or some abnormal events that should be further investigated. If the updates are caused by malicious attacks, administrators can take remedial counter measures for the new cluster. Furthermore, additional compromised hosts can be identified by checking if the new cluster expands later and if other hosts have updated the same set of candidate files.

## 4 Experiments

We have developed a multi-platform (Linux and Windows) prototype of Seurat that consists of a light-weight data collection tool and a correlation module. The data collection tool scans the file system of the host where it is running and generates a daily summary of file update attributes. Seurat harvests the summary reports from multiple hosts in a network system and the correlation module uses the reports for anomaly detection.

We have installed the Seurat data collection tool on a number of campus office machines and a teaching cluster that are used by students daily. By default, the tool scans the attributes of all system files on a host. Personal files under user home directories are exempted from scanning for privacy protection. The attributes of a file include the file name, type, device number, permissions, size, I-node number, important timestamps, and a 16-byte MD5 checksum of file content. The current Seurat uses binary information about file updates, but the next version of Seurat may make use of other attributes reported by the data collection tool. Everyday, each host compares the newly scanned disk snapshot against that from the previous day and generates a file update summary report. In the current prototype, all the reports are daily uploaded to a centralized server where system administrators can monitor and correlate the file updates using the correlation module.

In this section, we study the effectiveness of Seurat's pointillist approach for detecting aggregated anomalous events. We use the daily file update reports from our real deployment to study the *false positive rate* and the corresponding causes in Section 4.1. We evaluate the *false negative rate* with simulated attacks in Section 4.2. In order to verify the effectiveness of our approach on real malicious attacks, we launched a real Linux worm into an isolated cluster and report the results in Section 4.3.

### 4.1 False Positives

The best way to study the effectiveness of our approach is to test it with real data. We have deployed Seurat on a teaching cluster of 22 hosts and have been collecting the daily file update reports since Nov 2003. The teaching cluster is mostly used by students for their programming assignments. They are also occasionally used by a few graduate students for running network experiments.

For this experiment, we use the file update reports from Dec 1, 2003 until Feb 29, 2004 to evaluate the false positive rate. During this period, there are a few days when a couple of hosts failed to generate or upload reports due to system failure or re-configurations. For those small number of missing reports, we simply ignore them because they do not affect the aggregated file update patterns.


We set the correlation window to 32 days in order to accommodate monthly file update patterns. That is, we correlate the update pattern from day 1 - day 32 to identify abnormal events on day 32, and correlate the update pattern from day 2 - day 33 to detect anomaly on day 33, etc. Thus our detection starts from Jan 1, 2004 since we do not have enough long correlation windows for the days in Dec 2003.

**Dimension Reduction** Once we fixed the correlation window of a particular day, we perform a single-level wavelet transform on each file and use the low frequency signal

as its long term trend. We select relevant files using a constant threshold  $\alpha$  of 2 to define the feature vector space for simplicity. We then perform PCA to reduce the data dimensionality further by picking the first several principle components that accounts for 98% of the input data variances.

Throughout the entire period of 91 days, there were, in total, 772 file with unique file names updated by at least two different hosts. Figure 7 (a) shows the number of hosts that updated each file during the data collection period. We observe that only a small number files (e.g., `/var/adm/syslog/mail.log`) are updated very regularly by all of the hosts, while most other files (e.g., `/var/run/named.pid`) are updated irregularly depending on the system usage or the applications running.


Figure 7 (b) shows the results of feature selection. There were, on average, 140 files updated by at least two different hosts during each correlation window. After wavelet-based selection, the average number of feature dimensions is 17. PCA further reduces the vector space dimension to below 10.



**Fig. 7. Feature selection and dimension reduction:** (a) File update patterns. Files are sorted by the cumulative number of hosts that have updated them throughout the 91 days. The darker the color is, the more hosts updated the corresponding file. (b) The number of feature vector dimensions after wavelet-based selection and PCA consecutively.

**False Alarms** After dimension reduction, we perform clustering of feature vectors and identify new clusters for each day. Figure 8 illustrates the clustering results of 6 consecutive days from Jan 19, 2004 to Jan 24, 2004. There are two new clusters identified on Jan 21 and Jan 23, which involve 9 hosts and 6 hosts, respectively. Since Seurat outputs a list of suspicious files as the candidate root cause of each alarm, system administrators can tell if the new clusters are caused by malicious intrusions.

Based on the list of files output by Seurat, we could figure out that the new clusters on Jan 21 and Jan 23 reflect large scale file updates due to a system re-configuration at the beginning of the spring semester. For both days, Seurat accurately pinpoints the exact hosts that are involved. The re-configuration started from Jan 21, when a large



**Fig. 8. Clustering feature vectors for anomaly detection:** Each circle represents a cluster. The number at the center of the figure shows the total number of clusters. The radius of a circle corresponds to the number of points in the cluster, which is also indicated beside the circle. The squared dots correspond to the new points generated on the day under detection. New clusters are identified by a thicker circle.

number of binaries, header files, and library files were modified on 9 out of the 22 hosts. Since the events are known to system administrators, we treat the identified vectors as normal for future anomaly detection. Thus no alarm is triggered on Jan 22, when the same set of library files were modified on 12 other hosts. On Jan 23, the re-configuration continued to remove a set of printer files on 6 out of the 22 hosts. Again, administrators can mark this event as normal and we spot no new cluster on Jan 24, when 14 other hosts underwent the same set of file updates.

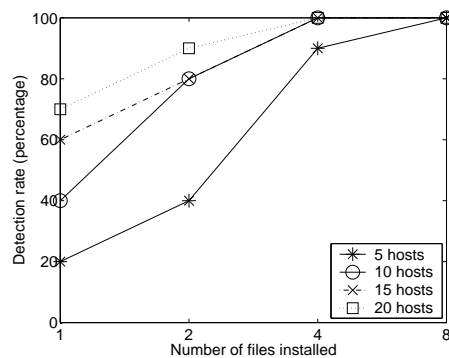
In total, Seurat raises alarms on 9 out of the 60 days under detection, among which 6 were due to system re-configurations. Since the system administrators are aware of such events in advance, they can simply suppress these alarms. The 3 other alarms are generated on 3 consecutive days when a graduate student performed a network experiment that involved simultaneous file updates at multiple hosts. Such events are normal but rare, and should alert the system administrators.

## 4.2 False Negatives

The primary goal for this experiment is to study the false negative rate and detection latency of Seurat as the stealthiness of the attack changes. We use simulated attacks by manually updating infected files on the selected host reports.

We first examine the detection rate of Seurat by varying the degree of attack aggressiveness. We model the attack propagation speed as the number of hosts infected

on each day (i.e., the detection window), and model the attack stealthiness on a local host as the number of new files installed by this attack. Our simulation runs on the same teaching cluster that we described in Section 4.1. Since the aggregated file update patterns are different for each day, we randomly pick ten days in Feb 2004, when there is no intrusions. On each selected day, we simulate attacks by manually inserting artificial new files into a number of host reports on only that day, and use the modified reports as input for detection. We then remove those modified entries, and repeat the experiments with another day. The detection rate is calculated as the number of days that Seurat spots new clusters over the total ten days.



**Fig. 9. Detection rate:** We vary the number of hosts infected and the number of files inserted on each host by the simulated attacks.

Figure 9 shows the detection rate of Seurat by varying the number of files inserted on each host and the number of hosts infected. The detection rate monotonically increases as we increase the number of files inserted on each host by an attack. Since the inserted files do not exist before, each of them will be selected as a feature dimension by the wavelet-based selection, leading to larger distances between the points of infected host state changes and the points of normal host state changes. Therefore, the more new files are injected by an attack, the higher the detection rate gets. On another aspect, as we increase the number of infected hosts, the number of points for abnormal host state changes becomes large enough to create an independent new cluster. Thus rapidly propagating attacks are more likely to be caught. Accordingly, we need longer latency to have enough number of infected hosts for detecting a slowly propagating attack, which we will further discuss in Section 5.

We further evaluate the detection rate of Seurat on six Linux worms with simulated attacks. To do so, we compile a subset of files modified by each worm based on the descriptions from public Web sites such as Symantec [17] and F-Secure information center [18]. We then manually modify the described files in a number of selected host reports to simulate the corresponding worm attacks. Again, for each worm, we vary the number of infected hosts, and run our experiments on the teaching cluster with ten randomly selected days.

Worms	Adore	Ramen-A	Ramen-B	Slapper-A	Slapper-B	Kork
Files modified	10	8	12	3	4	5
2 infected hosts	80%	80%	90%	30%	40%	30%
4 infected hosts	100%	100%	90%	70%	80%	70%
8 infected hosts	100%	100%	100%	100%	100%	100%

**Fig. 10. Detection rate of emulated worms:** We vary the number of hosts compromised by the attacks.

Table 10 shows the number of files modified by each worm and the detection rate of Seurat. In general, the more files modified by a worm, the more likely the worm will be detected. But the position of a file in the file system directory tree also matters. For example, both Slapper-B worm and Kork worm insert 4 new files into a compromised host. However, Kork worm additionally modifies `/etc/passwd` to create accounts with root privileges. Because there are many hosts that have updated `/etc/passwd` during a series system re-configuration events, the inclusion of such files in the feature vector space reduces the distances from abnormal points to normal points, resulting in higher false negative rates. We further discuss this in Section 5 *mimicry attacks*.


### 4.3 Real Attacks

Now we proceed to examine the efficacy of Seurat during a real worm outbreak. The best way to show this would be to have Seurat detect an anomaly caused by a new worm propagation. Instead of waiting for a new worm’s outbreak, we have set up an isolated computer cluster where, without damaging the real network, we can launch worms and record file system changes. This way, we have full control over the number of hosts infected, and can repeat the experiments. Because the isolated cluster have no real users, we merge the data acquired from the isolated cluster to the data we have collected from the teaching cluster for anomaly detection.

We could obtain the binaries and source codes of a few popular worms from public Web sites such as whitehats [19] and packetstormes [20]. Extensively testing Seurat, with various real worms in the isolated cluster, requires tremendous effort in setting up each host with the right versions of vulnerable software. As a first step, we show the result with Lion worm [21] in this experiment.

Lion worm was found in early 2001. Lion exploits a vulnerability of BIND 8.2, 8.2-P1, 8.2.1, 8.2.2-Px. Once Lion infects a system, it sets up backdoors, leaks out confidential information (`/etc/passwd`, `/etc/shadow`) via email, and scans the Internet to recruit vulnerable systems. Lion scans the network by randomly picking a first 16 bits of an IP address, and then sequentially probing all the  $2^{16}$  IP addresses in the space of the block. After that, Lion selects another such address block randomly to continue scanning. As a result, once a host is infected by Lion, all the vulnerable hosts nearby (in the same IP address block) will be infected soon. Lion affects file systems: the worm puts related binaries and shell scripts under `/dev/.lib` directory, copies itself into `/tmp` directory, changes system files under `/etc` directory, and tries to wipe out some log files.

We configure the isolated cluster to have three machines that are vulnerable to Lion and one machine that launches Lion worm. The vulnerable machines are running Red-Hat 6.2 including the vulnerable BIND 8.2.2-P5.<sup>3</sup> The cluster uses one C class network address block. Every machine in the cluster is connected to an 100Mbps ethernet and is running *named*. The Seurat data collection tool daily checks the file system updates and generates a summary report on every machine.



**Fig. 11. Intrusion detection by Seurat:** Seurat identifies a new cluster of three hosts on Feb 11, 2004, when we manually launched Lion worm

After we launched the Lion worm, all three vulnerable hosts in the isolated cluster got infected quickly one after another. We merge the file update report by the each compromised host with a different normal host report generated on Feb 11, 2004, when we know there is no anomaly. Figure 11 shows the clustering results of three consecutive days from Feb 10, 2004 to Feb 12, 2004 using the merged reports.

On the attack day, there are 64 files picked by the wavelet-based selection. The number of feature dimensions is reduced to 9 after PCA. Seurat successfully detects a new cluster consisting of the 3 infected hosts. Figure 12 lists the 22 files selected by Seurat as the root causes of the alarm. These files provide enough hints to the administrators to confirm the existence of Lion worm. Once detected, these compromised hosts as well as the list of suspicious files can be marked for future detection. If, in the following days, there are more hosts that are clustered together with the already infected machines, or experience the same file updates, then we may conclude they are infected by the same attack.

## 5 Discussion

### 5.1 Vulnerabilities and Limitations

*Stealthy attack.* Attackers may try to evade the detection by slowing down attack propagations. If an attacker is patient enough to infect only one host a day in the monitored network system, Seurat will not notice the intrusion with the current one-day detection

<sup>3</sup> We tried RedHat 7.1 with the same version of BIND, but it was not vulnerable.

File ID.	File name	File ID.	File name
1	/sbin/asp	12	/var/spool/mail
2	/dev/.lib	13	/dev/.lib/bindx.sh
3	/dev/.lib/star.sh	14	/tmp/ramen.tgz
4	/var/spool/mail/root	15	/dev/.lib/scan.sh
5	/dev/.lib/bind	16	/dev/.lib/pscan
6	/etc/hosts.deny	17	/var/spool/mqueue
7	/dev/.lib/randb	18	/dev/.lib/hack.sh
8	/sbin	19	/dev/.lib/.hack
9	/var/log	20	/dev/.lib/index.html
10	/dev/.lib/bindname.log	21	/dev/.lib/asp62
11	/dev/.lib/index.htm	22	/var/log/sendmail.st

**Fig. 12. Suspicious files for the new cluster on Feb 11, 2004**

window because Seurat focuses only on anomalous file changes across multiple hosts. To catch slow, stealthy attacks, Seurat can employ a larger detection window such as a couple of days or a week. Note, however, that Seurat can notice the attacks only after multiple hosts in the network system are compromised. In other words, if an attack propagate slowly, Seurat may not recognize the attack for the first few days after the initial successful compromise. There is thus a tradeoff between detection rate and detection latency.

*Mimicry attack.* An attacker can carefully design his attack to cause file updates that look similar to regular file changes, and mounts a successful *mimicry attack* [22]. There are two ways to write a mimicry attack against the current prototype. First, an attacker may try to fool Seurat's feature selection process by camouflaging all intrusion files as frequently, regularly updated files. Those concealed files, even when they are modified in an unexpected way (e.g., entries removed from append-only log files), will not be selected as feature vector dimensions because of current use of the binary feature representation. Note that Seurat's data collection tool provides additional information on file system changes, such as file size, file content digest, and permissions. By incorporating the extra information in representing host state transition, Seurat can make such mimicry attacks harder to evade detection. Second, an attacker may find a way to cloak abnormal file updates with many normal but irregular changes during Seurat's clustering process. For example, in Section 4.2, we observed that the false negative of Kork worm detection was relatively high due to the interference of irregular system re-configuration. We leave it as future work to quantify this type of mimicry attack and the effectiveness of possible counter measures.

*Random-file-access attack.* Seurat correlates file updates based on their complete path names. Thus attackers can try to evade Seurat by installing attack files with random file names, placing them under different directories at different hosts, or replacing randomly chosen existing files with attack files. Many recent email viruses already change the virus file names when they propagate to a new host; we envision similar techniques could be employed by other types of attacks soon. Note, however, that even the random-file-access attack may need a few anchor files at fixed places, in which case



Seurat still has the opportunity to detect such attack. A more robust representation of a file, for examples, MD5 checksum, could help Seurat detect random-file-access attack more easily.

*TSR (Terminate and Stay Resident) attack.* Memory-resident and BIOS-resident only attacks make no file system changes. Thus Seurat will not be able to detect TSR attacks by examining host file updates, nor those attacks that erase disk evidence before the Seurat data collection tool performs the scheduled disk scan.

*Kernel/Seurat modification attack.* The effectiveness of Seurat relies on the correctness of reports from the data collecting tools running on distributed hosts. So the host kernels and the Seurat data collection tools should run on machines protected by trusted computing platforms [23].

## 5.2 Enhancement

*Collaboration with other types of IDS.* By correlating host file update events across both space and time, Seurat makes use of the system collective information from a different angle. It can not only be used as an independent tool for anomaly detection, but can co-work with existing IDS systems. For instance, Seurat can be used to identify other infected hosts that are similar to an already known compromised machine that is detected by other IDS. Or other IDS can take advantage of the list of root cause files, which are identified by Seurat in anomaly detection, as a new detection rule for catching more future attacks. Thus Seurat is complementary to existing security tools.

*Distributed correlation module.* Currently, Seurat moves the daily reports from distributed data collection tools to a centralized server, where the correlation module computes and clusters the host vectors. The centralized deployment simplifies the correlation algorithm and provides the correlation module with the detail of every host change. However, the centralized approach exposes Seurat to problems in scalability and reliability. First, one obvious drawback is the large amount of data to be transferred to and stored at the centralized server. In our experience, a host generates a file update report of 3K-140KBytes daily in a compressed format, so the aggregated report size from hundreds or thousands of hosts with a long comparison window will be large. The report size will be larger when Seurat's data collection tool checks the host state changes more frequently, or when hosts experience more active file updates everyday. Second, the monitored hosts could be in different administrative domains (i.e., hosts managed by different academic departments or labs) and it may be often impractical to transfer the detailed reports from all the hosts to one centralized server due to privacy and confidentiality issues. Third, the centralized server can be a single point-of-failure. It is important for Seurat to work even when one correlation server is broken or a part of network is partitioned. A distributed correlation module will cope with those issues. We are now investigating the way to extend our correlation algorithm into a distributed fashion leveraging the existing architectures [1, 24, 25] proposed to enable global observation shared among distributed participants.

## 6 Related Work

Seurat uses file system integrity to represent a host state. The file system integrity has been known to be a useful information for intrusion detection. Tripwire [4], AIDE [9], Samhain [26] are well-known intrusion detection systems that use the file system integrity to find intrusions. Recently proposed systems such as the storage-based intrusion detection system [8] and some commercial tools [27] support real-time integrity checking. All of them rely on the pre-defined rule set to detect anomalous integrity violation, while Seurat diagnoses the anomaly using learning and correlation across time and space.

Leveraging the information gathered from distributed multiple measurement points is not a new approach. Many researchers have noticed the potential of the collective approaches for intrusion detection or anomaly detection. Graph-based Intrusion Detection System (GrIDS) [28] project at UC Davis has developed graph engines that build a graph representation of network activity using the information reported by each host in the network. Different from Seurat, GrIDS uses the *TCP/IP network activity* between hosts in the network to infer patterns of intrusive or hostile activities based on *pre-defined* rules. GrIDS is known to be effective to detect massive, fast propagating worms. Other systems such as Cooperative Security Managers (CSM) [29], Distributed Intrusion Detection System (DIDS), also take advantage of collective approach to intrusion detection. They orchestrate multiple monitors watching multiple network links and track user activity across multiple machines.

EMERALD (Event Monitoring Enabling Responses to Anomalous Live Disturbances) [1] project of SRI and Autonomous Agents For Intrusion Detection (AAFID) [24] project at Purdue University have independently proposed distributed architectures for intrusion detection and response capability. Both of them use local monitors or agents to collect interesting events and anomaly reports (from a variety of sources; audit data, network packet trace, SNMP traffic, application logs, etc.). The architectures provide the communication methods to exchange the locally detected information and the easy way to manage components of the systems. AAFID performs statistical profile-based anomaly detection and EMERALD supports a signature-based misuse analysis in addition to the profile-based anomaly detection.

Note that Seurat starts with similar motivation. But Seurat focuses more on the technique for correlating the collective reports for anomaly detection, and infers interesting information on the system state from learning, rather than relying on pre-defined set of events or rules. We envision Seurat as a complementary technique, not as a replacement of the existing architectures that provide global observation sharing.

Correlating different types of audit logs and measurement reports is another active area in security research. Many researchers have proposed to correlate multiple heterogeneous sensors to improve the accuracy of alarms [3, 30, 31, 2, 32]. In this work, we attempt to correlate information gathered by homogeneous monitors (especially, the file system change monitors) but we may enhance our work to include different type of measurement data to represent individual host status.

Wang et al. [33] also have noticed the value of spatial correlation of multiple system configurations and applied a collective approach to tackle misconfiguration troubleshooting problems. In their system, a malfunctioning machine can diagnose its prob-

lem by collecting system configuration information from other similar and friendly hosts connected via a peer-to-peer network. The work does not target automatic detection of the anomaly, and rather it aims at figuring out the cause of a detected problem.

## 7 Conclusions

In this paper, we presented a new “pointillist” approach for detecting aggregated anomalous events by correlating collective information about host file updates across both space and time. Our approach explores the temporal and spatial locality of system state changes through learning and correlation. It requires neither prior knowledge about normal host activities, nor system specific rules.

A prototype implementation, called *Seurat*, suggests that the approach is effective in detecting rapidly propagating attacks that modify host file systems. The detection rate degrades as the stealthiness of attacks increases. By trading off detection latency, we are also able to identify hosts that are compromised by slowly propagating attacks. One important advantage of the *Seurat* prototype is its capability to diagnose the root cause of an anomaly by associating suspicious hosts and files with each alarm. Such information is extremely helpful to dismiss false alarms, which traditionally impede the wide deployment of learning based anomaly detection techniques.

While *Seurat* shows great promise in detecting attacks that involve file system changes, the representation of host state changes based on file paths limits the detection capability of *Seurat* on other types of attacks that incur no disk access. We are exploring other representations as our ongoing work. Other future work includes extensively testing *Seurat* with various real attacks, and generalizing the pointillist approach to other applications such as detecting illegal file sharing events for copyright protection.

## References

1. Porras, P.A., Neumann, P.G.: Emerald: Event monitoring enabling responses to anomalous live disturbances. In: Proceedings of the 20th National Information Systems Security Conference. (1997)
2. Abad, C., Taylor, J., Sengul, C., Zhou, Y., Yurcik, W., Rowe, K.: Log correlation for intrusion detection: A proof of concept. In: Proceedings of the 19th Annual Computer Security Applications Conference, Las Vegas, Nevada, USA (2003)
3. C.Kruegel, T.T., C.Kerer: Decentralized event correlation for intrusion detection. In: International Conference on Information Security and Cryptology (ICISC). (2001)
4. Tripwire, Inc.: Tripwire. (<http://www.tripwire.com>)
5. U. Kieseyer, M. Hulse, T.: George seurat: 1859 - 1891, the master of pointillism (2001)
6. Center, C.C.: Overview of attack trends. [http://www.cert.org/archive/pdf/attack\\_trends.pdf](http://www.cert.org/archive/pdf/attack_trends.pdf) (2002)
7. Moore, D., Paxson, V., Savage, S., Shannon, C., Staniford, S., Weaver, N.: Inside the slammer worm. *IEEE Security and Privacy* **1** (2003) 33–39
8. Pennington, A., Strunk, J., Griffin, J., Soules, C., Goodson, G., Ganger, G.: Storage-based intrusion detection: watching storage activity for suspicious behavior. In: Proceedings of 12th USENIX Security Symposium, Washington, DC (2003)
9. Lehti, R., Virolainen, P.: AIDE - Advanced Intrusion Detection Environment. (<http://www.cs.tut.fi/rammer/aide.html>)

10. Berry, M., Drmac, Z., Jessup, E.: Matrices, vector spaces, and information retrieval. *SIAM Review* **41** (1999)
11. Kamber, M.: Data mining: Concepts and techniques. Morgan Kaufmann Publishers (2000)
12. J.Zhang, F.Tsui, M.M.Wagner, Hogan, W.: Detection of outbreaks from time series data using wavelet transform. In: AMIA Fall Symp., Omni Press CD (2003) 748–752
13. I.T.Jolliffe: Principle component analysis. Springer-Verlag, New York (1986)
14. E.Forgy: Cluster analysis of multivariate data: Efficiency vs. interpretability of classifications. *Biometrics* **21** (1965)
15. Gersho, A., Gray, R.: Vector quantization and signal compression. Kluwer Academic Publishers (1992)
16. Moore, A.: K-means and Hierarchical Clustering. <http://www.cs.cmu.edu/awm/tutorials/kmeans09.pdf> (available upon request) (2001)
17. Symantec: Symantec security response. (<http://securityresponse.symantec.com>)
18. F-Secure: F-secure security information center. (<http://www.f-secure.com/virus-info>)
19. Whitehats, Inc.: Whitehats Network Security Resource. (<http://www.whitehats.com>)
20. PacketStorm: Packet Storm. (<http://www.packetstormsecurity.org>)
21. SANS Institute: Lion Worm. <http://www.sans.org/y2k/lion.htm> (2001)
22. Wagner, D., Dean, D.: Intrusion detection via static analysis. In: Proceedings of 2001 IEEE Symposium on Security and Privacy. (2001)
23. Trusted Computing Platform Alliance: Trusted Computing Platform Alliance. (<http://www.trustedcomputing.org>)
24. Balasubramanian, J.S., Garcia-Fernandez, J.O., Isacoff, D., Spafford, E., Zamboni, D.: An architecture for intrusion detection using autonomous agents. In: Proceedings of the 14th IEEE Computer Security Applications Conference. (1998)
25. Xie, Y., O'Hallaron, D.R., Reiter, M.K.: A secure distributed search system. In: Proceedings of the 11th IEEE International Symposium on High Performance Distributed Computing. (2002)
26. Samhain Labs: Samhain. (<http://la-samhna.de/samhain/>)
27. Pedestal Software: INTACT™. (<http://www.pedestalsoftware.com/products/intact/>)
28. Cheung, S., Crawford, R., Dilger, M., Frank, J., Hoagland, J., Levitt, K., Rowe, J., Staniford-Chen, S., Yip, R., Zerkle, D.: The design of grids: A graph-based intrusion detection system. Technical Report CSE-99-2, U.C. Davis Computer Science Department (1999)
29. White, G., Fisch, E., Pooch, U.: Cooperating security managers: A peer-based intrusion detection system. *IEEE Network* **10** (1994)
30. Valdes, A., Skinner, K.: Probabilistic alert correlation. In: Recent Advances in Intrusion Detection, Volume 2212 of Lecture Notes in Computer Science, Springer-Verlag (2001)
31. Andersson, D., Fong, M., Valdes, A.: Heterogeneous sensor correlation: A case study of live traffic analysis. Presented at IEEE Information Assurance Workshop (2002)
32. Ning, P., Cui, Y., Reeves, D.S.: Analyzing intensive intrusion alerts via correlation. In: Recent Advances in Intrusion Detection, Volume 2516 of Lecture Notes in Computer Science, Springer-Verlag (2002)
33. Wang, H., Hu, Y., Yuan, C., Zhang, Z.: Friends troubleshooting network: Towards privacy-preserving, automatic troubleshooting. In: Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS). (2004)