

UISCOPE: Accurate, Instrumentation-free, Deterministic and Visible Attack Investigation

Abstract—Existing attack investigation solutions suffer from a few limitations such as inaccuracy (because of the dependence explosion problem), requiring instrumentation, using non-deterministic approaches and providing very low visibility. Such limitations have hindered their widespread and practical deployment. In this paper, we present UISCOPE, a novel accurate, instrumentation-free and deterministic attack investigation system, which also provides high visibility. The core idea of UISCOPE is to perform causality analysis on both UI elements/events which represent users’ perspective and low level system events which provide detailed information of what happens under the hood, and then correlate system events with UI events to provide high accuracy and visibility. Long running processes are partitioned to individual UI transitions, to which low level system events are attributed, making the results accurate. The produced graphs contain (causally related) UI elements with which users are very familiar, making them easily accessible. We deployed UISCOPE on 7 machines for a week, and also utilized UISCOPE to conduct an investigation of 6 real-world attacks. Our evaluation shows that compared to existing works, UISCOPE introduces negligible extract runtime overhead (0.2%) and extra space overhead (3.05 MB event logs per hour on average) while UISCOPE can precisely identify attack provenance while offering users thorough visibility into the attack context.

I. INTRODUCTION

When security alerts are raised, a swift attack investigation should be conducted to determine the cause and scope of the attack so that the impact of the attack can be minimized and such attacks in the future can be prevented. Provenance tracking and causality analysis is an important technique for efficient attack investigation. Starting from a compromised system entity (e.g., file, socket, or process), investigators perform causality analysis on the collected provenance to figure out: 1) the root or origin of the entity, i.e., all the external entities (e.g., a socket connection) affecting the target entity, and 2) the causal path from the root to the entity. Such analysis facilitates identifying attack root cause, assessing damage incurred, and developing countermeasures. System event (e.g., system calls) auditing is a built-in feature in mainstream operating systems and can be used for such investigation. Existing work [29], [16], [13], [18], [32], [40], [30], [25] has demonstrated their great potential, but they suffer from a few major limitations.

1) Inaccurate analysis results. In many causality analysis, when a long-running process interacts with many input and output objects, each output object will be conservatively considered causally dependent on all the preceding input objects. This is known as the *dependency explosion* problem. Such problems lead to significantly inaccurate analysis results when there are long-running processes involved and makes the investigation impossible to move forward.

2) Requiring instrumentation on end-user systems. Some approaches try to solve the dependency explosion prob-

lem using instrumentation on source code or binary. However, instrumentation is generally not practical and prohibited in real world production environments. Firstly, most COTS software only provides executable binaries and do not provide source code. Secondly, intrusive modification of binary code can make applications and operating system unstable. Even more, it can introduce new vulnerabilities which can be leveraged by malware [8], [6], [2]. As such, Microsoft integrated the Kernel Patch Protection (KPP) into Windows to prevent patching the kernel [5]. Thirdly, the party which instruments COTS executables or operating systems has to take full responsibility for all potential accidents, no matter whether the accidents are caused by instrumentation or not. Hence instrumentation is mostly prohibited in the enterprise environment.

3) Non-deterministic causality inference results. Another line of work [14], [48] relies on temporal proximity to establish causal relations or uses statistical analysis [17], [29] (e.g., only considering causal edges that rarely happened in the past) to guide causal dependency graph pruning to address the inaccuracy in causality analysis. However, these solutions are usually non-deterministic. Namely, they have to depend on the observed data distribution to calculate threshold values or design rules. Thus the quality of the results heavily depend on the prior observed data distribution.

4) Lack of visibility. Attack investigation inevitably involves human efforts. Thus generating human-perceivable results is a key factor to help security experts understand attacks, perform future investigations and also determine needed actions. Existing work just provides system level information (e.g., process id, file name) and can not fully recover what happened from the user’s perspective, which plays a vital role in the forensic analysis [41].

Our solution. We propose UISCOPE, a novel instrumentation-free, accurate and deterministic attack investigation system which provides meaningful contextual information to enhance the visibility of forensics analysis. The basic idea of UISCOPE is to combine low level causality analysis with high level UI elements and events analysis to gain the advantages of both. On one hand, we leverage detailed low level system events to fully recover what actually happens in the system. On the other hand, we attribute the system events to high level UI elements and UI events to provide better visibility for attack forensics and solve the dependence explosion problem (i.e., attributing system events to individual UI elements instead of a single long-running process to avoid dependency explosion). Such analysis is a deterministic algorithm based approach and it does not require instrumentation on the end user systems.

In UISCOPE, there are two types of event collectors: the UI elements and events collector and the system events collector. For the UI elements and events collector, we leverage the

accessibility service shipped with major operating systems, and for the system event collector, we leverage built-in audit systems such as Event Tracing for Windows (ETW). These systems are provided by OS vendors and thus usually have very low runtime and storage overhead. As such, we avoid instrumenting end user systems.

After collecting all UI events and system events, UISCOPE performs causality analysis on UI elements and events (through the UI event analyzer) as well as system events (the system event analyzer) to generate causal graphs for both types of logs. We devise an additional correlation analyzer which will correlate UI events and system events. This process is a deterministic algorithm based on timestamp alignment and resource attribution (e.g., attributing background file accesses to the UI operation where the file resource was initialized). The details of the algorithm including how UISCOPE deals with software background activities are presented in Section III-F2. Through this deterministic process, we partition a long-running process into individual UI transitions, making the result more accurate. Also, the final graph represents information in a way closely coupled with UI interactions which users are very familiar with, allowing to reconstruct the attack story from user’s perspective and provide high visibility.

We deployed UISCOPE on 7 machines for a week, and utilized UISCOPE to conduct investigation of 6 real-world attacks. Our evaluation shows that UISCOPE can not only accurately identify the attack path but also provide fine-grained human-comprehensible contextual semantics to users. In addition, compared to existing works, UISCOPE introduces negligible extract runtime overhead (0.2%) and extra space overhead (3.05 MB event logs per hour on average). UISCOPE does not require any end system change or instrumentation, and can be deployed as an add-on to any existing threat detection system in production environments.

In summary, we make the following contributions:

- We identify a few limitations of existing attack forensics techniques in practice and propose a novel accurate, instrumentation-free, deterministic attack investigation system, UISCOPE, which also provides high visibility in attack forensics. It leverages both UI information in the user space and system events in the kernel space to achieve our aforementioned design goals.
- We devise a novel UI event analyzer which analyzes UI events causality, and a correlation analyzer which correlates UI events with system events to produce accurate and highly visible causal graphs.
- Based on our design, we build a prototype on Windows and evaluate on 7 machines and 6 real world attacks. Results show that compared to existing works, UISCOPE introduces negligible extract runtime overhead (0.2%) and extra space overhead (3.05 MB event logs per hour on average) while provides more accurate and highly visible attack forensics.

II. MOTIVATION

A. Motivating Example

One day, an office clerk Bob tried to download a piece of software (`WinSCP.exe`) from a website `benign.com`. While waiting for the download to complete, he received an email informing that he had been selected as a winner for

an iPad, and asking him to claim the award on a website. It was actually a phishing email leading to a malicious website. The website uses the WordPress free host service and uses the domain name `well-known.wordpress.com`. It also leverages login detection techniques [7] to test if the user has logged into any well-known bank website. In our attack story, Bob happened to login `www.bank.com` to manage his company’s bank account. Hence, the malicious website launched a clickjacking attack [4] by creating a transparent frame containing a transfer form with the receiver being the attacker’s own bank account on top of the “Get iPad” button. Bob clicked the “Get iPad” button, which actually triggered a transfer from his bank account to the attack’s bank account, without being noticed by Bob.

The good news was that Bob’s company deployed protection techniques. The security system raised an abnormal transfer alert. Following this alert, forensics investigation began from the abnormal transfer socket communication, and tried to find how/when/where this happened and investigate if Bob was an insider attacker who stole money from the company.

B. Existing Attack Investigation Solutions

1) **Low Level Events Causality Analysis:** Traditional causality analysis [22], [23] tracks the lineage of system objects (e.g., files and sockets) and subjects (e.g., processes) via system events (e.g., syscalls) and analyzes the causal relations among system objects and subjects to generate causal graphs. With such graphs, investigators can perform *backward* provenance queries with the symptom (i.e., the abnormal transfer socket) to understand the root cause of the attack or *forward* provenance queries to identify the effects of the attack.

Fig. 1 (I) shows a typical dependence graph generated by traditional causality analysis for this investigation. Note that in this figure (and also the rest of the paper), we use *diamond nodes* to denote *sockets*, *box nodes* to denote *processes*, and *oval nodes* to denote *files*. Also, nodes in red represent the symptom event(s), the events which investigation starts from. Specifically, when investigators perform backward tracking to find which website is related to the abnormal transfer, they would be substantially distracted, because Chrome is a long running process and could have been used to visit hundreds of websites during an attack window (e.g., `mail.google.com`, `benign.com`, `drive.google.com` and many other sites in the figure), and thus the suspicious file is conservatively related to hundreds of preceding websites. This is known as the *dependence explosion problem*, which makes any further tracking attempt nearly impossible. This is a major limitation of traditional causality analysis.

Many previous works tried to address this problem via program analysis [25], [31], [32], [24]. Some of them require source code or binary level instrumentation [25], [31], [32], which is intrusive and not practical in enterprise environments, as discussed in Section I. Taint analysis [20], [39], [21] is another way to solve the dependency explosion problem. However, it causes tremendous runtime and space overhead, which makes it rarely used in production environments.

2) **Statistics-based Graph Pruning:** Observed that most attack-related events are abnormal and rarely occur in historical event logs of an enterprise network, PrioTracker [29] and NoDoze [17] proposed statistics-based attack investigation

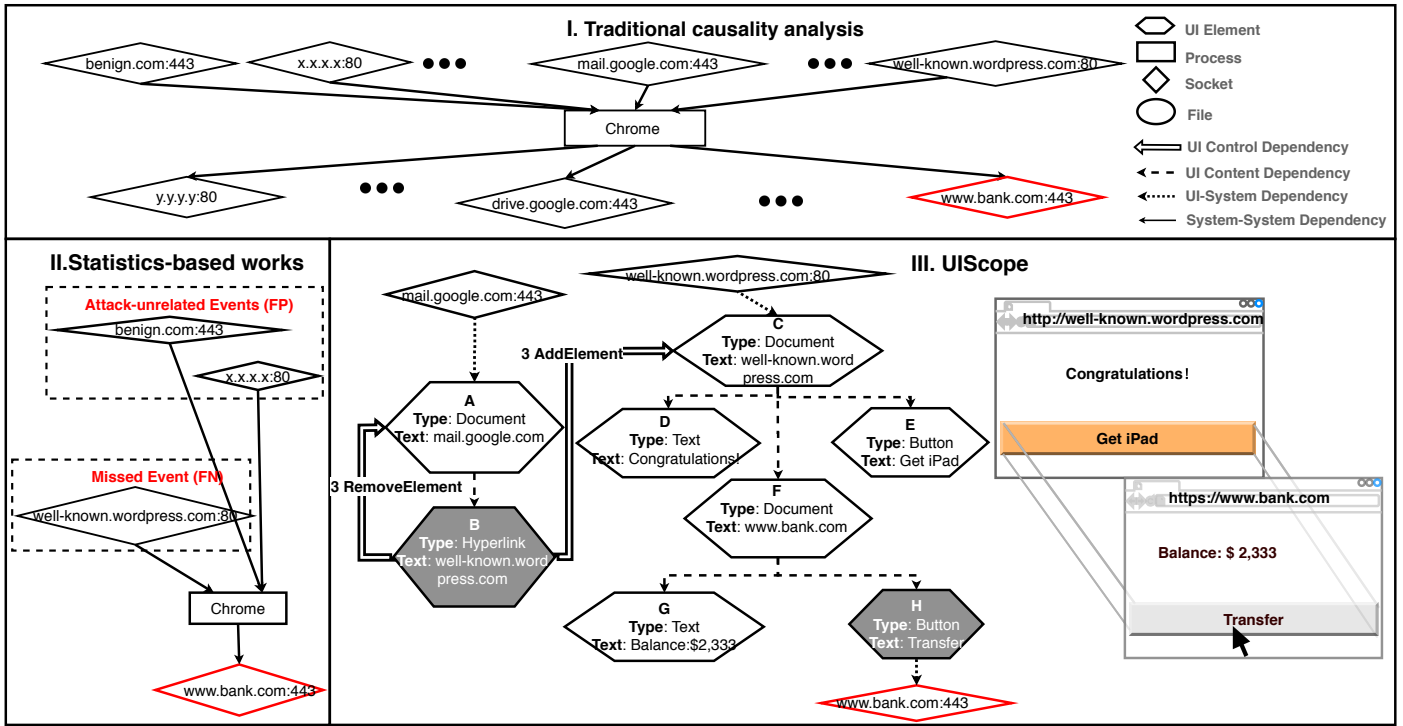


Fig. 1. Comparison between existing causality analysis solutions and our UISCOPE on the motivating example.

approaches to prioritize abnormal events and causal dependencies. They introduce quantitative metrics (i.e., frequency and topological features) to distinguish normal and abnormal events, and present a pruned causal graph to investigators. Pri-oTracker only considers the abnormality of individual events while NoDoze takes the abnormality of event chains into consideration and thus can generate more precise causal graphs.

Fig. 1 (II) presents the causal graph generated by NoDoze for the investigation, which is more concise than the graph generated by traditional techniques (Fig. 1 (I)). However, NoDoze cannot accurately locate the IP address which is the real source of the abnormal transfer. As shown, NoDoze traces back to two benign sockets, `benign.com` and `x.x.x.x`. The reason is that visiting `well-known.wordpress.com` which is hosted by WordPress is a common and normal behavior (because WordPress is one of the largest blog host websites) while `benign.com:80` and `x.x.x.x:80` are rarely visited websites and thus deemed abnormal instead. As such, statistics based approach may produce unstable results. In other words, the quality of the results heavily depends on the dataset used to calculate the observed distribution (e.g., which website gets more visits).

C. Problem Statement and UISCOPE

As discussed in the previous section, many low level event causality analysis based approaches are not applicable in real world systems as they may suffer from the dependence explosion problem, causes heavy runtime overhead or requires instrumentation. Statistics based graph pruning has difficulty handling attacks leveraging popular (and benign) applications and/or websites. Furthermore, graphs generated by both approaches lack the user understandable high level semantic information. Fig. 1 (II) shows a simplified causal graph generated by NoDoze (Note that using traditional low

level events causality analysis will introduce too many irrelevant sockets and we preclude it from our discussion). With such graphs, even if NoDoze may find the real malicious socket `well-known.wordpress.com:80` by tuning thresholds or using high quality data to calculate the observed distribution, it still cannot provide sufficient information to answer the key question in the investigation: is Bob an insider (i.e., did he intentionally initialize the transfer) or an victim of social engineering (e.g., fooled by clickjacking)? Without knowing more contextual information, it is impossible to answer this type of questions.

The aforementioned limitations motivate the following design goals for our own system:

- **Accurate:** It should produce accurate investigation results, i.e., low false positive and false negative rates.
- **Applicable:** It requires **NO** extra end system change and causes low overhead so that it is deployable in an enterprise environment.
- **Stable:** It should produce **STABLE** investigation results, i.e., results do not depend on prior observed data.
- **Visible:** Its output should be visible to human investigators so that they can easily understand what happened with the application level context information. *Visibility* plays a vital role in the forensic analysis, as pointed out by [41].

As far as we know, none of the existing approaches can achieve all of the four goals. Therefore, we propose a new investigation system, UISCOPE, which combines low level event causality analysis with high level user interface (UI) elements/events analysis to achieve all the aforementioned goals. While the detailed design will be discussed in Section III, Fig. 1 (III) shows the UISCOPE graph for the aforementioned investigation. In this graph, we use the same shape and color to represent low level system events as before. We also introduce

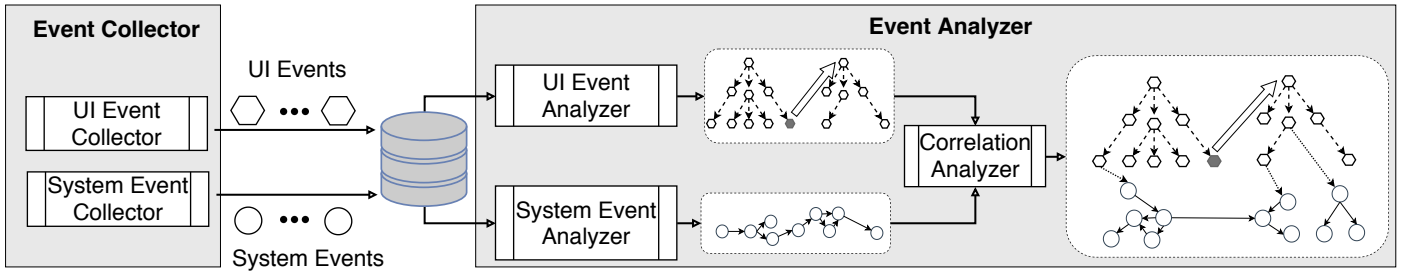


Fig. 2. System Architecture.

a new set of nodes and edges to represent UI elements and new relations, which is defined in Section III. Specifically, hexagons represent UI elements and nodes with gray background denote UI elements operated by the user. Dashed lines show the tree structure of different UI elements pointing from the parent node to the child node, which denotes causality as well. A dotted line represents connection between a UI element and a system event and the arrow is from the action initiator. A double solid line represents connection between two UI elements and the destination is affected by the source via some user operation. For example, in Fig. 1 (III), hexagon H represents the *Transfer* button that was actually clicked by Bob; hexagon E is the *Get iPad* button which Bob intended to click, and hexagons C, D, E, F, G are Document Object Model (DOM) elements in the same page with the button. Bob first clicked the link in his email (node B), which led him to a new web page (whose UI elements rooted in C). After Bob clicking the transparent button over the *Get iPad* button (node H), a socket was created and was used to perform the unintended money transfer. The UI element nodes in this graph tell us that in the same web page, there were two clickable buttons in an identical position. If we reconstruct what was seen by Bob, we can get a graph like Fig. 1 (III), indicating that Bob was fooled by a clickjacking attack.

Additionally, Fig. 1 (III) attributes low level system objects (e.g., detailed socket addresses and file names) to the corresponding high level UI elements (i.e., web pages in this case), that are well partitioned and denote autonomous sub-executions. Doing so, we can accurately associate low level system information with individual autonomous and high-level user actions. This avoids attributing all low level events to the same process, achieving low false positive/negative rates in causality analysis, as shown by our results in Section IV-B. Details of how the attribution is done including how background events are processed will be discussed in Section III-F.

D. Threat Model

We assume that both Event Tracing for Windows (ETW) and Accessibility libraries provided by OSes (introduced in Section III) are trusted, and audit logs (i.e., system events captured with ETW and UI events by Accessibility libraries) cannot be tampered with. Attacks that can compromise these two auditing systems are beyond the scope of this study. This assumption is consistent with previous literature [17], [29], [18]. In this paper, we focus on GUI-related attacks, in which user involvement is needed to initiate/trigger an attack (e.g., phishing attacks, driven-by downloads and insider attacks). For attacks which do not entail any user interactions, we apply the same methods as previous works [17], [29], [18].

III. SYSTEM DESIGN

A. System Design Overview

The overall workflow of UISCOPE is shown in Fig. 2. UISCOPE has two major components: the event collector and the event analyzer. Specifically, the event collector consists of the UI event collector and the system event collector, and the event analyzer contains the UI event analyzer, the system event analyzer and the correlation analyzer.

The overarching idea of UISCOPE is to perform both UI event causality analysis and low level system event causality analysis independently, and then attribute groups of low level system events to high level UI event. We do not use statistical approaches as their results have substantial non-determinism. Instead, all our analyzers use deterministic algorithms. In UISCOPE, we use UI events to deliver visibility to investigators and solve the dependency explosion problem in low level causality analysis by attributing low level events to high level UI elements and events, which are well partitioned. Doing so, UISCOPE can generate accurate (i.e., no dependence explosion problem) and visible (i.e., through human understandable UI elements and events) results. Moreover, we do not require instrumentation or heavyweight runtime monitoring so that our technique is applicable in production systems.

B. Event Collector

1) *UI Element and Event Collector*: We develop a UI event collector based on *Windows UI Automation* [34], to efficiently monitor interesting UI events triggered by user interactions. *Windows UI Automation* is an accessibility library developed by Microsoft. There are similar accessibility libraries on other platforms as well including *NSAccessibility* for Mac OS X [11], *ATK* and *XAutomation* for Linux [28] and so on. Enforced by the federal *IT Accessibility Laws and Policies* [15], mainstream OSes have to develop these libraries to support accessibility features. With these features, disabled people can have more control of the user interface of electronic devices, such as 1) retrieving information about UI elements (e.g., title of a window, name of a button), 2) adjusting UI elements (e.g., customizing a screen's color, zooming in and out on a web page), and 3) getting notification of changes to UI (e.g., website content loaded, a hyperlink clicked). Essentially, *Windows UI Automation* provides UISCOPE the capability of logging UI events triggered by any changes to UI elements.

UI Automation supports subscribing to around 100 different UI events, falling into 5 categories, including property change, element action, structure change, global desktop change and notification [35]. Among all these

events, three type of events, namely `Focus`, `Invoke`, and `SelectedItem_ElementSelected`, are used to capture common user interactions. Specifically,

- `Focus` events indicate that the focus of the user shifts from one element to another. For example, either clicking a hyperlink or yielding a new web page would trigger a `Focus` event. And at a specific time in the system, only one element can gain focus, and only that element can receive user inputs (e.g., via keyboard or mouse).
- `Invoke` events indicate when a UI element is triggered, such as clicking a button.
- `SelectedItem_ElementSelected` events indicate that an item (i.e., a group of elements) or an element is selected, such as an email in the Outlook or a browser tab is selected.

Besides, each UI event has a uniform format with 170 fields [36], out of which, eight fields provide wealthy semantics helpful to understand user interactions:

- `ProcessId` represents the process identifier (ID) of a UI element.
- `EventType` specifies the UI event type (e.g., `Focus`).
- `ControlType` specifies the type of the UI element which triggers the event, such as a hyperlink, a document or a tab.
- `ClassName` is the class name of the UI element assigned by software developers. Most of the time, `ClassName` together with `ControlType` can determine what kind of element triggers an event.
- `RuntimeId` is a unique identifier (ID) of a UI element, used to identify the UI element which triggers an event.
- `BoundingBox` is the point coordinates of a UI element's enclosing rectangle. It represents the element's position on the screen.
- `Name` is the name of UI element and is usually the same as the human-perceivable text on the screen. For example, the hypertext of a hyperlink or the title of a web page is often recorded in the `Name` field of an event.
- `Text` contains the hidden value of a UI element, such as URL of a web page or a hyperlink.

Example. Fig. 3 (A) presents an example of our collected UI information of the Outlook email composition page, in which important UI elements are marked in different colored boxes. Fig. 3 (B) shows the corresponding UI tree, and each node in the tree denotes a UI element. Our UI element collector will collect such information at runtime, and the UI can be reconstructed (during analysis) with such information. In Fig. 3, we showcase a UI event of clicking the attachment button. As shown in the figure, we will collect event related information (C-1), process related information (C-2) and also UI element related information (C-3).

2) **System Event Collector:** UISCOPE uses OS built-in audit systems for system event collection. These systems are pre-installed on target systems. Without any additional configuration or optimization, they can be used in production runs. They are also of high quality and have technical support from official providers. UISCOPE leverages Event Tracing for Windows (ETW) [44] to collect system events on Windows. It allows users to collect system-level events (e.g., system calls) with negligible overhead [44]. ETW has been widely in both academia [17], [32], [13], [12], [30] and industry [45]. Similar

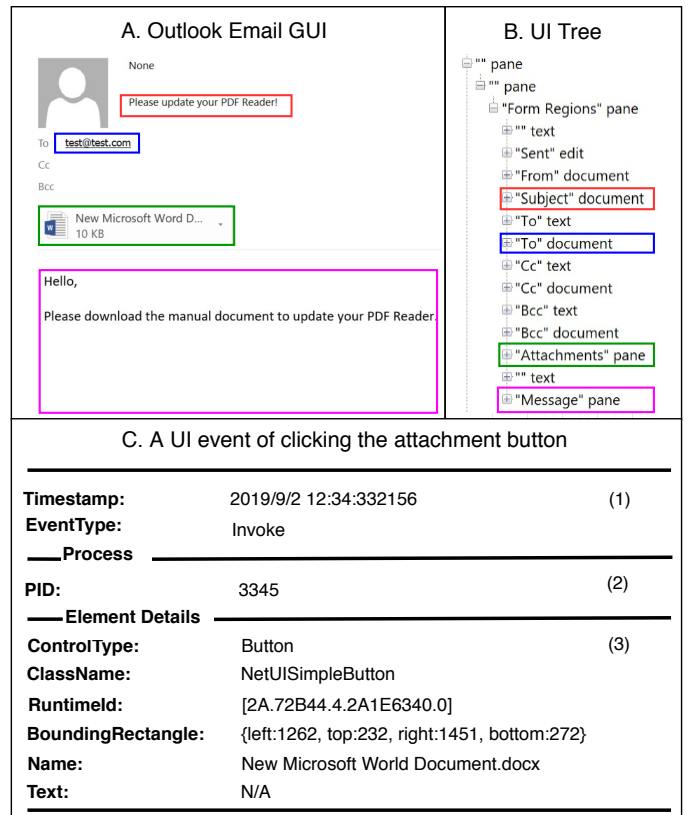


Fig. 3. Example of our collected UI information of the Outlook email composition page. (A) is the Outlook GUI which was seen by users, and (B) is the corresponding UI tree. (C) is a UI event raised by clicking the email attachment button highlighted in green in (A) and (B).

to previous attack investigation tools [17], [29], UISCOPE only monitors security-relevant events.

| | | |
|----------------------|-----------------------|-----|
| Timestamp: | 2019/9/2 12:32:538364 | (1) |
| EventType: | ProcessCreate | |
| Process | | |
| Parent PID: | 102 | (2) |
| PID: | 25336 | |
| TID: | 322 | |
| Event Details | | |
| Command: | MicrosoftEdge.exe | (3) |
| Env: | OS=Windows_NT | |
| Return Values | | |
| PID: | 32248 | (4) |

Fig. 4. Example of an ETW event.

Example. Fig. 4 shows an example ETW event entry. The first block (1) includes the basic information of the event such as the timestamp of the event (`Timestamp`), the event type (`EventType`) and so on. The second block (2) shows the process and thread triggering the event and its parent process. The third block (3) lists details of the event. For instance, for a `ProcessCreate` event, we will have the command of the created process and the environment where the created process was executed. The last block (4) has return values. In this case, it returns the child process PID value.

C. Event Analyzer Definitions

The event analyzers analyze the collected UI and system logs to derive dependencies between event entries (to construct

the final causal graph). In this section, we first define a number of dependencies and then explain how they are derived.

UI Control Dependency. When a UI element change directly affects other UI elements, we say they have UI control dependencies. A UI element can affect another one directly by adding and removing operations. Thus there are two types of UI Control Dependencies: AddElement (for adding a UI element) and RemoveElement (for removing a UI element). For example, after the user clicks a hyperlink on a web page and a new page is loaded to replace the current one. We will create a RemoveElement edge from the hyperlink element to the root element of the current page and a AddElement edge from the hyperlink element to the root element of the new page. It implies that all elements in this tree will be (directly or transitively) affected. In our paper, we use double solid lines to represent this type of dependency.

UI Content Dependency. UI elements are organized in a tree structure. A parent tree node usually represent a larger element such as a web page, and all the child nodes represent sub-elements in this web page. The affiliation relations of UI elements introduce *content dependencies*. In our paper, we use dashed lines to represent this type of dependency.

UI-System Dependency. A UI-system dependency is introduced if the UI element lead to the system event (e.g., creating a new socket for page loading in browsers) or the UI element is affected by the system event (e.g., loading data from a socket to refresh a web page). Such dependencies are critical for attributing low level system events to execution structures enforced by UI elements (to avoid dependence explosion). In our paper, we use dotted lines to represent this type of dependency.

System-System Dependency. Previous literature [22] clearly defined dependencies between system objects (i.e., Files, Sockets) and subjects (i.e., Processes). We adapt the same definition and use the term System-System Dependency to represent this category of dependency in our paper. In our paper, we use solid lines to represent this type of dependency.

In addition, we use the term *initial event* to denote system object creation events.

D. UI Event Analyzer

The UI analyzer is to find UI control dependencies and content dependencies. The detailed analysis procedure is presented in Algorithm 1. For a list of given UI events in the chronological order and also all the related UI elements organized in trees. Our goal is to compute a provenance graph, with UI elements as nodes, connected by UI control dependency and content dependency. After initializing the variables (lines 2 to 6), the algorithm starts to examine individual UI events in the queue (lines 7 to 20). For each event, after getting the related UI tree (line 10), it finds all the associated UI elements in the newly added or removed tree and update the graph accordingly (lines 11 to 18). Specifically, it introduces AddElement and RemoveElement types of control dependency for added and removed UI elements, respectively. Note that the source of those dependencies are the element operated by the previous UI event U_{prev} . Lastly, it updates U_{prev} and T_{prev} to the current UI event and the current UI tree respectively (lines 19 to 20). For UI events like *focus*, as there is no new

Algorithm 1 UI Event Analyzer

Input: L_U - List of UI events in the chronological order
Output: *Graph* - Provenance graph whose vertexes are UI elements and edges UI Control or Content Dependency.
Functions: *GetUITree*(U) - Returns the entire UI tree when a UI event U happens
FindAddedUITrees(T_{cur}, T_{prev}), *FindRemovedUITrees*(T_{cur}, T_{prev}) - Returns a set of added or removed sub-trees by comparing two UI trees
Graph.vertexes.add(E) - Add a UI element to the graph if it is included
Graph.edges.add($E_{src}, E_{dst}, DependencyType$) - Add a edge from E_{src} to E_{dst} if the edge does not exist and the type of edge is set by *DependencyType*
GetChildren(E) - Returns a set of child nodes of UI element E
Variable: $T.root$ - the root element of the UI Tree T
 $U.element$ - The UI element operated by a UI event U
 $U.timestamp$ - The timestamp of a UI event U

```

1: function UIEVENTANALYZER( $L_U$ )
2:    $U_{cur} \leftarrow null$ 
3:    $T_{cur} \leftarrow null$ 
4:    $U_{prev} \leftarrow L_U[0]$ 
5:    $T_{prev} \leftarrow GetUITree(U_{prev})$ 
6:   AddTreeToGraph( $T_{prev}, Graph$ )
7:   for  $i = 1; i < Size(L_U); i++$  do
8:      $U = L_U[i]$ 
9:      $U_{cur} \leftarrow U$ 
10:     $T_{cur} \leftarrow GetUITree(U)$ 
11:    /* Build UI Control Dependency */
12:    for each  $T_{add} \in FindAddedUITrees(T_{cur}, T_{prev})$  do
13:      AddTreeToGraph( $T_{add}, Graph$ )
14:       $Graph.edges.add(U_{prev}.element, T_{add}.root, AddElement)$ 
15:      Set timestamp of above new added edge as  $U_{prev}.timestamp$ 
16:    for each  $T_{remove} \in FindRemovedUITrees(T_{cur}, T_{prev})$  do
17:       $Graph.edges.add(U_{prev}.element, T_{remove}.root, RemoveElement)$ 
18:      Set timestamp of above new added edge as  $U_{prev}.timestamp$ 
19:     $U_{prev} \leftarrow U_{cur}$ 
20:     $T_{prev} \leftarrow T_{cur}$ 
21:
22: function ADDTREETOGRAPH( $Tree, Graph$ )
23:   for each  $E_{parent} \in Tree$  do
24:      $Graph.vertexes.add(E_{parent})$ 
25:     for each  $E_{child} \in GetChildren(E_{parent})$  do
26:       /* Build UI Content Dependency */
27:        $Graph.edges.add(E_{parent}, E_{child}, UIContentDependency)$ 

```

or removed elements, it simply updates the T_{prev} and U_{prev} without adding new edges to the graph. We use the function *AddTreeToGraph*(\cdot) to add a new tree to the existing graph (lines 22 to 27). In this function, the algorithm traverses the whole UI tree associated with the existing UI element through UI content dependency and add them all to the graph.

Example. Fig. 5 presents an example of how Algorithm 1 works on the motivating example where Bob clicked a hyperlink in the phishing email and was attacked by clickjacking. Fig. 5 (I) and (II) are UI events in temporal order and their corresponding UI trees. UI elements with gray background represent the ones operated by users. Fig. 5 (III) is the generated graph by the event analyzer. Before timestamp 3, the UI element of the `mail.google.com` has been added to the graph. At timestamp 3, Bob clicked the hyperlink (element B), then the previous web page `mail.google.com` (Elements A and B) is removed from the tree and a new web page `well-known.wordpress.com` (Elements C, D, E, F, G, H) is loaded. The new page is thus added to the graph, together with a *RemoveElement* edge from Element B to Element A and an *AddElement* edge from element B to element C. Note the timestamp of the edges are 3. At timestamp 6, Bob was tricked to click the *Transfer* button (element H) masqueraded as a *Get iPad* button (element E). Once the transfer finished, the embedded web page `www.bank.com` (elements F,G,H) was removed, thus a *RemoveElement*

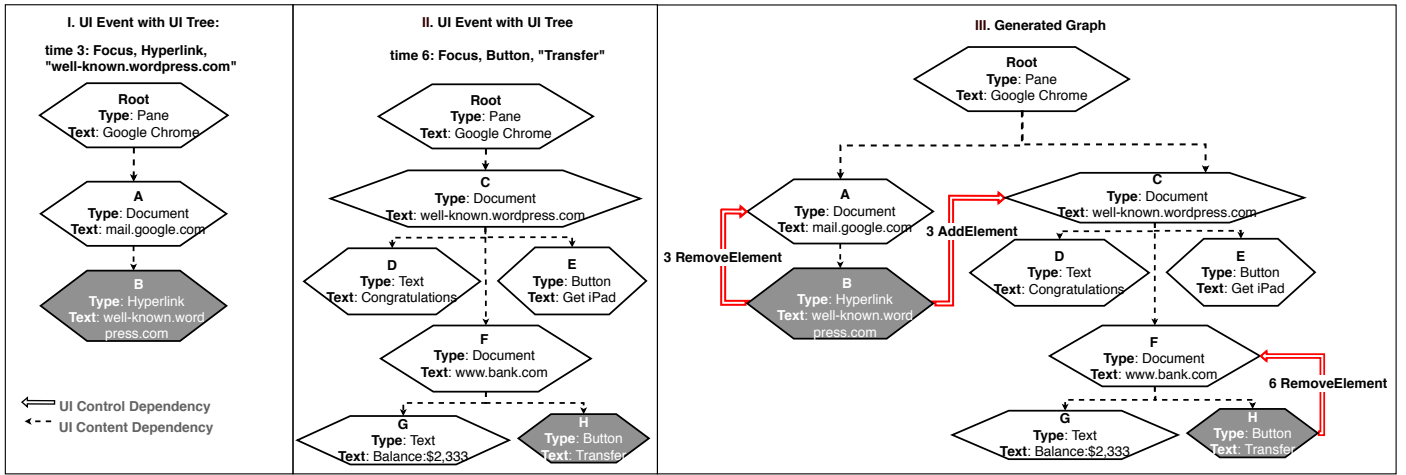


Fig. 5. A log example of Event Analyzer.

pointing from element H to the root element of the removed page (element F) is added with timestamp 6.

E. System Event Analyzer

We follow the standard causality analysis method used in previous literature [25], [31], [32], which returns a causal graph with system subjects or objects as nodes and System-System Dependencies as edges.

F. Correlation Analyzer

1) *UI-System Dependency*: UI events and system events capture behaviors of the same attack from two different levels: *foreground* with visibility and *background* with fine-grained information. After acquiring the UI causal graph from the UI event analyzer and the low level event causal graph from the system event analyzer, we analyze the correlations of the two so that low level system events can be attributed to individual UI elements (instead of the whole process) which makes the graph both visible and accurate (more accurate than using a single node to represent the whole graph like in traditional graphs, which leads to dependency explosion).

We devise a timestamp-based attribution approach based on two observations: 1) an application has only one currently focused UI tree at a certain time; 2) system events and UI events are typically triggered by the same attack behavior at the same time. That is, the two categories of events are mostly time-aligned. There are also many background events and we will discuss them in III-F2.

System events and UI events are correlated based on two mechanisms. Firstly, for a system object that is associated with a sequence of system events such as a socket with many socket read/write events, we use the creation of system object to correlate with UI events. This is known as the **initial event based correlation**. Secondly, we mainly use a timestamp based alignment to correlate general system events and UI elements/events, and this is referred as the **timestamp based correlation**. The details of the attribution process is presented in Algorithm 2. After initializing variables (line 2), the algorithm generates two graphs using the UI Event Analyzer and the System Event Analyzer, respectively, and

add them to the result graph (lines 3 to 6). Then it starts to examine individual processes occurring in the graph (lines 7 to 15). For each process, it first checks whether the process has GUI. If so, it extracts two sub-lists of events related to the process and attributes system objects or subjects to UI elements by invoking the function *Attribute* (lines 10 to 12).

The function *Attribute()* first orders all system and UI events chronologically in list L_U (line 18). For each event in the list, it first checks the event type. If it is a UI event, it checks whether the user's focus has been changed and mark the new UI tree as the new focus if so (lines 24 to 29). If the event is a system event, it first tests whether the event is an initial event (line 32). If so, it records the relation between the initialized resource and the current active element in the map M_{init} (line 34). Then it adds a new edge pointing from the active element to the sink (line 35). Otherwise, it retrieves the UI element related to the operated resource from previous records in M_{init} and conducts the same operations as before (lines 38 to 42). Note that the direction of the new created edge is the same as the one of the system event. The essence of lines 38 to 42 is as follows. We observe that the usage of a system object (e.g., socket read) may not be in the same time window of a UI event, rendering the timestamp based alignment (lines 33-35) ineffective, while its initialization (e.g., socket creation) can be correctly aligned with the trigger of the UI event (e.g., button clicking) in most cases. Thus, for a (background) system event related to some system object, we trace back to the time when the system object was initialized and attribute the event to the active UI element at that moment.

2) *Background Activities*: One may question if our technique (based on timestamps and system object initialization) provides a sound solution for background activities. In the following, we discuss the possible background behaviors of popular applications and argue that our technique is highly effective in practice, which is also demonstrated by our experiment in Section IV-B.

We classify applications into two categories: *static* and *dynamic* based on if there are activities triggered in the background. The behaviors of static applications are only triggered by user interactions. For example, there are no background activities in Notepad PlusPlus when there are




Fig. 6. The number of active time units is small for most websites.

Algorithm 2 Correlation Analyzer

Input: L_S - A list of system events in the chronological order
 L_U - A list of UI events in the chronological order
Output: $Graph$ - A complete provenance graph combining both UI and system events.
Variable: E_{active} - An active UI element
 M_{init} - A key-value map of initialized objects and UI elements
 $U_{element}$ - The UI element operated by a UI event U
 S_{sink}, S_{source} - The system object pointed to or from by a system event S
 S_{type} - The event type of a system event S
Function: $GetProcessVertexes(Graph)$ - Return all vertexes of processes
 $AddGraphToGraph(G_{from}, G_{to})$ - add the vertexes and edges in G_{from} to G_{to} if those vertexes and edges do not exist in the G_{to}
 $ReorderInChronologicalOrder(List_{one}, List_{two})$ - merge two lists into a list in the chronological order
 $FindRootElementOfProcess(process)$ - return the root element of a process.
If a process has multiple ones (e.g., Firefox with 2 windows), return the earliest one.
 $FindAddedElement(E, Graph)$ - check whether there is $AddElement$ dependency pointing from the element E , if yes, return the sink element of the dependency, otherwise return $null$.

```

1: function CORRELATE( $L_S, L_U$ )
2:    $Graph \leftarrow$  a new empty graph
3:    $G_{system} \leftarrow$  SystemEventAnalyzer( $L_S$ )
4:    $G_{ui} \leftarrow$  UIEventAnalyzer( $L_U$ )
5:   AddGraphToGraph( $G_{system}, Graph$ )
6:   AddGraphToGraph( $G_{ui}, Graph$ )
7:   for each  $process \in GetProcessVertexes(G_{system})$  do
8:     if HasGUI( $process$ ) then
9:       /* For GUI applications*/
10:       $SubL_{system} \leftarrow$  ExtractEventsByProcess( $L_S$ )
11:       $SubL_{ui} \leftarrow$  ExtractEventsByProcess( $L_U$ )
12:      Attribute( $process, SubL_{system}, SubL_{ui}, Graph$ )
13:     else
14:       /* For non-GUI applications, we follow traditional methods*/
15:       continue
16:
17: function ATTRIBUTE( $process, L_{system}, L_{ui}, Graph$ )
18:    $L_{event} \leftarrow$  ReorderInChronologicalOrder( $L_{ui}, L_{system}$ )
19:   for  $i = 0; i < Size(L_{event}); i++$  do
20:      $event \leftarrow L_{event}[i]$ 
21:     if IsUIEvent( $event$ ) then
22:        $U \leftarrow event$ 
23:        $E_{add} \leftarrow$  FindAddedElement( $U_{element}, Graph$ )
24:       if  $E_{add} = null$  then
25:         /* Activate operated element if no new element occurs */
26:          $E_{active} \leftarrow U_{element}$ 
27:       else
28:         /* Activate the root of the new added element */
29:          $E_{active} \leftarrow E_{add}$ 
30:     else if IsSystemEvent( $event$ ) then
31:        $S \leftarrow event$ 
32:       if IsInitialEvent( $S$ ) then
33:         /* Timestamp-based approach */
34:          $M_{init}.put(S_{sink}, E_{active})$ 
35:          $Graph.edges.add(E_{active}, S_{sink}, S_{type})$ 
36:       else
37:         /* Initial event-based approach */
38:          $E_{init} \leftarrow M_{init}.get(S_{object})$ 
39:         if IsProcess( $S_{source}$ ) then
40:            $Graph.edges.add(E_{init}, S_{sink}, S_{type})$ 
41:         else
42:            $Graph.edges.add(S_{source}, E_{init}, S_{type})$ 

```

no user inputs. Dynamic applications contains executions that are triggered by background activities. For example, Chrome can download videos or other files in the background while the user is browsing in the foreground. Browser is the most common and complex dynamic application type. Hence in the following, we will use browsers as an example to demonstrate how UISCOPE handles background activities.

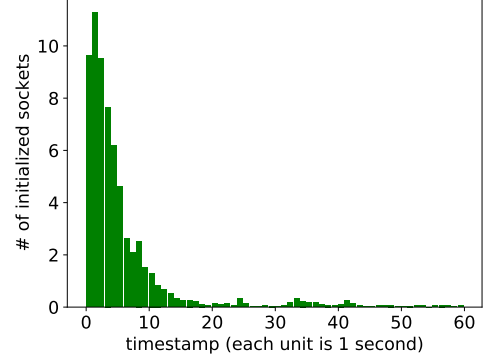


Fig. 7. The number of initialized socket sessions within the first 60 second.

The accuracy of UISCOPE for browsers is affected by 1) the types of visited websites, 2) the duration of web sessions, 3) the number of website with background activities. We use Chrome to monitor the behaviors of 500 popular websites in an hour after being loaded. Those website are selected from Alexa top websites [1]. In order to simulate a real-world environment, for each website, we manually select a common visited web page which can reflect the typical behavior of the website, e.g., playing video on `www.youtube.com`. Then, we use a GUI automation tool [3] to automatically visit them and collect trace for 1 hour. Furthermore, we only consider network events because websites do not have permissions to access the file system. Finally 436 websites successfully monitored and used in this experiment is listed in Table VI. After analyzing the data, we make a few important observations.

Observation 1: Most new websites activities are narrowed in a short period of time and activities in the background are not common. We monitor website activities in every 10 seconds as a unit and count the total number of units the website initializes at least one new socket. The result is shown in Fig. 6. The Y-axis in the graph represents the number of websites and the X-axis represent the number of active units. For example, the first bar denote that 106 websites are active in only one unit (10 seconds). From the graph, we know that 64.28% of the 436 websites finish all the work within 1 minute and 88.49% websites are idle for at least 55 minutes during the whole time (1 hour).

Observation 2: Most socket sessions are initialized during web page loading. Fig. 7 shows the average number of initialized socket sessions within 60 seconds after a website is visited. We observe that around 87.13% of all socket sessions are initialized within 10 seconds and 94.11% socket sessions are initialized within 20 second. Together with the analysis results in observation 1, it verifies the usefulness of our timestamp-based event attribution approach (Section III-F).

Observation 3: Background activities are usually repeated behaviors. We analyzed the socket connection destinations of all sockets, including both existing sockets and new sockets,




Fig. 8. The number of initialized socket for each time slice starting from 1 minute to 1 hour.

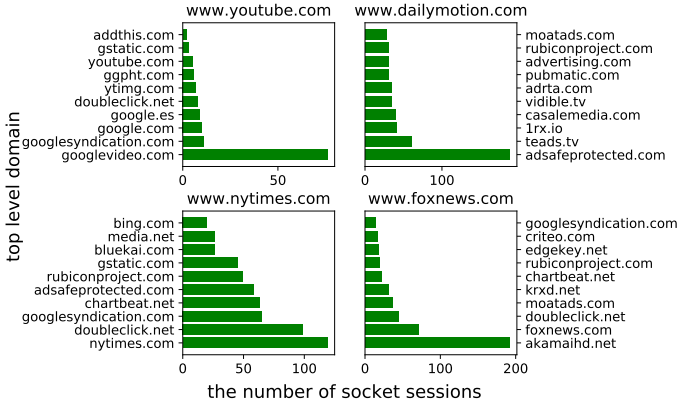


Fig. 9. Long-term active website examples.

and the results are shown in Fig. 8. Green bars are the number of both types of sockets and red bars are the number of newly created sockets. The X-axis is time, and one hour period is divided to 360 units. As we can see, there are not many new sockets after the pages are loaded, which implies that websites in the background connect to a limited set of domain names repeatedly. To further verify our finding, we manually check with source code of some of these sites. Fig. 9 shows 4 websites that have the maximum number of background activities. The Y-axis represents the top-level domain name and X-axis shows the number of socket sessions that are related to a certain domain name. As we can see, they all have a limited number of top-level domain names as the socket connection destination indicating that background behaviors are largely repeated connections. This allows UISCOPE to use pattern based filters to remove such background activities.

In an adaptive attack scenario in which the attacker is aware of the presence of UISCOPE, the attacker may intentionally use delayed background activities that do not have the resources initialized when the UI event occurs (e.g., using a timer to postpone the creation of a socket to download a payload). Such attacks may lead to incorrect attribution of low level events to UI operations. We argue that all the existing forensics techniques similar to ours (e.g., [22], [23], [29]) have difficulty dealing with adaptive attacks. From our earlier analysis, UISCOPE is highly effective in practical attack scenarios.

Example. Fig. 10 presents an example of how Algorithm 2 works on the motivating example. Fig. 10 (I) and (II) show the graphs generated by the UI Event Analyzer and the System Event Analyzer. Fig. 10 (III) shows the final output graph. We use logic timer in this case to denote the order of all the events. At the timestamp 1, Explorer creates a new process Chrome. Thus there is an edge (from Explorer to the root node of Chrome) to represent this dependency. Elements

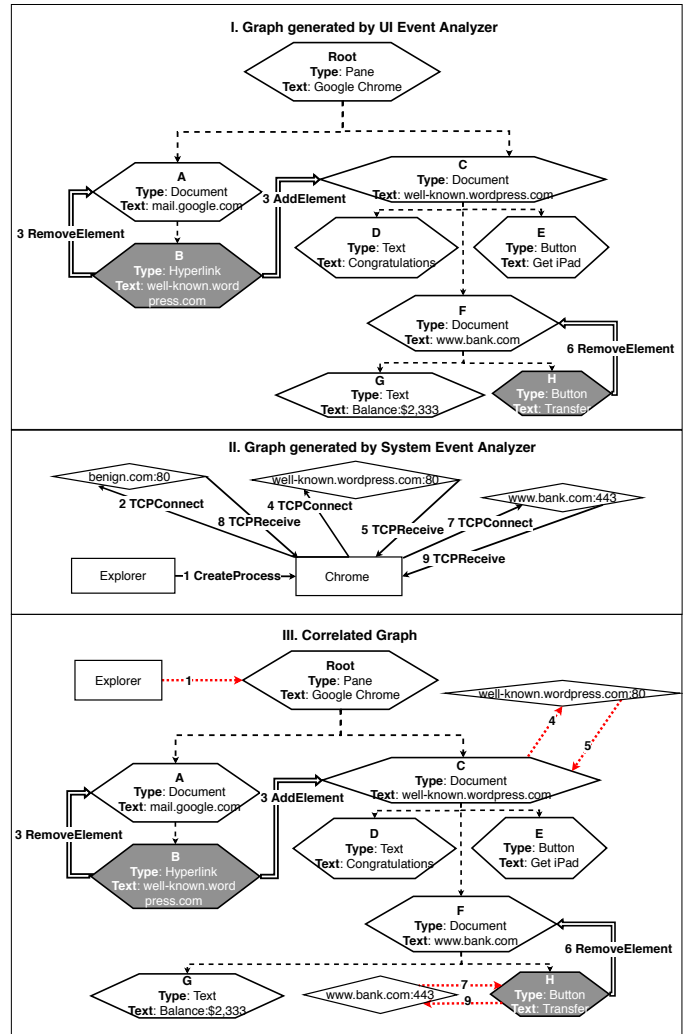


Fig. 10. A log example of correlation analyzer.

B and H were clicked at timestamps 3 and 6 respectively, which lead to two sockets being initialized at timestamps 4 and 7. According to the algorithm, the socket events are attributed to elements B and H, respectively. System events occurring at timestamps 5 and 9 are non-initial events and we attribute them to elements B and H through the tracing-back to initial system object method. Note that `benign.com:80` was initialized before those two UI events for downloading the software `WinSCP.exe`, thus it was attributed to previous UI elements. Furthermore, the non-initial system event related to `benign.com:80` occurring at time 7 was attributed to the previous element to which the socket creation belongs to. *Observe that by separating a process to different autonomous*

TABLE I. SPACE OVERHEAD EVALUATION RESULTS

| | S-0 | S-1 | S-2 | S-3 | S-4 | S-5 | S-6 | Total |
|-----------------------------|---------|----------|---------|---------|----------|----------|---------|---------|
| CPU | I5-9400 | I7-6500U | I7-8700 | I7-8700 | I7-7500U | I7-8550U | I5-7400 | N/A |
| RAM (GB) | 8 | 8 | 16 | 16 | 16 | 16 | 8 | N/A |
| Windows Version | 7 | 10 | 10 | 10 | 10 | 10 | 10 | N/A |
| Duration (hours) | 16.7 | 108.7 | 95.4 | 54.6 | 37.4 | 28.7 | 121.9 | 463.4 |
| # of UI Events/Trees | 1,131 | 32,824 | 18,687 | 34,010 | 22,766 | 8,704 | 354,907 | 473,029 |
| Size (MB) | 14.84 | 251.44 | 55.18 | 179.86 | 81.16 | 23.78 | 753.35 | 1416.62 |




Fig. 11. Runtime overhead imposed by the UI collector.

UI trees and attributing low level system events to these trees, our technique achieves the effect of partitioning a long running execution to autonomous units, which is the key to avoiding dependence explosion.

IV. EVALUATION

We evaluate UISCOPE by answering following questions.

- **Q1:** How much runtime and space overhead does UISCOPE incur in production environments? (Section IV-A)
- **Q2:** How accurately can UISCOPE cohere UI elements and System objects/subjects? (Section IV-B)
- **Q3:** How effective is UISCOPE in conducting real-world attack investigation and how does it compare to the state-of-the-art? (Section IV-C)

Experiment Setup. We deploy UISCOPE on 7 Windows computers, and run experiments for a week to collect UI events and system events triggered. Note that we do not require those 7 users to perform any specific actions, and our purpose is to collect event traces generated by their daily usage. In addition, we simulate six real-world attacks listed in Table V to evaluate the effectiveness of UISCOPE in attack investigation.

A. Performance Overhead

UISCOPE includes two event logging tools, the UI collector and the system event collector, for trace collection on end hosts. System event traces are deemed to be the major data source for most existing attack investigation systems. UISCOPE uses the Windows built-in audit system ETW as its system event collector. ETW has been evaluated to be quite lightweight, only imposing 0.4% ~ 2.5% runtime overhead [30]. The space overhead caused by ETW mainly depends on what event types ETW is configured to capture. We use ETW to monitor the same system event types (i.e., only the security-relevant types) as existing works [26], [52], [37], [16], [19], and thus the space overhead by ETW in UISCOPE is around 210 MB per week after deploying their system, which

is reasonable. Therefore we focus on evaluating the run-time overhead and space overhead introduced by the UI collector.

1) Runtime overhead: We determine the execution time of an application handling the generated UI events, such as processing the events and storing them, as the runtime overhead. Generally, more frequent UI events incur higher runtime overhead. That is, the runtime overhead hinges on the number of UI events generated per second. We then build a benchmark to simulate various system workloads by triggering UI events of a program. In particular, we use the benchmark to run popular GUI applications, including Chrome, Notepad PlusPlus, Outlook, Skype and so on. We test the benchmark on a computer equipped with 16GB RAM and Intel i7 CPU. Note that unlike system event collector, UI collector does not affect the overhead of non-GUI applications.

Fig. 11 shows the runtime overhead of the UI collector under different workloads. The X-axis in the graph represents the number of UI events per second generated by a program, and the Y-axis shows the average processing time (in milliseconds) used to handle these UI events. The figure shows that 100 UI events only introduce about 20 ms overhead for a program. According to our experiments, typically, a regular application with human interactions triggers less than 10 UI events per second which correspond to around 2 ms overhead; and an application without any human involvement does not trigger any UI events. Thus we believe the runtime overhead introduced by UI collector is negligible (less than 0.2%).

2) Space overhead: Table I presents the details about the collected UI event logs on each host. For each user machine, Table I lists the hardware configuration (i.e., CPU and RAM), Windows OS installed, duration of running the event collectors, number of UI events triggered by daily user behavior, and storage space for the UI events. Note that the users are free to enable or disable UISCOPE, so the duration of event logging varies from user to user. In addition, we did not apply any compression technique for data reduction, so the storage size could be reduced further.

In summary, being deployed on 7 user machines for a week (with a total active duration of 463.4 hours), UISCOPE collected 1416.62 MB UI log including 473,029 UI events and the corresponding UI trees. On average, the UI collector component on one machine generates 3.05 MB event logs per hour. Hence, the space overhead is negligible, compared to system logging.

B. Accuracy of Events Correlation

1) Static Applications: To evaluate the accuracy of event correlation, it is necessary to obtain the ground truth. We did the following experiments with 6 popular applications.

TABLE II. ATTRIBUTION ACCURACY OF STATIC APPLICATIONS

| App | Timestamp | | Timestamp + Initial Event | |
|------------------|--------------|-------------|---------------------------|-------------|
| | TPR | FPR | TPR | FPR |
| Notepad PlusPlus | 100% | 0% | 100% | 0% |
| Adobe Reader | 92.59% | 0% | 100% | 0% |
| Foxit Reader | 100% | 0% | 100% | 0% |
| WinSCP | 94.44% | 17.16% | 100% | 0% |
| Outlook | 99.49% | 2.80% | 98.64% | 2.80% |
| Explorer | 66.67% | 0% | 100% | 0% |
| Average | 92.8% | 3.0% | 98.3% | 0.7% |

For each application, we use a UI automation tool [42] to trigger a typical application specific behavior and collect system events, e.g., file opening and editing for Notepad PlusPlus, Adobe Reader and Foxit Reader, file transfer using WinSCP, email editing on Outlook and file copy-and-paste with Explorer. We repeat this process for three times in isolated environments and extract a common set of system events from these three runs. The pairings of the used UI events and the corresponding common set of system events are considered our ground truth of one behavior.

For each application, we performed its specific behavior 10 times with different settings (e.g., opening and editing 10 different files), and obtain the ground truth for the 10 instances. Next, for each application, we trigger all the 10 behavior instances, collect the generated system events and use UISCOPE to perform the correlation. Then we compare the results with the ground truth to evaluate UISCOPE. As mentioned in Section III-F, we develop two attribution methods: *timestamp-based* and *initial event-based*. We apply different combinations of these attribution methods to the collected traces. Table II shows the attribution accuracy results on the 6 popular applications. The last row shows the average value of True Positive Rate (TPR) and False Positive Rate (FPR) for all the 6 applications.

We can see that on average 92.8% system events can be attributed to the correct UI events only based on timestamps, which is consistent with our observation in Section III-F2. When we apply both methods at the same time, the average TPR improves from 92.8% to 98.3%, and the average FPR reduced from 3.0% to 0.7%. Overall, UISCOPE achieves high attribution accuracy (98.3%) and negligible false positives.

2) *Dynamic Application*: We evaluate UISCOPE on Chrome which is the most complex dynamic application. As mentioned in Section III-F2, the accuracy of UISCOPE for browsers is affected by 1) the types of visited websites, 2) the duration of web sessions, 3) the number of website with background activities. We evaluated UISCOPE on Chrome over these three variables and summarised the results in Table IV. To calculate the TPR and FPR, we use the same method in Section III-F2 to obtain the ground truth for 463 websites listed in Table VI. Table III shows an overview of 1 hour network traffic for those websites.

In this experiment, we randomly select a given number of websites (**NumberOfWebsites**, column 1) from 463 website and load them one by one. The interval of loading two different web pages is a constant number (**StayTime**, column 2). Then we apply UISCOPE to the collected trace and calculate the TPR and FPR by comparing with the collected ground truth data.

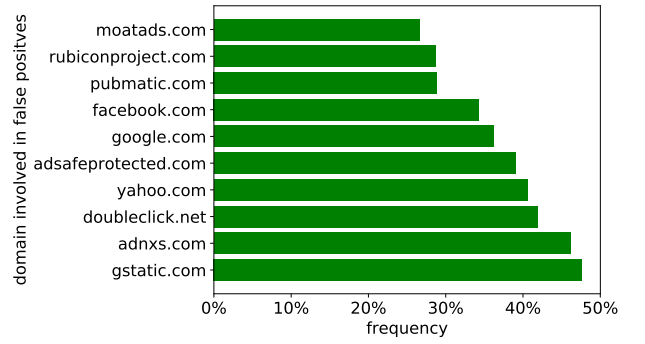


Fig. 12. Frequency of domains occurring in the false positives.

To evaluate the significance of our two approaches (timestamp-based and initial-event based), we apply the timestamp based approach (columns 3 to 5) and the timestamp plus initialization based method (columns 6 to 8). Besides FPR and TPR, we calculate the average number of unique domain involved in the false positives of each tested website (columns 5 and 8). In order to cover more types of websites, we repeat the above process 1000 times and calculate average TPR and FPR for each setting, and Table IV shows the average value for each setting in each row.

From the table, we can see that 1) given a certain **NumberOfWebsites**, increasing **StayTime** would decrease the TPR and FPR. The decreased FPR is reasonable because the longer time users stay on a web page, the more events of background websites would be attributed to the foreground website. TPR is decreased because as time goes by, background websites would initialize new sockets which cannot be attributed to the original websites but rather the current foreground websites. 2) given a certain **StayTime**, increasing **NumberOfWebsites** also decrease the TPR and FPR. With more background websites, the more false-dependencies would be introduced to the foreground website. 3) given a certain **StayTime** and **NumberOfWebsites**, we can see that initial-event based approach could significantly increase the TPR and decrease the FPR, which conforms to the observation 2 mentioned in Section III-F2. That is, most socket sessions are initialized during web page loading.

Although UISCOPE cannot handle newly created socket in the background, the number of unique domains of false positives per website is relatively small (0.61 ~ 1.64). We further explore those false positives of each website and find that they are related to many common domains, which conforms to the observation 3 mentioned in section III-F2. Fig. 12 shows the top 10 common attributed domains by background communication. Y-axis represents domains and X-axis represents the frequency of wrong attribution. From the figure, we can see that these domains are related to advertisement (moatads.com, rubiconproject.com, pubmatic.com, adnxs.com, adsafeprotected.com, and doubleclick.net), Google services (gstatic.com, google.com) and social media (facebook.com and yahoo.com). For those frequent domains, we can remove them to further reduce the FPR.

C. Attack Investigation

We demonstrate the effectiveness of UISCOPE in attack provenance tracking by applying it to 6 real-world attacks,

TABLE III. ONE HOUR NETWORK TRAFFIC FOR DIFFERENT WEBSITE CATEGORIES

| | # of websites | # of socket session per website | # of non-initial events per socket session | Active duration per socket session (second) |
|--------------------------|---------------|---------------------------------|--|---|
| E-commerce and Shopping | 47 | 88.72 | 31.37 | 124.96 |
| Restaurants and Delivery | 34 | 79.55 | 33.57 | 167.12 |
| Accommodation and Hotels | 44 | 57.68 | 41.59 | 170.23 |
| TV Movies and Music | 10 | 260.8 | 265.28 | 128.01 |
| Consumer Electronics | 41 | 57.02 | 41.92 | 155.61 |
| Social Networks | 41 | 125.53 | 23.99 | 77.20 |
| Search Engines | 45 | 21.28 | 33.39 | 183.18 |
| News and Media | 57 | 402.59 | 32.92 | 86.41 |
| Investing | 34 | 105.82 | 32.60 | 142.41 |
| Finance | 10 | 375.3 | 35.76 | 110.19 |
| Sports | 17 | 435.41 | 70.37 | 124.38 |
| Maps | 34 | 86.94 | 29.74 | 145.90 |
| Others | 22 | 134.22 | 209.92 | 87.73 |
| Average | N/A | 146.95 | 54.75 | 112.30 |

TABLE IV. ATTRIBUTION ACCURACY OF DYNAMIC APPLICATIONS

| NumberOfWebsites | StayTime | Timestamp | | | Timestamp + Initial Event | | |
|------------------|----------|-----------|-------|--------------------------------------|---------------------------|-------|--------------------------------------|
| | | TPR | FPR | # of unique domain of FP per website | TPR | FPR | # of unique domain of FP per website |
| 10 | 10 | 84.13 | 1.76% | 2.23 | 93.19% | 0.75% | 0.61 |
| 10 | 20 | 83.56% | 1.82% | 2.35 | 94.07% | 0.76% | 0.61 |
| 10 | 30 | 80.38% | 2.17% | 3.05 | 92.40% | 0.84% | 0.71 |
| 20 | 10 | 77.33% | 1.19% | 3.17 | 89.76% | 0.53% | 0.86 |
| 20 | 20 | 73.11% | 1.41% | 4.38 | 87.98% | 0.63% | 1.01 |
| 20 | 30 | 66.14% | 1.78% | 5.15 | 86.22% | 0.72% | 1.22 |
| 30 | 10 | 72.41% | 0.95% | 4.34 | 87.07% | 0.45% | 1.07 |
| 30 | 20 | 64.46% | 1.22% | 5.56 | 83.68% | 0.56% | 1.37 |
| 30 | 30 | 56.46% | 1.50% | 6.21 | 80.93% | 0.72% | 1.64 |

including Phishing email [49], Remote Code Execution [43], MS Office Macro Attack [46], Credential-based Attack [51], Watering Hole Attack [47], and Insider Attack [50]. In the following, we will use the Remote Code Execution attack to present a case study and other five attacks are summarised in Table V. The first column shows the attack name and reference. The second column summarizes the attack scenario and the last column indicates if UISCOPE can find the root cause of the attack. And we can see that, UISCOPE is capable of finding all root causes of different types of attacks.

1) *Case Study: Remote Code Execution:* Remote Code Execution (RCE) is a vulnerability that can provide an attacker with the ability to execute malicious code and take complete control of an affected host, no matter where the host is geographically located.

Scenario. In this attack, UISCOPE and a threat detector have been installed and enabled on a user’s machine. The user accidentally navigates to a malicious Flappy Bird game website with the browser Edge by clicking on a hyperlink returned from Google Search. The website asks the user to press the ‘enter’ key to control the bird. The website leverages CVE-2018-8495 [43] to perform an attack. Once the enter key is held, the website will invoke a pop-up window asking if the user wants to start, and as soon as the enter key is released, a malicious PowerShell script gets executed on the victim’s machine. While the user is thinking he is playing the game by using the enter key, he actually has been tricked to execute malicious code on his own machine.

Threat alert. Soon, the threat detector installed detects malicious PowerShell script running and thus raises an alert.

TABLE V. ATTACK INVESTIGATION SUMMARY ON THE REST 5 REAL-WORLD ATTACKS

| Attacks | Short Description | Root cause by UISCOPE |
|------------------------------|--|-----------------------|
| Phishing email [49] | Motivating example discussed in Section II-A. | ✓ |
| MS Office Macro Attack [46] | An malicious document was downloaded and executed as an Outlook attachment and the enclosed macro was triggered by Excel to perform malicious behaviors. | ✓ |
| Watering Hole Attack [47] | An malicious file was uploaded to a popular forum. A victim visited the forum through Google search, and downloaded and executed the malicious file. | ✓ |
| Insider Attack [50] | An insider attacker downloaded sensitive files from a FTP server and compress and send out such files through Outlook. | ✓ |
| Credential-based Attack [51] | An attacker accessed the machine with the stolen VNC credential and transfer bank money through passwords automatically saved by Chrome. | ✓ |

Investigation. An incident investigator starts investigation with our tool. With the given threat alert and the audit logs (UI events and system events) collected, UISCOPE could efficiently yield a semantics-rich human-comprehensible causal graph shown in Fig. 13. By provenance tracking in the graph, the investigator can quickly find that the malicious PowerShell script was from the website <http://FakeWebsite.com>, which was opened by clicking a link return by google.com (element B). The inves-

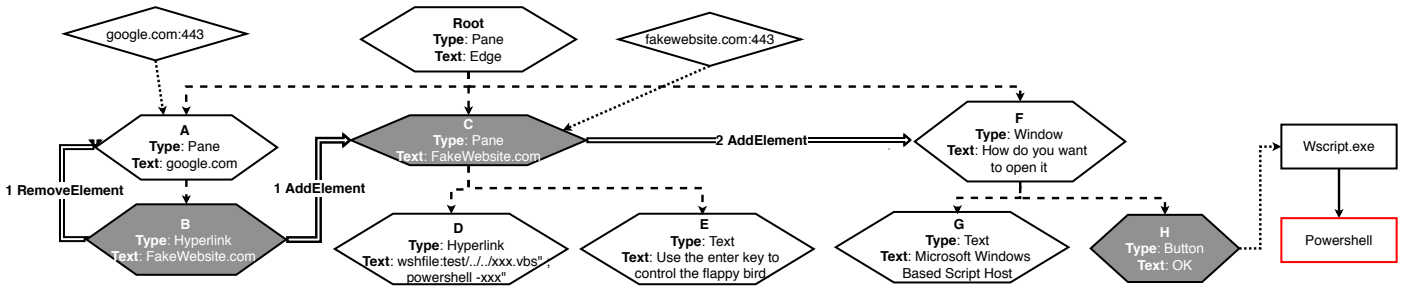


Fig. 13. Causality Graph generated by UISCOPE for the RCE scenario.




Fig. 14. Causality Graph generated by NoDoze for the RCE scenario.

tigator then examines the FakeWebsite and finds that a suspicious hyperlink (element D) contains suspicious keywords (i.e., PowerShell and vbs) in its URL string. Close scrutiny of the suspicious hyperlink would reveal that the hyperlink exploits an Edge vulnerability CVE-2018-8495 [43] to execute a vulnerable benign VBS program xxx.vbs which further allows the attacker to execute any PowerShell code remotely to have full control over the victim’s machine. However, the vulnerability would open a pop-up window to explicitly ask for user’s consent. No user would be fooled into clicking ‘OK’ and run the program. Then the investigator looks into other behaviors of FakeWebsite and finds that after the website was loaded, a pop-up window (element F) with the text “How do you want to open it” showed up, and an OK button (element H) in the window was focused by default. And the website asks the user to press the enter key to control the game character (element E), which causes the OK button to be pressed. This explains how the attacker acquired the permission.

This case demonstrates that UISCOPE can not only assist identifying the attack provenance but also provide fine-grained human-comprehensible contextual semantics to users. Comparatively, Fig. 14 shows the graph generated by NoDoze. Although NoDoze contains the malicious socket fakewebsite.com:443 in the graph because fakewebsite.com was rarely visited, it misses the google.com which is frequently visited and it cannot provide context information to understand the attack or prevent similar attacks from happening again.

V. RELATED WORKS

Along with the previous works discussed in Section II-B, there exist other studies that are related to UISCOPE from different aspects, as discussed below.

Web-based Forensic. JSgraph [27] enables a detailed reconstruction of ephemeral JS-based web attacks by recording fine-grained audit logs related to the execution of JS programs. However, it is application-dependent and they need to instrument the source code of Chromium. Furthermore, JSgraph only

focuses on the provenance tracking of JS execution. Web-browser Record & Replay systems [38], [33] have similar limitations as JSgraph. UISCOPE supports any GUI applications without instrumentation and tracks from low-level system events to high-level UI elements.

Log Reduction. Collectors can easily generate millions of system events by just visiting a website. Thus there exists a large number of works [26], [52], [37], [16], [19] focusing on reducing the log size of system events while preserving the dependency for provenance tracking. UISCOPE collect the same types of system events as those works and our system only reply on must-have fields of system events (e.g., timestamp), thus UISCOPE is orthogonal to these approaches and can incorporate those them to reduce the storage of system events.

Other Digital Forensic Tools. Many other digital forensic tools focus on analyzing digital artifacts (e.g., memory [10] and disk [9]) left in computers after an attack happened. Then security analysts manually analyze and identify causality-related events (e.g., timeline analysis [48], [14]). Comparatively, UISCOPE starts to monitor the system before attacks and automatically correlate system and UI events to construct a dependency graph for provenance tracking. Some forensic tools could extract high-level semantics by analyzing application-specific logs, e.g., email forensic tools. However, those tools highly rely on the format of application logs. If logs are stored with a unknown format, those tools cannot works. Comparatively, UISCOPE supports capture high-level semantics of all GUI applications by monitoring UI events.

VI. CONCLUSION

We develop UISCOPE, an accurate, instrumentation-free, and deterministic attack investigation system with high visibility. UISCOPE highlights the critical role of user interaction with the system through GUI in partitioning system operations and tracing the provenance of attack. By leveraging user-space user interaction logs to complement kernel-level system events with human-perceivable contextual semantics, UISCOPE is able to provide deterministic, accurate and human-comprehensible investigation results, even to ordinary users who are not tech-savvy. UISCOPE is lightweight and efficient, incurring negligible performance overhead. UISCOPE is applicable in production environments, requiring no end system change or instrumentation. Our evaluation shows that UISCOPE can efficiently and precisely identify the provenance of real-world attacks.

REFERENCES

- [1] “Alexa top websites,” <https://www.alexacom/topsites>, accessed on 2019-09-10.
- [2] “Captain hook:pirating avs to bypass exploit mitigations,” <https://goo.gl/zVyUAL>, accessed on 2019-09-10.
- [3] “Chromedriver - webdriver for chrome,” <https://chromedriver.chromium.org/getting-started>, accessed on 2019-09-10.
- [4] “Clickjacking,” <https://en.wikipedia.org/wiki/Clickjacking>, accessed on 2019-09-10.
- [5] “Kernel Patch Protection — Wikipedia, The Free Encyclopedia,” <https://goo.gl/S4idr7>, accessed on 2019-09-10.
- [6] “Khobe 8.0 earthquake for windows desktop security software,” <https://goo.gl/5UhzpQ>, accessed on 2019-09-10.
- [7] “Login detection,” <https://robinlinus.github.io/socialmedia-leak/>, accessed on 2019-09-10.
- [8] “Plague in (security) software drivers,” <https://goo.gl/kmycvb>, accessed on 2019-09-10.
- [9] “Sleuth kit,” <https://www.sleuthkit.org/>, accessed on 2019-09-10.
- [10] “Volatility: An advanced memory forensics framework,” <https://github.com/volatilityfoundation/volatility>, accessed on 2019-09-10.
- [11] A. Developer, “NSAccessibility,” <https://apple.co/2w2z8Ww>, 2019, accessed on 2019-09-10.
- [12] P. Gao, X. Xiao, D. Li, Z. Li, K. Jee, Z. Wu, C. H. Kim, S. R. Kulkarni, and P. Mittal, “Saq: A stream-based query system for real-time abnormal system behavior detection,” in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 639–656.
- [13] P. Gao, X. Xiao, Z. Li, F. Xu, S. R. Kulkarni, and P. Mittal, “AIQL: Enabling efficient attack investigation from system monitoring data,” in *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, 2018, pp. 113–126.
- [14] Google, “Timesketch: Collaborative forensic timeline analysis,” <http://bit.ly/2JmCnAT>, 2019, accessed on 2019-09-10.
- [15] T. U. S. Government, “It accessibility laws and policies,” <http://bit.ly/2VFOwaD>, 2019, accessed on 2019-09-10.
- [16] W. U. Hassan, L. Aguse, N. Aguse, A. Bates, and T. Moyer, “Towards scalable cluster auditing through grammatical inference over provenance graphs,” in *Network and Distributed Systems Security Symposium (NDSS 18)*, 2018.
- [17] W. U. Hassan, S. Guo, D. Li, Z. Chen, K. Jee, Z. Li, and A. Bates, “Nodoze: Combatting threat alert fatigue with automated provenance triage,” in *Network and Distributed Systems Security Symposium (NDSS 19)*, 2019.
- [18] M. N. Hossain, S. M. Milajerdi, J. Wang, B. Eshete, R. Gjomemo, R. Sekar, S. D. Stoller, and V. Venkatakrishnan, “Sleuth: Real-time attack scenario reconstruction from cots audit data,” in *26th USENIX Security Symposium (USENIX Security 17)*, 2017, pp. 487–504.
- [19] M. N. Hossain, J. Wang, R. Sekar, and S. D. Stoller, “Dependence-preserving data compaction for scalable forensic analysis,” in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 1723–1740.
- [20] K. Jee, G. Portokalidis, V. P. Kemerlis, S. Ghosh, D. I. August, and A. D. Keromytis, “A general approach for efficiently accelerating software-based dynamic data flow tracking on commodity hardware,” in *Network and Distributed Systems Security Symposium (NDSS 12)*, 2012.
- [21] Y. Ji, S. Lee, E. Downing, W. Wang, M. Fazzini, T. Kim, A. Orso, and W. Lee, “Rain: Refinable attack investigation with on-demand inter-process information flow tracking,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS 17)*, 2017, pp. 377–390.
- [22] S. T. King and P. M. Chen, “Backtracking intrusions,” *ACM SIGOPS Operating Systems Review*, vol. 37, no. 5, pp. 223–236, 2003.
- [23] S. T. King, Z. M. Mao, D. G. Lucchetti, and P. M. Chen, “Enriching intrusion alerts through multi-host causality,” in *Network and Distributed Systems Security Symposium (NDSS 05)*, 2005.
- [24] Y. Kwon, F. Wang, W. Wang, K. H. Lee, W.-C. Lee, S. Ma, X. Zhang, D. Xu, S. Jha, G. Ciocarlie *et al.*, “Mci: Modeling-based causality inference in audit logging for attack investigation,” in *Network and Distributed Systems Security Symposium (NDSS 18)*, 2018.
- [25] K. H. Lee, X. Zhang, and D. Xu, “High accuracy attack provenance via binary-based execution partitioning,” in *Network and Distributed Systems Security Symposium (NDSS 13)*, 2013.
- [26] —, “Loggc: garbage collecting audit log,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS 13)*, 2013, pp. 1005–1016.
- [27] B. Li, P. Vadrevu, K. H. Lee, and R. Perdisci, “Jsgraph: Enabling reconstruction of web attacks via efficient tracking of live in-browser javascript executions,” in *NDSS*, 2018.
- [28] B. Linux, “ATK Package,” <http://bit.ly/2WJM92Y>, 2019, accessed on 2019-09-10.
- [29] Y. Liu, M. Zhang, D. Li, K. Jee, Z. Li, Z. Wu, J. Rhee, and P. Mittal, “Towards a timely causality analysis for enterprise security,” in *Network and Distributed Systems Security Symposium (NDSS 18)*, 2018.
- [30] S. Ma, K. H. Lee, C. H. Kim, J. Rhee, X. Zhang, and D. Xu, “Accurate, low cost and instrumentation-free security audit logging for windows,” in *Proceedings of the 31st Annual Computer Security Applications Conference (ACSAC 15)*, 2015.
- [31] S. Ma, J. Zhai, F. Wang, K. H. Lee, X. Zhang, and D. Xu, “MPI: Multiple perspective attack investigation with semantic aware execution partitioning,” in *26th USENIX Security Symposium (USENIX Security 17)*, 2017, pp. 1111–1128.
- [32] S. Ma, X. Zhang, and D. Xu, “Protracer: Towards practical provenance tracing by alternating between logging and tainting,” in *Network and Distributed Systems Security Symposium (NDSS 16)*, 2016.
- [33] J. W. Mickens, J. Elson, and J. Howell, “Mugshot: Deterministic capture and replay for javascript applications,” in *NSDI*, vol. 10, 2010, pp. 159–174.
- [34] Microsoft, “UI Automation - Windows applications,” <http://bit.ly/2Q4Vn7v>, 2018, accessed on 2019-09-10.
- [35] —, “UI Automation Events Overview,” <http://bit.ly/2WJM8yQ>, 2018, accessed on 2019-09-10.
- [36] Microsoft, “Property identifiers,” <https://docs.microsoft.com/zh-cn/windows/win32/winauto/uiauto-entry-propids>, 2019, accessed on 2019-09-10.
- [37] K.-K. Muniswamy-Reddy, D. A. Holland, U. Braun, and M. I. Seltzer, “Provenance-aware storage systems,” in *2006 USENIX Annual Technical Conference (USENIX ATC 06)*, 2006, pp. 43–56.
- [38] C. Neasbitt, B. Li, R. Perdisci, L. Lu, K. Singh, and K. Li, “Webcapsule: Towards a lightweight forensic engine for web browsers,” in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 133–145.
- [39] Panda, “Platform for architecture-neutral dynamic analysis,” <http://bit.ly/2W5KaJc>, 2019, accessed on 2019-09-10.
- [40] K. Pei, Z. Gu, B. Saltaformaggio, S. Ma, F. Wang, Z. Zhang, L. Si, X. Zhang, and D. Xu, “HERCULE: attack story reconstruction via community discovery on correlated log graph,” in *Proceedings of the 32nd Annual Computer Security Applications Conference (ACSAC 16)*, 2016.
- [41] S. Rekhis and N. Boudriga, “A hierarchical visibility theory for formal digital investigation of anti-forensic attacks,” *Computers & Security*, vol. 31, no. 8, pp. 967–982, 2012.
- [42] TinyTask, “Simple + fast + free: Automation for everyone,” <http://bit.ly/2vXwrWm>, 2019, accessed on 2019-09-10.
- [43] Wikipedia contributors, “Cve-2018-8495,” <http://bit.ly/30nvxjN>, 2019, accessed on 2019-09-10.
- [44] —, “Event tracing for windows,” <http://bit.ly/2EbzKxM>, 2019, accessed on 2019-09-10.
- [45] —, “Hidden treasure: Intrusion detection with etw,” <http://bit.ly/2VGODUZ>, 2019, accessed on 2019-09-10.
- [46] —, “Macro virus — Wikipedia, the free encyclopedia,” <http://bit.ly/2w3W9Z2>, 2019, [Online; accessed 13-May-2019].
- [47] —, “Many watering holes, targets in hacks that netted facebook, twitter and apple,” <http://bit.ly/30otNqs>, 2019, accessed on 2019-09-10.
- [48] —, “Plaso: super timeline all the things,” <http://bit.ly/2WP3MyC>, 2019, accessed on 2019-09-10.

- [49] —, “Spear phishing campaign targets ukraine government and military,” <https://bit.ly/2XaiUGu>, 2019, accessed on 2019-09-10.
- [50] —, “Target to pay \$18.5m for 2013 data breach that affected 41 million consumers,” <http://bit.ly/2W77ZQU>, 2019, accessed on 2019-09-10.
- [51] —, “Trickbot adds remote application credential-grabbing capabilities to its repertoire,” <http://bit.ly/2VubtZk>, 2019, accessed on 2019-09-10.
- [52] Z. Xu, Z. Wu, Z. Li, K. Jee, J. Rhee, X. Xiao, F. Xu, H. Wang, and G. Jiang, “High fidelity data reduction for big data security dependency analyses,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS 16)*, 2016, pp. 504–516.

APPENDIX

TABLE VI. TESTED WEBSITE LIST

| | | | | |
|---------------------------------|-----------------------------|-----------------------------|------------------------------------|------------------------------|
| Accommodation and Hotels | www.vrbo.com | www.pegipegi.com | www.ihg.com | www.thetrainline.com |
| | www.marriott.com | www.airbnb.com.au | www.hotels.com | www.hyatt.com |
| | www.airbnb.ru | www.trivago.com.tr | www.airbnb.de | www.travelzoo.com |
| | www.expedia.it | www.airbnb.co.uk | www.choicehotels.com | hotelscan.com |
| | www.abritel.fr | www.airbnb.fr | www.caesars.com | www.costcotravel.com |
| | www.oyorooms.com | www.airbnb.com.br | www.trivago.de | secure.accorhotels.com |
| | magazine.trivago.it | www.wyndhamhotels.com | www.agoda.com | www.fewo-direkt.de |
| | www.airbnb.ca | www.jalan.net | magazine.trivago.es | www.booking.com |
| | www.trivago.co.uk | www.homeaway.com | blog.couchsurfing.com | hiltonhonors3.hilton.com |
| | www.expedia.co.uk | www.airbnb.it | www.hostelworld.com | www.secretescapes.com |
| www.trivago.com.br | www.airbnb.es | www.gosur.com | discover.expediapartnercentral.com | |
| Consumer Electronics | sharp.cn | www.made-in-china.com | consumer.huawei.com | www.sony.jp |
| | www.samsung.com | www.lavamobiles.com | na.panasonic.com | www.boulangier.com |
| | www.mediamarkt.de | www.micromaxinfo.com | www.apple.com | www.gearbest.com |
| | www.lg.com | www.mediaexpert.pl | shop.huawei.ru | www.darty.com |
| | www.pccomponentes.com | www.mediamarkt.es | www.vivo.com.cn | www.vatanbilgisayar.com |
| | www.onlinetrade.ru | www.uscellular.com | www.oneplus.com | www.mvideo.ru |
| | www.saturn.de | www.kimovil.com | www.newegg.com | unity.com |
| | www.gadgetsnow.com | www.shutterfly.com | www8.hp.com | hd.oppo.com |
| | www.sonymobile.com | www.roku.com | www.currys.co.uk | www.bhphotovideo.com |
| | www.euro.com.pl | www.mi.com | www.vmall.com | eshop.htc.com |
| E-commerce and Shopping | www.ebay-kleinanzeigen.de | www.wildberries.ru | search.jd.com | hz.58.com |
| | shopping.yahoo.co.jp | www.olx.ua | www.amazon.es | www.target.com |
| | shopee.co.id | www.leboncoin.fr | www.ebay.com | articulo.mercadolibre.com.ar |
| | www.flipkart.com | sale.aliexpress.com | www.amazon.de | www.amazon.co.jp |
| | www.ebay.co.uk | www.olx.pl | www.ebay.com.au | www.amazon.fr |
| | err.tmall.com | market.yandex.ru | www.hepsiburada.com | www.sahibinden.com |
| | listado.mercadolibre.com.mx | www.amazon.co.uk | www.tokopedia.com | www.etsy.com |
| | allegro.pl | list.tmall.com | www.alibaba.com | login.tmall.com |
| | kakaku.com | www.amazon.com | slickdeals.net | www.bukalapak.com |
| | www.ebay.de | produto.mercadolivre.com.br | world.taobao.com | www.amazon.ca |
| Finance | www.amazon.it | hongkong.craigslist.org | miao.tmall.com | www.dmm.com |
| | www.amazon.in | www.walmart.com | www.avito.ru | |
| | www.binance.com | cn.investing.com | www.tradingview.com | www.fidelity.com |
| Investing | www.investopedia.com | www.schwab.com | www.boursorama.com | iqoption.com |
| | www.moneycontrol.com | www.alipay.com | | |
| | olymptrade.com | pit.blockchain.com | www.advfn.com | www.ig.com |
| | new.nasdaq.com | etherscan.io | www.rakuten-sec.co.jp | jfinfo.com |
| | www.etoro.com | us.etrade.com | www.aastocks.com | finance.eastmoney.com |
| | yuanchuang.10jqka.com.cn | www.forexfactory.com | www.rico.com.vc | www.kitco.com |
| | robinhood.com | news.cnyes.com | finviz.com | www.fool.com |
| | nifty50etf.nseindia.com | www.xm.com | stocktwits.com | www.morningstar.com |
| | www.finanzen.net | minkabu.jp | nikkei225jp.com | www.tdameritrade.com |
| | www.ml.com | zerodha.com | about.vanguard.com | www1.oanda.com |
| Maps | xueqiu.com | www.clear.com.br | | |
| | www.viamichelin.it | www.falk.de | wikiroutes.info | 2gis.ru |
| | fr.mappy.com | www.driveplaza.com | www.onefivefive.com | actualite.lachainemeteo.com |
| | www.meteofrance.com | economia.uol.com.br | moscow.flamp.ru | mapfan.com |
| | www.mapquest.com | www.openstreetmap.org | www.targeo.pl | www.langenscheidt.com |
| | www.city-data.com | www.bustime.ru | www.mapion.co.jp | www.arasickackm.com |
| | www.marinetraffic.com | www.viamichelin.fr | en.mapy.cz | www.here.com |
| www.geoportail.gouv.fr | www.tuttocitta.it | map.goo.ne.jp | www.komoot.de | |
| News and Media | www.trendsmat.com | geoguessr.com | www.worldatlas.com | ekitan.com |
| | www.viamichelin.de | ditu.amap.com | | |
| | www.naver.com | news.finance.yahoo.co.jp | sohu.com | www.ifeng.com |
| | www.ukr.net | timesofindia.indiatimes.com | www.rambler.ru | edition.cnn.com |
| | matome.naver.jp | www.nytimes.com | www.toutiao.com | sina.cn |
| | elpais.com | www.goo.ne.jp | finance.yahoo.com | www.yidiazixun.com |
| | www.globo.com | news.yahoo.com | www.iqiyi.com | www.sina.com.cn |
| | www.yahoo.co.jp | www.livedoor.com | finance.sina.com.cn | www.163.com |
| | drudgereport.com | map.yahoo.co.jp | www.bbc.com | www.tribunnews.com |
| | sports.news.naver.com | www.onet.pl | news.google.com | economictimes.indiatimes.com |
| | section.blog.naver.com | www.ndtv.com | www.qq.com | www.indiatimes.com |
| | news.yahoo.co.jp | news.mail.ru | search.yahoo.co.jp | www.foxnews.com |
| | www.wp.pl | www.bild.de | www.t-online.de | headlines.yahoo.co.jp |
| | www.interia.pl | www.washingtonpost.com | www.msn.com | www.dailymail.co.uk |
| | us.yahoo.com | www.yahoo.com | vnexpress.net | www.rt.com |
| www.uol.com.br | lenta.ru | map.naver.com | www.detik.com | |
| www.theguardian.com | | | | |

| | | | | |
|---------------------------------|------------------------------|----------------------------|-------------------------|---------------------------------|
| Restaurants and Delivery | www.odysys.net | www.just-eat.co.uk | www.swiggy.com | www.doordash.com |
| | www.wongnai.com | popslotscasino.com | www.inmoment.com | www.subway.com.hk |
| | www.seamless.com | www.ubereats.com | www.skipthedishes.com | www.mcdonalds.com |
| | www.chick-fil-a.com | www.zomato.com | www.opentable.com | dominos.com.mx |
| | menu.wendys.com | www.lieferando.de | sg.theasianparent.com | www.bk.com |
| | www.pizzahut.com.hk | zmenu.com | www.chowhound.com | www.olivegarden.com |
| | postmates.com | www.icomera.com | order.littlecaesars.com | hds-streaming.com |
| | www.pyszne.pl | www.starbucks.com | www.singleplatform.com | deliveroo.co.uk |
| www.grubhub.com | www.mcdonalds.co.jp | | | |
| Search Engines | www.google.ca | www.google.co.in | special-offers.online | www.so.com |
| | www.baidu.com | www.justdial.com | www.startpage.com | www.google.co.th |
| | www.google.fr | www.google.com.mx | www.ask.com | www.google.com.au |
| | www.sogou.com | www.google.be | www.google.ru | search.daum.net |
| | www.google.com.ar | www.google.com.vn | www.bing.com | www.google.de |
| | www.google.nl | yandex.ru | www.google.com.pe | map.baidu.com |
| | www.google.ro | www.google.pl | duckduckgo.com | www.google.co.jp |
| | www.google.com.tr | www.google.com.tw | www.google.com.hk | www.wjx.cn |
| | m.sm.cn | www.google.es | www.google.com.ua | yandex.com.tr |
| | search.myway.com | www.google.it | www.google.co.uk | www.google.cl |
| www.google.com.br | y2mate.com | whatismyip.li | www.hao123.com | |
| search.seznam.cz | | | | |
| Social Networks | ostrovok.ru | www.reddit.com | namu.wiki | incorrect-good-omens.tumblr.com |
| | programminghumour.tumblr.com | story.snapchat.com | www.pinterest.com.mx | www.pinterest.de |
| | www.slideshare.net | www.shafa.com | ask.fm | www.pinterest.fr |
| | bakusai.com | www.pinterest.co.uk | www.znds.com | enterprise.foursquare.com |
| | ok.ru | www.whatsapp.com | www.weibo.com | vk.com |
| | www.messenger.com | www.pp.cn | perple.exblog.jp | lihkg.com |
| | twitter.com | www.facebook.com | www.ptt.cc | www.dangbei.com |
| | zhuanlan.zhihu.com | www.linkedin.com | ameblo.jp | www.meetup.com |
| | medium.com | www.instagram.com | www.tagged.com | www.fiverr.com |
| | www11.eyny.com | www.pinterest.com | www.pinterest.es | zalo.me |
| www.dcard.tw | | | | |
| Sports | sports.yahoo.com | www.yr.no | lequipe.fr | www.americanas.com.br |
| | www.gazzetta.it | www.kooora.com | www.mundodeportivo.com | www.espnricinfo.com |
| | www.marca.com | www.goal.com | bbs.hupu.com | as.com |
| | global.espn.com | www.championat.com | www.mlb.com | www.skysports.com |
| | www.cricbuzz.com | | | |
| TV Movies and Music | www.nicovideo.jp | www.crunchyroll.com | soundcloud.com | www.dailymotion.com |
| | www.imdb.com | www.google.com | www.youtube.com | www.netflix.com |
| | www.vesti.ru | vimeo.com | | |
| Others | stackoverflow.com | Xinhuanet.com | Yahoo.co.jp | www.bestbuy.ca |
| | www.trivago.hk | bbs.tianya.cn | www.twitch.tv | outlook.live.com |
| | www.papajohnschina.com | airregi.jp | en.wikipedia.org | www.kawauchisyun.com |
| | templates.office.com | www.gojek.com | www.fishbowl.com | www.okezone.com |
| | www.360.cn | acomputerblog.blogspot.com | www.csdn.net | www.att.com |
| | www.microsoft.com | Babytree.com | | |