(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
22 January 2015 (22.01.2015)

WIPO | PCT

(10) International Publication Number
WO 2015/009411 A1

(51) International Patent Classification:
G06F 21/56 (2013.01) *H04L 29/06* (2006.01)

(21) International Application Number:
PCT/US2014/043724

(22) International Filing Date:
23 June 2014 (23.06.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
13/945,394 18 July 2013 (18.07.2013) US

(71) Applicant: **FIREEYE, INC.** [US/US]; 1440 McCarthy Boulevard, Milpitas, California 95305 (US).

(72) Inventors: **PIDATHALA, Vinay**; 1440 McCarthy Boulevard, Milpitas, California 95305 (US). **UYENO, Henry**; 1440 McCarthy Boulevard, Milpitas, California 95305 (US).

(74) Agents: **SCHAAL, William W.** et al.; BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, 1279 Oakmead Parkway, Sunnyvale, California 94085 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

(54) Title: SYSTEM AND METHOD FOR DETECTING MALICIOUS LINKS IN ELECTRONIC MESSAGES

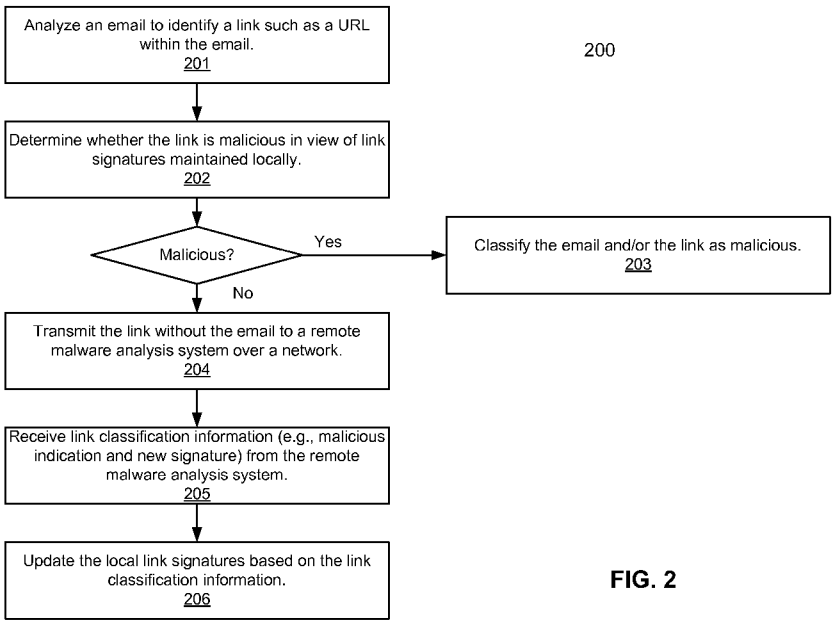

FIG. 2

(57) Abstract: According to one embodiment, in response to receiving a plurality of uniform resource locator (URL) links for malicious determination, any known URL links are removed from the URL links based on a list of known link signatures. For each of remaining URL links that are unknown, a link analysis is performed on the URL link based on link heuristics to determine whether the URL link is suspicious. For each of the suspicious URL links, a dynamic analysis is performed on a resource of the suspicious URL link. It is classified whether the suspicious URL link is a malicious link based on a behavior of the resource during the dynamic analysis.

WO 2015/009411 A1

SYSTEM AND METHOD FOR DETECTING MALICIOUS LINKS IN ELECTRONIC MESSAGES

RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Patent Application No. 61/841,210, filed June 28, 2013, which is incorporated by reference herein in its entirety.

FIELD OF THE INVENTION

[0001] Embodiments of the present invention relate generally to malware detection. More particularly, embodiments of the invention relate to detecting malicious links in electronic messages.

BACKGROUND

[0002] Malicious software, or malware for short, may include any program or file that is harmful by design to a computer. Malware includes computer viruses, worms, Trojan horses, adware, spyware, and any programming that gathers information about a computer or its user or otherwise operates without permission. The owners of the computers are often unaware that these programs have been added to their computers and are often similarly unaware of their function.

[0003] Malicious network content is a type of malware distributed over a network via websites, e.g., servers operating on a network according to a hypertext transfer protocol (HTTP) standard or other well-known standard. Malicious network content distributed in this manner may be actively downloaded and installed on a computer, without the approval or knowledge of its user, simply by the computer accessing the web site hosting the malicious network content (the "malicious web site"). Malicious network content may be embedded within objects associated with web pages hosted by the malicious web site. Malicious network content may also enter a computer on receipt or opening of email. For example, email may contain an attachment, such as a PDF document, with embedded malicious executable programs. Furthermore, malicious content may exist in files contained in a computer memory or storage device, having infected those files through any of a variety of attack vectors.

[0004] Various processes and devices have been employed to prevent the problems associated with malicious content. For example, computers often run antivirus scanning software that scans a particular computer for viruses and other forms of malware. The scanning typically involves automatic detection of a match between content stored on the computer (or attached media) and a library or database of signatures of known malware. The scanning may be initiated manually or based on a schedule specified by a user or system administrator associated with the particular computer. Unfortunately, by the time malware is detected by the scanning software, some damage on the computer or loss of privacy may have already occurred, and the malware may have propagated from the infected computer to other computers. Additionally, it may take days or weeks for new signatures to be manually created, the scanning signature library updated and received for use by the scanning software, and the new signatures employed in new scans.

[0005] Moreover, anti-virus scanning utilities may have limited effectiveness to protect against all exploits by polymorphic malware. Polymorphic malware has the capability to mutate to defeat the signature match process while keeping its original malicious capabilities intact. Signatures generated to identify one form of a polymorphic virus may not match against a mutated form. Thus polymorphic malware is often referred to as a family of virus rather than a single virus, and improved anti-virus techniques to identify such malware families is desirable.

[0006] Another type of malware detection solution employs virtual environments to replay content within a sandbox established by virtual machines (VMs). Such solutions monitor the behavior of content during execution to detect anomalies that may signal the presence of malware. One such system offered by FireEye®, Inc., the assignee of the present patent application, employs a two-phase malware detection approach to detect malware contained in network traffic monitored in real-time. In a first or "static" phase, a heuristic is applied to network traffic to identify and filter packets that appear suspicious in that they exhibit characteristics associated with malware. In a second or "dynamic" phase, the suspicious packets (and typically only the suspicious packets) are replayed within one or more virtual machines. For example, if a user is trying to download a file over a network, the file is extracted from the network traffic and analyzed in the virtual machine. The results of the analysis aids in determining whether the file is malicious. The two-phase malware detection solution may detect numerous types of malware and, even malware missed by other commercially available approaches. Through verification, the two-phase malware detection solution may also achieve a significant reduction of false positives relative to such other commercially available approaches. Dealing with false positives in malware detection may needlessly slow or interfere with

download of network content or receipt of email, for example. This two-phase approach has even proven successful against many types of polymorphic malware and other forms of advanced persistent threats.

[0007] Network traffic can be in a variety of forms. One type of network traffic is in a form of emails, where an email may contain an attachment and/or a link such as a universal resource locator (URL) link. An email may or may not be malicious dependent upon whether a link embedded therein is associated with a remotely located malicious resource. There has been a lack of efficient ways for detecting whether a link is malicious.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Embodiments of the invention are illustrated by way of example and not limitation in the figures of the accompanying drawings in which like references indicate similar elements.

[0009] Figure 1 is a block diagram illustrating a network system for malware detection according to one embodiment of the invention.

[0010] Figure 2 is a flow diagram illustrating a method of detecting malicious links according to one embodiment of the invention.

[0011] Figure 3 is a flow diagram illustrating a method of detecting malicious links according to another embodiment of the invention.

[0012] Figure 4 is a flow diagram illustrating a method for detecting malicious links according to another embodiment of the invention.

[0013] Figure 5 is a flow diagram illustrating a method for detecting malicious links according to another embodiment of the invention.

[0014] Figure 6 is a block diagram of a computer network system deploying a malicious content detection system according to one embodiment of the invention.

[0015] Figure 7 is a block diagram illustrating an example of a data processing system which may be used with one embodiment of the invention.

DETAILED DESCRIPTION

[0016] Various embodiments and aspects of the inventions will be described with reference to details discussed below, and the accompanying drawings will illustrate the various embodiments. The following description and drawings are illustrative of the invention and are not to be construed as limiting the invention. Numerous specific details are described to provide

a thorough understanding of various embodiments of the present invention. However, in certain instances, well-known or conventional details are not described in order to provide a concise discussion of embodiments of the present inventions.

[0017] Reference in the specification to "one embodiment" or "an embodiment" means that a particular feature, structure, or characteristic described in conjunction with the embodiment can be included in at least one embodiment of the invention. The appearances of the phrase "in one embodiment" in various places in the specification do not necessarily all refer to the same embodiment. Also, the term "email" generally denotes a communication message being digital data with a particular format such as one or more packet(s), frame(s), or any other series of bits having a prescribed format, which may include, but not limited or restricted to an electronic mail message, an instant message (IM), or another type of communication message.

[0018] According to some embodiments, one or more malware detection appliances, such as email malware detection systems, are deployed strategically to capture and detect emails that include a URL link. Each of the malware detection appliances is configured to monitor email traffic and to recognize a URL link within the emails. Once the malware detection appliance detects a URL link, the malware detection appliance extracts the URL link from an email and performs an analysis on the extracted link based on malicious link signatures maintained locally within the malware detection appliance. If it is determined that the extract URL link is considered as malicious in view of the malicious link signatures, malware detection appliance classifies the URL link as a malicious link and may alert an administrator or user accordingly. If the malware detection appliance cannot determine whether the URL link is malicious based on the locally maintained malicious link signatures, it sends the extracted URL link to a remote malware analysis system over a network such as the Internet (also referred to as a dedicated malware analysis system in the cloud). In one embodiment, the malware detection appliance only sends the URL link to the remote malware analysis system without sending the remaining content or any identifying information of the email since such information may be confidential. A malware detection appliance may constantly or periodically send the detected URL links to the malware analysis system(s) in the cloud, which may be implemented as a centralized or distributed malware analysis system.

[0019] In response to the URL links received from the malware detection appliances, according to one embodiment, the malware analysis system is configured to perform a link analysis, for example, within a dedicated analysis operating environment (e.g., a sandboxed environment or virtual machine), on the received URL links using link heuristics to determine whether the received URL link or links are likely malicious. If the malware analysis system

determines that the received URL link or links are likely malicious based on the link heuristics, according to one embodiment, the malware analysis system generates one or more new malicious link signatures and distributes the new malicious link signatures back to the malware detection appliances for future local malicious link detection. Since the link analysis is performed at a dedicated malware analysis system remotely, a malware detection appliance does not have to interact with the links, which may expose the identities of the email recipients and increase workload of the malware detection appliance. Note that throughout this application, techniques have been described to detect malicious URL links within emails. However, the techniques described herein can also be applied to other types of links that may be embedded within other types of electronic documents such as word documents or portable document format (PDF) documents.

[0020] Figure 1 is a block diagram illustrating a network system for malware detection according to one embodiment of the invention. Referring to Figure 1, system 100 includes, but is not limited to, one or more malware detection appliances 102-103 communicatively coupled to malware analysis system 101 over network 104. Malware detection appliances 102-103 may be strategically deployed as appliances at various locations or LANs for malicious content detection. Malware detection appliances may be associated with the same or different organizations. Each of malware detection appliances 102-103 may include a malware detection system or module that can perform a static analysis or detection, a dynamic analysis or detection, or both, which may be described in details further below. Malware detection appliances 102-103 may offload at least a portion of the malware detection processes to a dedicated malware analysis system, such as malware analysis system 101. Network 104 may be any type of networks such as a local area network (LAN), a wide area network (WAN) such as the Internet, or a combination thereof.

[0021] According to one embodiment, any of malware detection appliances 102-103 may operate as an email malware detection (EMD) system for detecting malicious emails. For example, any of malware detection appliances 102-103 may be communicatively coupled to an email server to monitor the email traffic and perform a malicious content analysis on the emails to determine whether any of the emails are malicious emails. Whether an email should be considered as a malicious email may be determined based in part on content of the email, the attachment of the email, and/or a URL link embedded within the email. The attachment can be analyzed using the dynamic content detection techniques while the text or content of the email can be analyzed using static or heuristics content detection techniques. The combination of the

static and dynamic detections on the emails may further improve the accuracy of the overall detection.

[0022] According to one embodiment, each of the malware detection appliances 102-103 includes a link extractor (e.g., link extractors 105-106) to recognize a URL link within an email and to analyze the extracted link in view of the malicious link signatures maintained locally (e.g., link signatures 107-108). Specifically, for the purpose of illustration, in response to an email received at malware detection appliance 102 for malware detection, link extractor 105 is configured to recognize a URL link within email and extract the link from the email. The URL link may be analyzed by link extractor 105 or by another analysis module (not shown) based on malicious link signatures 107 maintained locally. Link extractor 105 may match at least a portion of the URL characters against the malicious link signatures 107 to determine whether the URL link is similar or identical to the malicious links represented by malicious link signatures 107. If at least a portion of the URL link matches at least one of the malicious link signatures 107 (e.g., a domain name, uniform resource identifier or URI, or a particular character sequence or sequences), the URL link and/or the associated email may be classified as malicious and an alert may be issued to a user (e.g., administrator(s), email recipient(s)).

[0023] If link extractor 105 cannot determine whether the URL link is malicious based on malicious link signatures 107, according to one embodiment, link extractor 105 sends (e.g., via paths 121-122) the URL link to malware analysis system 101 over network 104, without sending any remaining content of the email to maintain the confidentiality of the email content. The links received from malware detection appliances 102-103 may be stored in link database 110. Malware analysis system 101 may periodically receive URL links from many malware detection appliances. The malware detection appliances may be deployed at various geographic locations and associated with different sites of the same corporate organization or different corporate organizations, while malware analysis system 101 provides malware analysis services to clients that employ the malware detection appliances.

[0024] According to one embodiment, in response to the URL links received from malware detection appliances 102-103 and stored in link database 110, link analysis module 109 is configured to perform a link analysis on each of the URL links using link heuristics 111. Link heuristics may be generated or compiled over a period of time based on the analysis of many different links. For each of links that has been determined by link analysis module 109 as a malicious link based on heuristics 111, link analysis module 109 generates a new malicious link signature as part of malicious signatures 112. Thereafter, malware analysis system 101 distributes (e.g., via paths 123-124) the new malicious link signatures to malware detection

appliances 102-103 to be integrated into malicious link signatures 107-108 for subsequent local detections.

[0025] Since there may be many links received from malware detection appliances 102-103, according to one embodiment, link analysis module 109 is configured to initially apply a white list (e.g., a list of domain names that are known to be non-malicious) to screen out or filter any non-malicious links, which may reduce a number of links to be examined. For each of the remaining links, link analysis module 109 examines the link based on heuristics to determine whether the link should be considered as suspicious. This may be based on whether the link or portion thereof has one or more characteristics associated with malware at a probability satisfying a set threshold. While the probability may not be sufficient to declare the link as malicious, the probability may be sufficient to classify the link as suspicious and requiring further analysis. In one embodiment, if a URL link contains an Internet protocol (IP) address instead of a host name (e.g., http//203.x.x.x instead of http//WebSiteName.com), optionally in view of other considerations provided by link heuristics 111, the link may be considered as a possible suspicious link. For example, other heuristics may factor in the frequency of encountering a particular URL or domain name as compared with a user specified or machine set threshold. Other heuristics may also be used instead or in addition to the foregoing in finding a link "suspicious".

[0026] Once a link is considered as a suspicious or a possibly malicious link, also referred to as a malicious link suspect, according to one embodiment, link analysis module 109 is configured to access the link in an attempt to further confirm that the link is more likely a malicious link. In one embodiment, link analysis module 109 may access the link to access a remote resource specified by the URI of the link. In the above example, the file of "invoice.zip," and examines at least a portion of the resource to determine whether the resource likely contains malicious content. For example, link analysis module 109 may examine the size of the resource (e.g., whether the size of the file is unusual such as too small or too large), the type of the resource (e.g., whether it is an executable file), metadata of the resource (e.g., file header), and/or the URI itself (e.g., an unusual filename such as having double file extensions, multiple consecutive spaces in the name, etc.). Based on these characteristics of the URL links and in view of link heuristics 111, certain link suspects may be classified as malicious links.

[0027] Once a link is considered as a malicious link, link analysis module 109 generates a new malicious link signature which may be a part of malicious link signatures 112. Thereafter, link analysis module 109 distributes malicious link signatures 112 to malware detection appliances 102-103 via paths 123-124 as part of link signatures 107-108. Note that in addition to

the malicious link detection, the malware detection appliances 102-103 may perform other types of malware detections. For example, a malware detection appliance may perform a dynamic content analysis on the attachment of an email or distribute the attachment to a dedicated malware analysis system for such a dynamic detection. Also note that a dedicated system (not shown) separated from malware analysis system 101 may be utilized to collect the links from malware detection appliances 102-103 and then transmit the collected links to malware analysis system 101 for link analysis.

[0028] Figure 2 is a flow diagram illustrating a method of detecting malicious links according to one embodiment of the invention. Method 200 may be performed by processing logic which may be implemented in software, hardware, or a combination thereof. For example, method 200 may be performed by any of malware detection appliances 102-103 of Figure 1. Referring to Figure 2, at block 201, processing logic analyzes an email to identify a URL link within the email. At block 202, processing logic determines whether the URL link is possible malicious based on malicious link signatures maintained locally. If the URL link is considered as malicious, at block 203, the email and/or the URL link is classified as malicious and an alert may be generated to a user such as an administrator or an email recipient. If the URL link cannot be determined whether it is malicious based on the local malicious link signatures, at block 204, processing logic extracts the URL link from the email and transmits the URL link to a remote malware analysis system over a network for malicious link analysis, without transmitting the remaining content of the email. As described above, the malware analysis system is configured to perform a link analysis on the link and generate a new malicious link signature if the link is determined to be a malicious link using at least some of the techniques described above. At block 205, link classification information, such as malicious indication and/or malicious link signatures, is received from the remote malware analysis system, and at block 206, the local malicious link signatures are updated.

[0029] Figure 3 is a flow diagram illustrating a method of detecting malicious links according to another embodiment of the invention. Method 300 may be performed by processing logic which may be implemented in software, hardware, or a combination thereof. For example, method 300 may be performed by any of malware analysis system 101 of Figure 1. Referring to Figure 3, at block 301, processing logic receives one or more links (e.g., URL links) from one or more malware detection appliances over a network. At least one of the malware detection appliances recognizes a link within an email, extracts the link from the email, and sends the link to a malware analysis system over the network without sending the remain content of the email. At block 302, processing logic performs a link analysis on each of the links

received from the malware detection appliances based on link heuristics to determine whether the link is likely malicious. If it is determined that the link is likely malicious based on the analysis, at block 303, the link is classified as a non-malicious link. Otherwise, at block 304, the link is classified as a malicious link and an alert may be issued to a user (e.g., an administrator or the email's recipient). At block 305, a new malicious link signature(s) is generated based on the newly classified malicious link and at block 306, the new malicious link signature(s) may be distributed to the malware detection appliances over the network for subsequent local malicious link detection.

[0030] Figure 4 is a flow diagram illustrating a method for detecting malicious links according to another embodiment of the invention. Method 400 may be performed by processing logic which may be implemented in software, hardware, or a combination thereof. For example, method 400 may be performed as part of operations involve in block 302 of Figure 3. Referring to Figure 4, at block 401, processing logic determines whether a particular domain name of a link occurs frequently (e.g., above a predetermined threshold) amongst the links received from the malware detection appliances. If so, at block 406, the link or links having that particular domain name may be considered suspicious. At block 403, processing logic determines whether a particular URI of a link occurs frequently (e.g., above a predetermined threshold) amongst the links received from the malware detection appliances. If so, at block 406, the link or links having that particular URI may be considered suspicious. At block 404, processing logic accesses a file or resource via a link and examines at least a portion of the file (e.g., file header, file extension, filename, and/or file size) to determine whether the file contains malicious content. If so, at block 402, the link or links referencing that resource may be classified as malicious links. Otherwise, at block 405, the link or links may be classified as non-malicious.

[0031] The configuration as shown in Figure 1 is described for the illustration purpose only; other configurations or settings may also be applied. For example, according to another embodiment, malware analysis system 101 may include multiple analysis systems or stations, each performing a portion of the overall URL link analysis process, in sequence or in parallel, where the analysis systems may be deployed in the cloud.

[0032] According to one embodiment, a first system performs pre-filtering using whitelists and/or blacklists to pass only specimens requiring further analysis to the second system. The blacklists and whitelists can be updated to the first system aperiodically or periodically (e.g., hourly) from a central or cloud based station. Such a pre-filtering operation can significantly screen out any URL links that are known to be either non-malicious links or malicious links.

[0033] The remaining URL links (e.g., unknown links) may then be processed by a second system. The second system is configured to perform certain link heuristic analysis as described above, for example, across multiple malware detection systems, including those of any vector type (e.g., web, email and file) as well as across multiple ones of any of these. The multiple malware detection systems which may reside at different locations within an enterprise over the cloud. For example, if the same URL is encountered widely (above a threshold count), the URL is deemed suspicious and passed to the third level. In addition, if the URLs are nearly the same but have somewhat different domain names, it may be deemed suspicious, even if the second system finds these URLs within the same enterprise or at different enterprises.

[0034] If a URL link is determined to be suspicious based on the link heuristics, the link is then transmitted to a third system for further analysis. For example, as described above, the third system determines if the URL corresponds to a "live" website and analyzes metadata for the resource at the website (e.g., the size and type of the resource, etc.). If alive and the size of the resource is larger than a predetermined threshold, but is zipped (or otherwise encoded), the zip file may also be parsed to identify anomalies in the headers in determining whether it contains malware.

[0035] Subsequent to these three levels of email malware detection systems, the content/object at the URL (now deemed to be malicious) is processed in a virtual environment provided, e.g., by a web malware detection system (e.g., a fourth system in the cloud), to verify that it is indeed malicious. Verification may require the download of the webpage at the URL. If verified, a signature may be generated and distributed to all the malware detection systems, which signature may contain the URL. Note that a malware detection system may be dedicated to web content or may also be capable of processing emails and web downloads for malware detection.

[0036] Figure 5 is a flow diagram illustrating a method for detecting malicious links according to another embodiment of the invention. Method 500 may be performed by processing logic which may be implemented in software, hardware, or a combination thereof. Referring to Figure 5, at block 501, a first system performs a pre-filtering operation on the URL links that have been collected from multiple malware detection appliances, using a whitelist and/or a blacklist to screen out any known URL links (e.g., known malicious links or known non-malicious links). At block 502, the remaining URL links may be transmitted to a second system to perform a link analysis based on link heuristics to determine whether the URL links are suspicious. At block 503, the suspicious links are analyzed and accessed by a third system to determine whether the links are alive, as well as other metadata of the links (e.g., size and type of

the associated resources). If a URL link is deemed to be alive, at block 504, the live link is processed at a fourth system to download the associated resource and a dynamic analysis is performed on the resource, for example, in a virtual operating environment, to verify that the URL link is indeed malicious. Thereafter, at block 505, the link signatures of the malicious links are generated and distributed to other systems. Note that some of the above operations may be merged and performed within a single system.

[0037] In an alternative embodiment, after the first system performs its pre-filtering, the activities described for the second system and third system are performed in a single malware detecting system, which determines whether the URL is malicious. Afterwards, malware verification may be performed either on the same malware detection system if equipped with a dynamic/virtual analyzer or on another system (e.g., in the cloud). In some implementations of this embodiment, some of the activity of the second system (as described in the preceding paragraph) cannot be performed, particularly those that require cross-enterprise analysis, which is normally performed remotely.

[0038] In view of the techniques described above, embodiments of the invention may use three filters: a first filter called pre-filtering; a second filter called heuristic filtering; and a third filter called object analysis. These three filters can be run in the order described above or in an overlapping/concurrent/parallel manner, though sequencing provides efficiencies. Each filter can be implemented in separate email detection systems, or combined in some fashion in a single or in two separate systems. After these three filtering steps, any detected malware is verified and reported. Verification may be performed in yet another system, a web detection system.

[0039] In a further embodiment, a single system or station may be employed. In this embodiment, a rule-based extractor within the single system is configured to extract one or more URLs from email bodies and attachments. The one or more URLs are extracted and matched against heuristic link signatures to filter the extracted ULR's to a manageable number for further analysis.

[0040] The "heuristic link signatures" used in this matching are generated through automatic and/or manual review of data collected at other network sites (e.g. beta sites) through machine learning techniques. "Machine learning" refers to a process or system that can learn from data, i.e., be trained to distinguish between malicious and non-malicious, and classify samples under test accordingly or develop indicators having a high correlation to those that are malicious. The core principals of machine learning deal with representation and generalization, that is, representation of data instances (e.g., reputation or anomaly information), and functions performed on those instances (e.g., weighting and scoring). The data used for weighting and/or

scoring the likelihood of samples under test including malware may be updated from time to time, whether on an aperiodic or periodic basis. The heuristic link signatures are validated to determine what is malicious and what is not and the most popular patterns are converted into rules. For example, patterns that appear more often, i.e., greater than a predetermined threshold, in the previously determined suspicious/malicious links or non-malicious links may be converted into rules. Patterns may include certain characters or sequence of characters in a link, such as, for example, domain names or URIs, etc. These heuristic link signatures also filter out the "auto opt-in" and "auto opt-out" links contained within emails. The heuristic link signatures may be generated, for example, at a remote or central facility, coupled to the single system and possibly the other standalone malware detection systems for distribution (and updating) of the heuristic link signatures from time to time.

[0041] The heuristics represented by the heuristic link signatures may be the same as described hereinabove for the other embodiments of the invention. For example, if the frequency of encountering a link specifying a particular IP address or particular domain is above a set threshold, the heuristic may assign to the link a higher probability that the link is associated with malware, and should be classified as at least suspicious and requiring further analysis.

[0042] In the further analysis, called the "dynamic" phase, the suspicious link may be analyzed in a virtual machine to determine whether it is actually malicious. For this, the single system is provided with a network interface and network (e.g., Internet) access to permit retrieval (e.g., download) of the content of the website referenced by the URL. The system may also configure the virtual machine with one or more computer programs or applications (e.g., PDF reader or media player) suitable for the particular type or types of content (e.g., PDF document or video file, respectively) retrieved from the website. The behavior of the content within the virtual machine during execution (e.g., loading and processing of the content within an application) is monitored and unexpected or anomalous behavior (spawning processes, attempted unauthorized communications, etc.) is evaluated. Based on the behavior, the content is classified or verified as malicious or non-malicious.

[0043] In one embodiment, if an email/attachment contains a suspicious URL, the system "delays" delivery of the email for a time period up to a maximum length of time (e.g., half a minute). This may require the system to pull the email from the wire, store the email, and effectively "resend" the email if no malware is found or send an alert to the recipient that a malicious email was detected and blocked. Additionally, besides matching against heuristic link signatures, the single system is configured to apply static analysis techniques to a header or other metadata of the email message to aid in determining whether the email or its attachments are

malicious. For instance, according to one embodiment of the disclosure, the email (SMTP) header is extracted and is sent to one or more static analysis modules to detect whether information within the email header is malicious or not. As an example, known malicious email headers or other metadata are analyzed to determine if the malicious content includes a pattern. The pattern is converted to a software rule, where email headers that match any of these software rules will be classified as malicious. The email rules may specify that the content of certain fields of the header (e.g., FROM, TO, or SUBJECT are associated with malware.

[0044] The above process can be implemented in real time and may use adaptable filters (updated from time to time), e.g., with changeable thresholds, may employ intelligence in the cloud to update based on experience of other malware detection systems, and/or based on machine learning or experiential techniques. On the other hand, use of dependable heuristics permits local systems to perform most processing, thus avoiding delays/latency associated with communication into the cloud, and avoid sending sensitive information "off network" into the cloud. Other configurations may also be implemented.

[0045] Figure 6 is a block diagram of an illustrative computer network system 800 having a malicious content detection system 850 in accordance with a further illustrative embodiment. The malicious content detection system 850 may represent any of the malicious content detection systems or malicious content analysis systems described above, such as, for example, systems 101-103 of Figure 1. In one embodiment, the malicious content detection system 850 includes link analysis module 109 to detect and determine whether a particular URL link is malicious using some or all of the techniques described above.

[0046] The malicious content detection system 850 is illustrated with a server device 810 and a client device 830, each coupled for communication via a communication network 820. In various embodiments, there may be multiple server devices and multiple client devices sending and receiving data to/from each other, and the same device can serve as either a server or a client in separate communication sessions. Although Figure 6 depicts data transmitted from the server device 810 to the client device 830, either device can transmit and receive data from the other.

[0047] Note that throughout this application, network content is utilized as an example of content for malicious content detection purposes; however, other types of content can also be applied. Network content may include any data transmitted over a network (i.e., network data). Network data may include text, software, images, audio, or other digital data. An example of network content includes web content, or any network data that may be transmitted using a Hypertext Transfer Protocol (HTTP), Hypertext Markup Language (HTML) protocol, or be transmitted in a manner suitable for display on a Web browser software application. Another

example of network content includes email messages, which may be transmitted using an email protocol such as Simple Mail Transfer Protocol (SMTP), Post Office Protocol version 3 (POP3), or Internet Message Access Protocol (IMAP4). A further example of network content includes Instant Messages, which may be transmitted using an Instant Messaging protocol such as Session Initiation Protocol (SIP) or Extensible Messaging and Presence Protocol (XMPP). In addition, network content may include any network data that is transferred using other data transfer protocols, such as File Transfer Protocol (FTP).

[0048] The malicious network content detection system 850 may monitor exchanges of network content (e.g., Web content) in real-time rather than intercepting and holding the network content until such time as it can determine whether the network content includes malicious network content. The malicious network content detection system 850 may be configured to inspect exchanges of network content over the communication network 820, identify suspicious network content, and analyze the suspicious network content using a virtual machine to detect malicious network content. In this way, the malicious network content detection system 850 may be computationally efficient and scalable as data traffic volume and the number of computing devices communicating over the communication network 820 increase. Therefore, the malicious network content detection system 850 may not become a bottleneck in the computer network system 800.

[0049] The communication network 820 may include a public computer network such as the Internet, in which case a firewall 825 may be interposed between the communication network 820 and the client device 830. Alternatively, the communication network may be a private computer network such as a wireless telecommunication network, wide area network, or local area network, or a combination of networks. Though the communication network 820 may include any type of network and be used to communicate different types of data, communications of web data may be discussed below for purposes of example.

[0050] The malicious network content detection system 850 is shown as coupled with the network 820 by a network tap 840 (e.g., a data/packet capturing device). The network tap 840 may include a digital network tap configured to monitor network data and provide a copy of the network data to the malicious network content detection system 850. Network data may comprise signals and data that are transmitted over the communication network 820 including data flows from the server device 810 to the client device 830. In one example, the network tap 840 monitors and copies the network data without an appreciable decline in performance of the server device 810, the client device 830, or the communication network 820. The network tap 840 may copy any portion of the network data, for example, any number of data packets from the

network data. In embodiments where the malicious content detection system 850 is implemented as an dedicated appliance or a dedicated computer system, the network tap 840 may include an assembly integrated into the appliance or computer system that includes network ports, network interface card and related logic (not shown) for connecting to the communication network 820 to non-disruptively "tap" traffic thereon and provide a copy of the traffic to the heuristic module 860. In other embodiments, the network tap 840 can be integrated into a firewall, router, switch or other network device (not shown) or can be a standalone component, such as an appropriate commercially available network tap. In virtual environments, a virtual tap (vTAP) can be used to copy traffic from virtual networks.

[0051] The network tap 840 may also capture metadata from the network data. The metadata may be associated with the server device 810 and/or the client device 830. For example, the metadata may identify the server device 810 and/or the client device 830. In some embodiments, the server device 810 transmits metadata which is captured by the tap 840. In other embodiments, a heuristic module 860 (described herein) may determine the server device 810 and the client device 830 by analyzing data packets within the network data in order to generate the metadata. The term, "content," as used herein may be construed to include the intercepted network data and/or the metadata unless the context requires otherwise.

[0052] The malicious network content detection system 850 may include a heuristic module 860, a heuristics database 862, a scheduler 870, a virtual machine pool 880, an analysis engine 882 and a reporting module 884. In some embodiments, the network tap 840 may be contained within the malicious network content detection system 850.

[0053] The heuristic module 860 receives the copy of the network data from the network tap 840 and applies heuristics to the data to determine if the network data might contain suspicious network content. The heuristics applied by the heuristic module 860 may be based on data and/or rules stored in the heuristics database 862. The heuristic module 860 may examine the image of the captured content without executing or opening the captured content. For example, the heuristic module 860 may examine the metadata or attributes of the captured content and/or the code image (e.g., a binary image of an executable) to determine whether a certain portion of the captured content matches a predetermined pattern or signature that is associated with a particular type of malicious content. In one example, the heuristic module 860 flags network data as suspicious after applying a heuristic analysis. This detection process is also referred to as a static malicious content detection. The suspicious network data may then be provided to the scheduler 870. In some embodiments, the suspicious network data is provided directly to the scheduler 870 with or without buffering or organizing one or more data flows.

[0054] When a characteristic of the packet, such as a sequence of characters or keyword, is identified that meets the conditions of a heuristic, a suspicious characteristic of the network content is identified. The identified characteristic may be stored for reference and analysis. In some embodiments, the entire packet may be inspected (e.g., using deep packet inspection techniques) and multiple characteristics may be identified before proceeding to the next step. In some embodiments, the characteristic may be determined as a result of an analysis across multiple packets comprising the network content. A score related to a probability that the suspicious characteristic identified indicates malicious network content is determined.

[0055] The heuristic module 860 may also provide a priority level for the packet and/or the features present in the packet. The scheduler 870 may then load and configure a virtual machine from the virtual machine pool 880 in an order related to the priority level, and dispatch the virtual machine to the analysis engine 882 to process the suspicious network content.

[0056] The heuristic module 860 may provide the packet containing the suspicious network content to the scheduler 870, along with a list of the features present in the packet and the malicious probability scores associated with each of those features. Alternatively, the heuristic module 860 may provide a pointer to the packet containing the suspicious network content to the scheduler 870 such that the scheduler 870 may access the packet via a memory shared with the heuristic module 860. In another embodiment, the heuristic module 860 may provide identification information regarding the packet to the scheduler 870 such that the scheduler 870, or virtual machine may query the heuristic module 860 for data regarding the packet as needed.

[0057] The scheduler 870 may identify the client device 830 and retrieve a virtual machine associated with the client device 830. A virtual machine may itself be executable software that is configured to mimic the performance of a device (e.g., the client device 830). The virtual machine may be retrieved from the virtual machine pool 880. Furthermore, the scheduler 870 may identify, for example, a Web browser running on the client device 830, and retrieve a virtual machine associated with the web browser.

[0058] In some embodiments, the heuristic module 860 transmits the metadata identifying the client device 830 to the scheduler 870. In other embodiments, the scheduler 870 receives one or more data packets of the network data from the heuristic module 860 and analyzes the one or more data packets to identify the client device 830. In yet other embodiments, the metadata may be received from the network tap 840.

[0059] The scheduler 870 may retrieve and configure the virtual machine to mimic the pertinent performance characteristics of the client device 830. In one example, the scheduler 870 configures the characteristics of the virtual machine to mimic only those features of the client

device 830 that are affected by the network data copied by the network tap 840. The scheduler 870 may determine the features of the client device 830 that are affected by the network data by receiving and analyzing the network data from the network tap 840. Such features of the client device 830 may include ports that are to receive the network data, select device drivers that are to respond to the network data, and any other devices coupled to or contained within the client device 830 that can respond to the network data. In other embodiments, the heuristic module 860 may determine the features of the client device 830 that are affected by the network data by receiving and analyzing the network data from the network tap 840. The heuristic module 860 may then transmit the features of the client device to the scheduler 870.

[0060] The virtual machine pool 880 may be configured to store one or more virtual machines. The virtual machine pool 880 may include software and/or a storage medium capable of storing software. In one example, the virtual machine pool 880 stores a single virtual machine that can be configured by the scheduler 870 to mimic the performance of any client device 830 on the communication network 820. The virtual machine pool 880 may store any number of distinct virtual machines that can be configured to simulate the performance of a wide variety of client devices 830.

[0061] The analysis engine 882 simulates the receipt and/or display of the network content from the server device 810 after the network content is received by the client device 110 to analyze the effects of the network content upon the client device 830. The analysis engine 882 may identify the effects of malware or malicious network content by analyzing the simulation of the effects of the network content upon the client device 830 that is carried out on the virtual machine. There may be multiple analysis engines 882 to simulate multiple streams of network content. The analysis engine 882 may be configured to monitor the virtual machine for indications that the suspicious network content is in fact malicious network content. Such indications may include unusual network transmissions, unusual changes in performance, and the like. This detection process is referred to as a dynamic malicious content detection.

[0062] The analysis engine 882 may flag the suspicious network content as malicious network content according to the observed behavior of the virtual machine. The reporting module 884 may issue alerts indicating the presence of malware, and using pointers and other reference information, identify the packets of the network content containing the malware. Additionally, the server device 810 may be added to a list of malicious network content providers, and future network transmissions originating from the server device 810 may be blocked from reaching their intended destinations, e.g., by firewall 825.

[0063] The computer network system 800 may also include a further communication network 890, which couples the malicious content detection system (MCDS) 850 with one or more other MCDS, of which MCDS 892 and MCDS 894 are shown, and a management system 896, which may be implemented as a Web server having a Web interface. The communication network 890 may, in some embodiments, be coupled for communication with or part of network 820. The management system 896 is responsible for managing the MCDS 850, 892, 894 and providing updates to their operation systems and software programs. Also, the management system 896 may cause malware signatures generated by any of the MCDS 850, 892, 894 to be shared with one or more of the other MCDS 850, 892, 894, for example, on a subscription basis. Moreover, the malicious content detection system as described in the foregoing embodiments may be incorporated into one or more of the MCDS 850, 892, 894, or into all of them, depending on the deployment. Also, the management system 896 itself or another dedicated computer station may incorporate the malicious content detection system in deployments where such detection is to be conducted at a centralized resource.

[0064] Further information regarding an embodiment of a malicious content detection system can be had with reference to U.S. Patent No. 8,171,553, the disclosure of which being incorporated herein by reference in its entirety.

[0065] As described above, the detection or analysis performed by the heuristic module 860 may be referred to as static detection or static analysis, which may generate a first score (e.g., a static detection score) according to a first scoring scheme or algorithm. The detection or analysis performed by the analysis engine 882 is referred to as dynamic detection or dynamic analysis, which may generate a second score (e.g., a dynamic detection score) according to a second scoring scheme or algorithm. The first and second scores may be combined, according to a predetermined algorithm, to derive a final score indicating the probability that a malicious content suspect is indeed malicious.

[0066] Furthermore, detection systems 850 and 892-894 may deployed in a variety of distribution ways. For example, detection system 850 may be deployed as a detection appliance at a client site to detect any suspicious content, for example, at a local area network (LAN) of the client. In addition, any of MCDS 892 and MCDS 894 may also be deployed as dedicated data analysis systems. Systems 850 and 892-894 may be configured and managed by a management system 896 over network 890, which may be a LAN, a wide area network (WAN) such as the Internet, or a combination of both. Management system 896 may be implemented as a Web server having a Web interface to allow an administrator of a client (e.g., corporation entity) to log in to manage detection systems 850 and 892-894. For example, an administrator may able to

activate or deactivate certain functionalities of malicious content detection systems 850 and 892-894 or alternatively, to distribute software updates such as malicious content definition files (e.g., malicious signatures or patterns) or rules, etc. Furthermore, a user can submit via a Web interface suspicious content to be analyzed, for example, by dedicated data analysis systems 892-894. As described above, malicious content detection includes static detection and dynamic detection. Such static and dynamic detections can be distributed amongst different systems over a network. For example, static detection may be performed by detection system 850 at a client site, while dynamic detection of the same content can be offloaded to the cloud, for example, by any of detection systems 892-894. Other configurations may exist.

[0067] Figure 7 is a block diagram illustrating an example of a data processing system which may be used with one embodiment of the invention. For example, system 900 may represents any of data processing systems described above performing any of the processes or methods described above. System 900 may represent a desktop, a tablet, a server, a mobile phone, a media player, a personal digital assistant (PDA), a personal communicator, a gaming device, a network router or hub, a wireless access point (AP) or repeater, a set-top box, or a combination thereof.

[0068] Referring to Figure 7, in one embodiment, system 900 includes processor 901 and peripheral interface 902, also referred to herein as a chipset, to couple various components to processor 901 including memory 903 and devices 905-908 via a bus or an interconnect. Processor 901 may represent a single processor or multiple processors with a single processor core or multiple processor cores included therein. Processor 901 may represent one or more general-purpose processors such as a microprocessor, a central processing unit (CPU), or the like. More particularly, processor 901 may be a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or processor implementing other instruction sets, or processors implementing a combination of instruction sets. Processor 901 may also be one or more special-purpose processors such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), a network processor, a graphics processor, a network processor, a communications processor, a cryptographic processor, a co-processor, an embedded processor, or any other type of logic capable of processing instructions. Processor 901 is configured to execute instructions for performing the operations and steps discussed herein.

[0069] Peripheral interface 902 may include memory control hub (MCH) and input output control hub (ICH). Peripheral interface 902 may include a memory controller (not shown) that

communicates with a memory 903. Peripheral interface 902 may also include a graphics interface that communicates with graphics subsystem 904, which may include a display controller and/or a display device. Peripheral interface 902 may communicate with graphics device 904 via an accelerated graphics port (AGP), a peripheral component interconnect (PCI) express bus, or other types of interconnects.

[0070] An MCH is sometimes referred to as a Northbridge and an ICH is sometimes referred to as a Southbridge. As used herein, the terms MCH, ICH, Northbridge and Southbridge are intended to be interpreted broadly to cover various chips who functions include passing interrupt signals toward a processor. In some embodiments, the MCH may be integrated with processor 901. In such a configuration, peripheral interface 902 operates as an interface chip performing some functions of the MCH and ICH. Furthermore, a graphics accelerator may be integrated within the MCH or processor 901.

[0071] Memory 903 may include one or more volatile storage (or memory) devices such as random access memory (RAM), dynamic RAM (DRAM), synchronous DRAM (SDRAM), static RAM (SRAM), or other types of storage devices. Memory 903 may store information including sequences of instructions that are executed by processor 901, or any other device. For example, executable code and/or data of a variety of operating systems, device drivers, firmware (e.g., input output basic system or BIOS), and/or applications can be loaded in memory 903 and executed by processor 901. An operating system can be any kind of operating systems, such as, for example, Windows® operating system from Microsoft®, Mac OS®/iOS® from Apple, Android® from Google®, Linux®, Unix®, or other real-time or embedded operating systems such as VxWorks.

[0072] Peripheral interface 902 may provide an interface to IO devices such as devices 905-908, including wireless transceiver(s) 905, input device(s) 906, audio IO device(s) 907, and other IO devices 908. Wireless transceiver 905 may be a WiFi transceiver, an infrared transceiver, a Bluetooth transceiver, a WiMax transceiver, a wireless cellular telephony transceiver, a satellite transceiver (e.g., a global positioning system (GPS) transceiver) or a combination thereof. Input device(s) 906 may include a mouse, a touch pad, a touch sensitive screen (which may be integrated with display device 904), a pointer device such as a stylus, and/or a keyboard (e.g., physical keyboard or a virtual keyboard displayed as part of a touch sensitive screen). For example, input device 906 may include a touch screen controller coupled to a touch screen. The touch screen and touch screen controller can, for example, detect contact and movement or break thereof using any of a plurality of touch sensitivity technologies, including but not limited to capacitive, resistive, infrared, and surface acoustic wave technologies, as well as other proximity

sensor arrays or other elements for determining one or more points of contact with the touch screen.

[0073] Audio IO 907 may include a speaker and/or a microphone to facilitate voice-enabled functions, such as voice recognition, voice replication, digital recording, and/or telephony functions. Other optional devices 908 may include a storage device (e.g., a hard drive, a flash memory device), universal serial bus (USB) port(s), parallel port(s), serial port(s), a printer, a network interface, a bus bridge (e.g., a PCI-PCI bridge), sensor(s) (e.g., a motion sensor, a light sensor, a proximity sensor, etc.), or a combination thereof. Optional devices 908 may further include an imaging processing subsystem (e.g., a camera), which may include an optical sensor, such as a charged coupled device (CCD) or a complementary metal-oxide semiconductor (CMOS) optical sensor, utilized to facilitate camera functions, such as recording photographs and video clips.

[0074] Note that while Figure 7 illustrates various components of a data processing system, it is not intended to represent any particular architecture or manner of interconnecting the components; as such details are not germane to embodiments of the present invention. It will also be appreciated that network computers, handheld computers, mobile phones, and other data processing systems which have fewer components or perhaps more components may also be used with embodiments of the invention.

[0075] Some portions of the preceding detailed descriptions have been presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the ways used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities.

[0076] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the above discussion, it is appreciated that throughout the description, discussions utilizing terms such as those set forth in the claims below, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0077] The techniques shown in the figures can be implemented using code and data stored and executed on one or more electronic devices. Such electronic devices store and communicate (internally and/or with other electronic devices over a network) code and data using computer-readable media, such as non-transitory computer-readable storage media (e.g., magnetic disks; optical disks; random access memory; read only memory; flash memory devices; phase-change memory) and transitory computer-readable transmission media (e.g., electrical, optical, acoustical or other form of propagated signals – such as carrier waves, infrared signals, digital signals).

[0078] The processes or methods depicted in the preceding figures may be performed by processing logic that comprises hardware (e.g. circuitry, dedicated logic, etc.), firmware, software (e.g., embodied on a non-transitory computer readable medium), or a combination of both. Although the processes or methods are described above in terms of some sequential operations, it should be appreciated that some of the operations described may be performed in a different order. Moreover, some operations may be performed in parallel rather than sequentially.

[0079] In the foregoing specification, embodiments of the invention have been described with reference to specific exemplary embodiments thereof. It will be evident that various modifications may be made thereto without departing from the broader spirit and scope of the invention as set forth in the following claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A computer-implemented method for detecting malicious links in electronic messages, comprising:

in response to receiving a plurality of uniform resource locator (URL) links for malicious determination, removing any known URL links from the URL links based on a list of known link signatures;

for each of remaining URL links that are unknown, performing a link analysis on the URL link based on link heuristics to determine whether the URL link is suspicious;

for each of the suspicious URL links, performing a dynamic analysis on a resource of the suspicious URL link; and

classifying whether the suspicious URL link is a malicious link based on a behavior of the resource during the dynamic analysis.

2. The method of claim 1, wherein the URL links are received from one or more malware detection appliances over the Internet, and wherein the one or more malware detection appliances are deployed in one or more local area networks (LANs).

3. The method of claim 2, wherein each of the malware detection appliances is configured to

recognize a URL link embedded within an email,

extract the URL from the email, and

transmit the extracted URL to the malware analysis system over the Internet as part of the plurality of URL links, without transmitting remaining content of the email.

4. The method of claim 2, wherein performing a link analysis comprises:

determining whether a first domain name of the URL link occurs frequently amongst the plurality of URL links; and

designating each of the URL links that includes the first domain name as suspicious, if the first domain name occurs frequently amongst the URL links.

5. The method of claim 2, wherein performing a link analysis comprises:

determining whether a first uniform resource identifier (URI) of the URL link occurs frequently amongst the plurality of URL links; and

designating each of the URL links that includes the first URI as suspicious, if the first URI occurs frequently amongst the URL links.

6. The method of claim 2, wherein performing a link analysis comprises:

accessing a remote resource of a remote resource location via the URL link; and

classifying whether the URL link is suspicious based on a response received from the remote resource location.

7. The method of claim 6, wherein the URL link is classified as suspicious based on a combination of a type and a size of the remote resource.

8. The method of claim 1, further comprising:

generating a link signature for each of the links that have been classified as malicious; and

transmitting link signatures of the classified malicious links to the one or more malware detection appliances, such that the malware detection appliances can detect subsequent malicious links locally based on the link signatures.

9. A system comprising:

a memory; and

a processor coupled to the memory, the processor to (i) determine one or more uniform resource locator (URL) links from a plurality of URL links received for malicious determination that are absent from a list of known link signatures, (ii) perform a link analysis on each of the one or more URL links based on link heuristics to determine whether any of the one or more URL links is suspicious, (iii) for each of the suspicious URL links, perform a dynamic analysis on a resource of the suspicious URL link, and (iv) classify whether the suspicious URL link is a malicious link based on a behavior of the resource during the dynamic analysis.

10. The system of claim 9, wherein the plurality of URL links are received from one or more malware detection appliances over the Internet, and wherein the one or more malware detection appliances are deployed in one or more local area networks (LANs).

11. The system of claim 10, wherein each of the malware detection appliances is configured to

recognize a URL link embedded within an email,

extract the URL from the email, and

transmit the extracted URL to the malware analysis system over the Internet as part of the plurality of URL links, without transmitting remaining content of the email.

12. The system of claim 10, wherein the processor to perform the link analysis by

determining whether a first domain name for each of the one or more URL links occurs frequently amongst the plurality of URL links; and

designating each of the one or more URL links that includes the first domain name as suspicious, if the first domain name occurs frequently amongst the plurality of URL links.

13. The system of claim 10, wherein the processor to perform the link analysis by

determining whether a first uniform resource identifier (URI) of the one or more URL links occurs frequently amongst the plurality of URL links; and

designating each of the one or more URL links that includes the first URI as suspicious, if the first URI occurs frequently amongst the plurality of URL links.

14. The system of claim 10, wherein the processor to perform the link by

accessing a remote resource of a remote resource location via a URL link from the one or more URL links; and

classifying whether the URL link is suspicious based on a response received from the remote resource location.

15. The system of claim 14, wherein the processor to classify the URL link is classified as suspicious based on a combination of a type and a size of the remote resource.

16. The system of claim 9, wherein the processor further (i) generates a link signature for each of the suspicious URL links that have been classified as malicious, and (ii) transmits link signatures of the classified malicious links to the one or more malware detection appliances, such that the malware detection appliances can detect subsequent malicious links locally based on the link signatures.

17. A computer-implemented method for detecting malicious links, comprising:
in response to receiving an email having a uniform resource locator (URL) link for malicious determination, extracting the URL link from the email;
comparing at least a portion of the extracted URL link with a list of heuristic link signatures that represents a list of patterns to determine whether the URL link is suspicious;
performing a dynamic analysis on the extracted URL link in a virtual machine (VM) if at least a portion of the extracted URL link matches at least one of the heuristic link signatures, including
accessing and downloading a resource from a remote site via the extracted URL link,
executing the resource within the VM using a software program that is associated with the resource, and
monitoring a behavior of the resource within the VM; and
classifying whether the extracted URL link is a malicious link based on the behavior of the resource during the execution of the resource within the VM.

18. The method of claim 17, wherein the heuristic link signatures are generated based on data previously collected from a plurality of malware detection systems using machine learning techniques.

19. The method of claim 18, wherein the heuristic link signatures are updated periodically based on data periodically collected from the malware detection systems.

20. The method of claim 18, wherein the heuristic link signatures represent patterns of likely malicious links based on previous malware detection operations.

21. The method of claim 17, further comprising performing a static analysis on metadata of the email to determine whether the email is a suspicious email, wherein the dynamic analysis is performed on the extracted URL link if the email is determined to be suspicious.

22. The method of claim 21, wherein the metadata of the email includes at least one of TO, FROM, and SUBJECT fields.

23. A non-transitory machine-readable medium having instructions stored therein, which when executed by a processor, cause the processor to perform a method for detecting malicious links, the method comprising:

in response to receiving an email having a uniform resource locator (URL) link for malicious determination, extracting the URL link from the email;

comparing at least a portion of the extracted URL link with a list of heuristic link signatures that represents a list of patterns to determine whether the URL link is suspicious;

performing a dynamic analysis on the extracted URL link in a virtual machine (VM) if at least a portion of the extracted URL link matches at least one of the heuristic link signatures, including

accessing and downloading a resource from a remote site via the extracted URL link,

executing the resource within the VM using a software program that is associated with the resource, and

monitoring a behavior of the resource within the VM; and

classifying whether the extracted URL link is a malicious link based on the behavior of the resource during the execution of the resource within the VM.

24. The medium of claim 23, wherein the heuristic link signatures are generated based on data previously collected from a plurality of malware detection systems using machine learning techniques.

25. The medium of claim 24, wherein the heuristic link signatures are updated periodically based on data periodically collected from the malware detection systems.

26. The medium of claim 24, wherein the heuristic link signatures represent patterns of likely malicious links based on previous malware detection operations.

27. The medium of claim 23, wherein the method further comprises performing a static analysis on metadata of the email to determine whether the email is a suspicious email, wherein the dynamic analysis is performed on the extracted URL link if the email is determined to be suspicious.

28. A data processing system for detecting malicious links, comprising:
a processor; and
a memory coupled to the processor for storing instructions, which when executed from the memory by the processor, cause the processor to
in response to receiving an email having a uniform resource locator (URL) link for malicious determination, extract the URL link from the email,
compare at least a portion of the extracted URL link with a list of heuristic link signatures that represents a list of patterns to determine whether the URL link is suspicious,
perform a dynamic analysis on the extracted URL link in a virtual machine (VM) if at least a portion of the extracted URL link matches at least one of the heuristic link signatures, including
accessing and downloading a resource from a remote site via the extracted URL link,
executing the resource within the VM using a software program that is associated with the resource, and
monitoring a behavior of the resource within the VM, and
classify whether the extracted URL link is a malicious link based on the behavior of the resource during the execution of the resource within the VM.

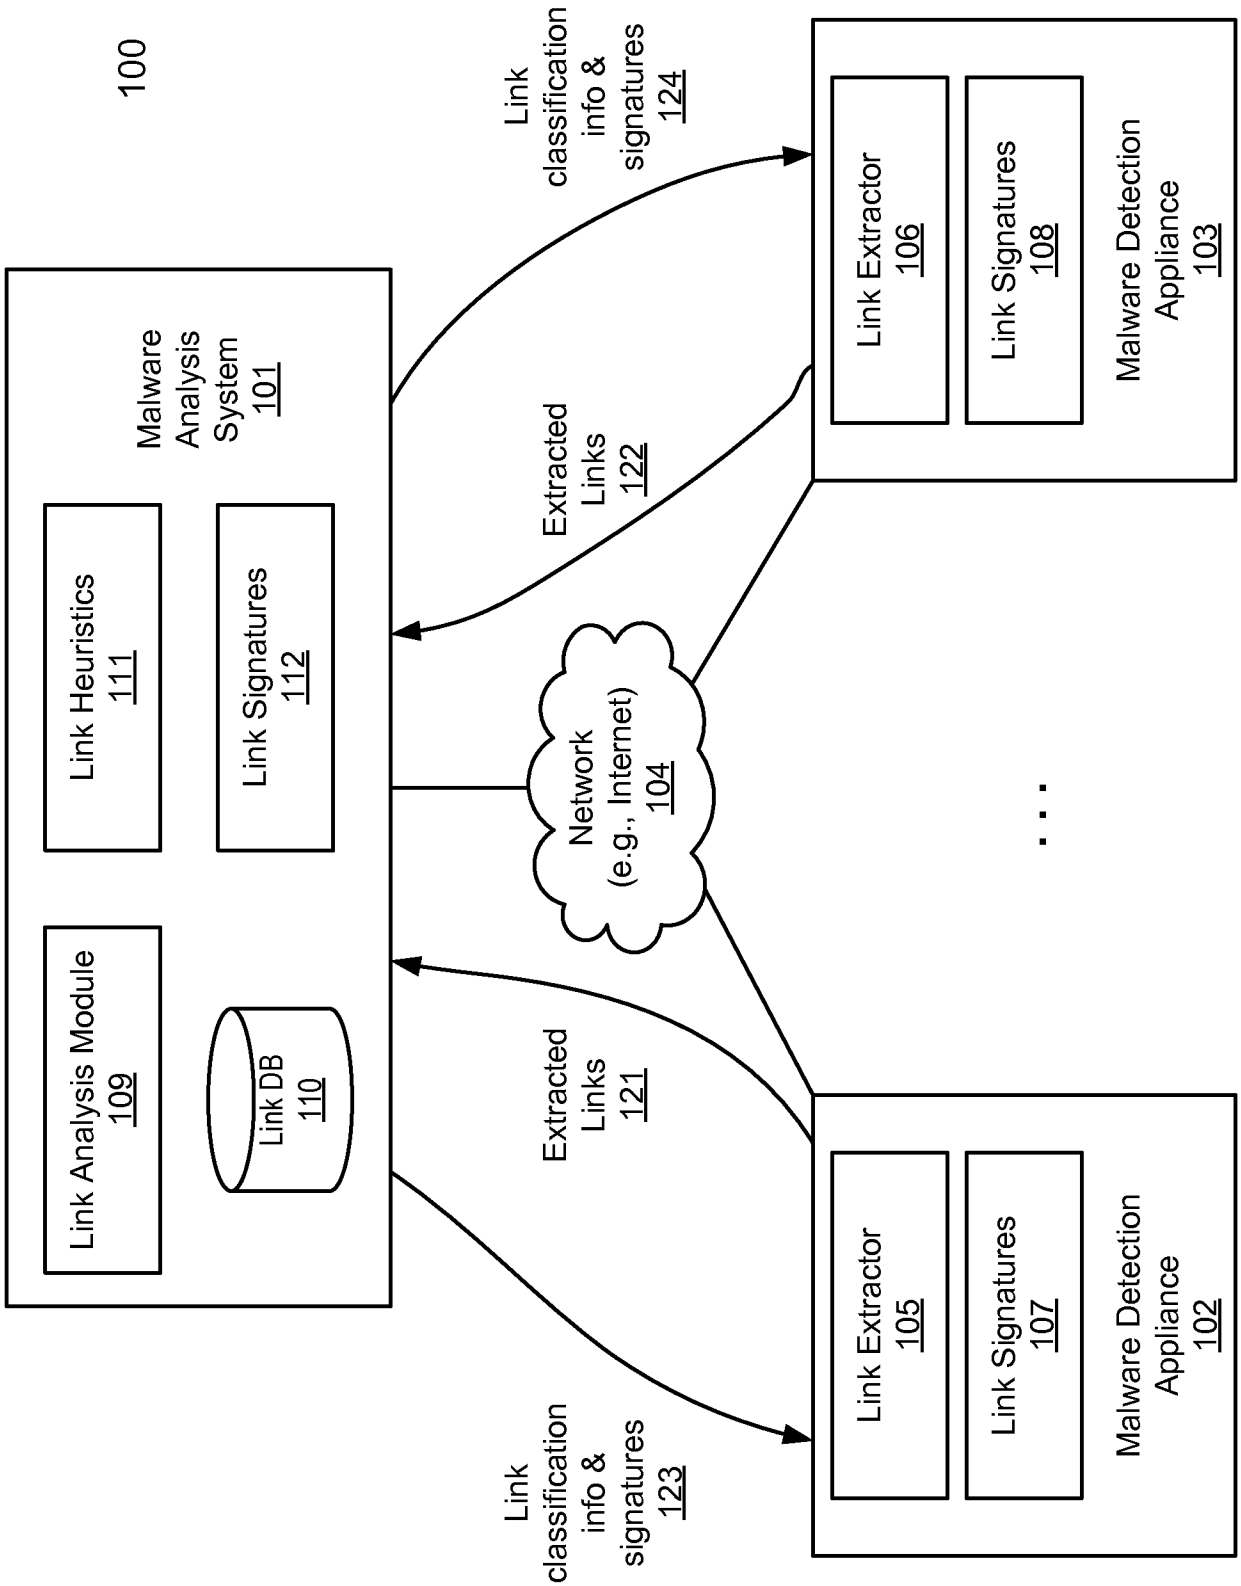
100
Malware Analysis System 101
Link Heuristics 111
Link Signatures 112
Link Analysis Module 109
Link DB 110
Network (e.g., Internet) 104
Link classification info & signatures 124
Extracted Links 122
Extracted Links 121
Link classification info & signatures 123
Link Extractor 106
Link Signatures 108
Malware Detection Appliance 103
Link Extractor 105
Link Signatures 107
Malware Detection Appliance 102
. . .

FIG. 1

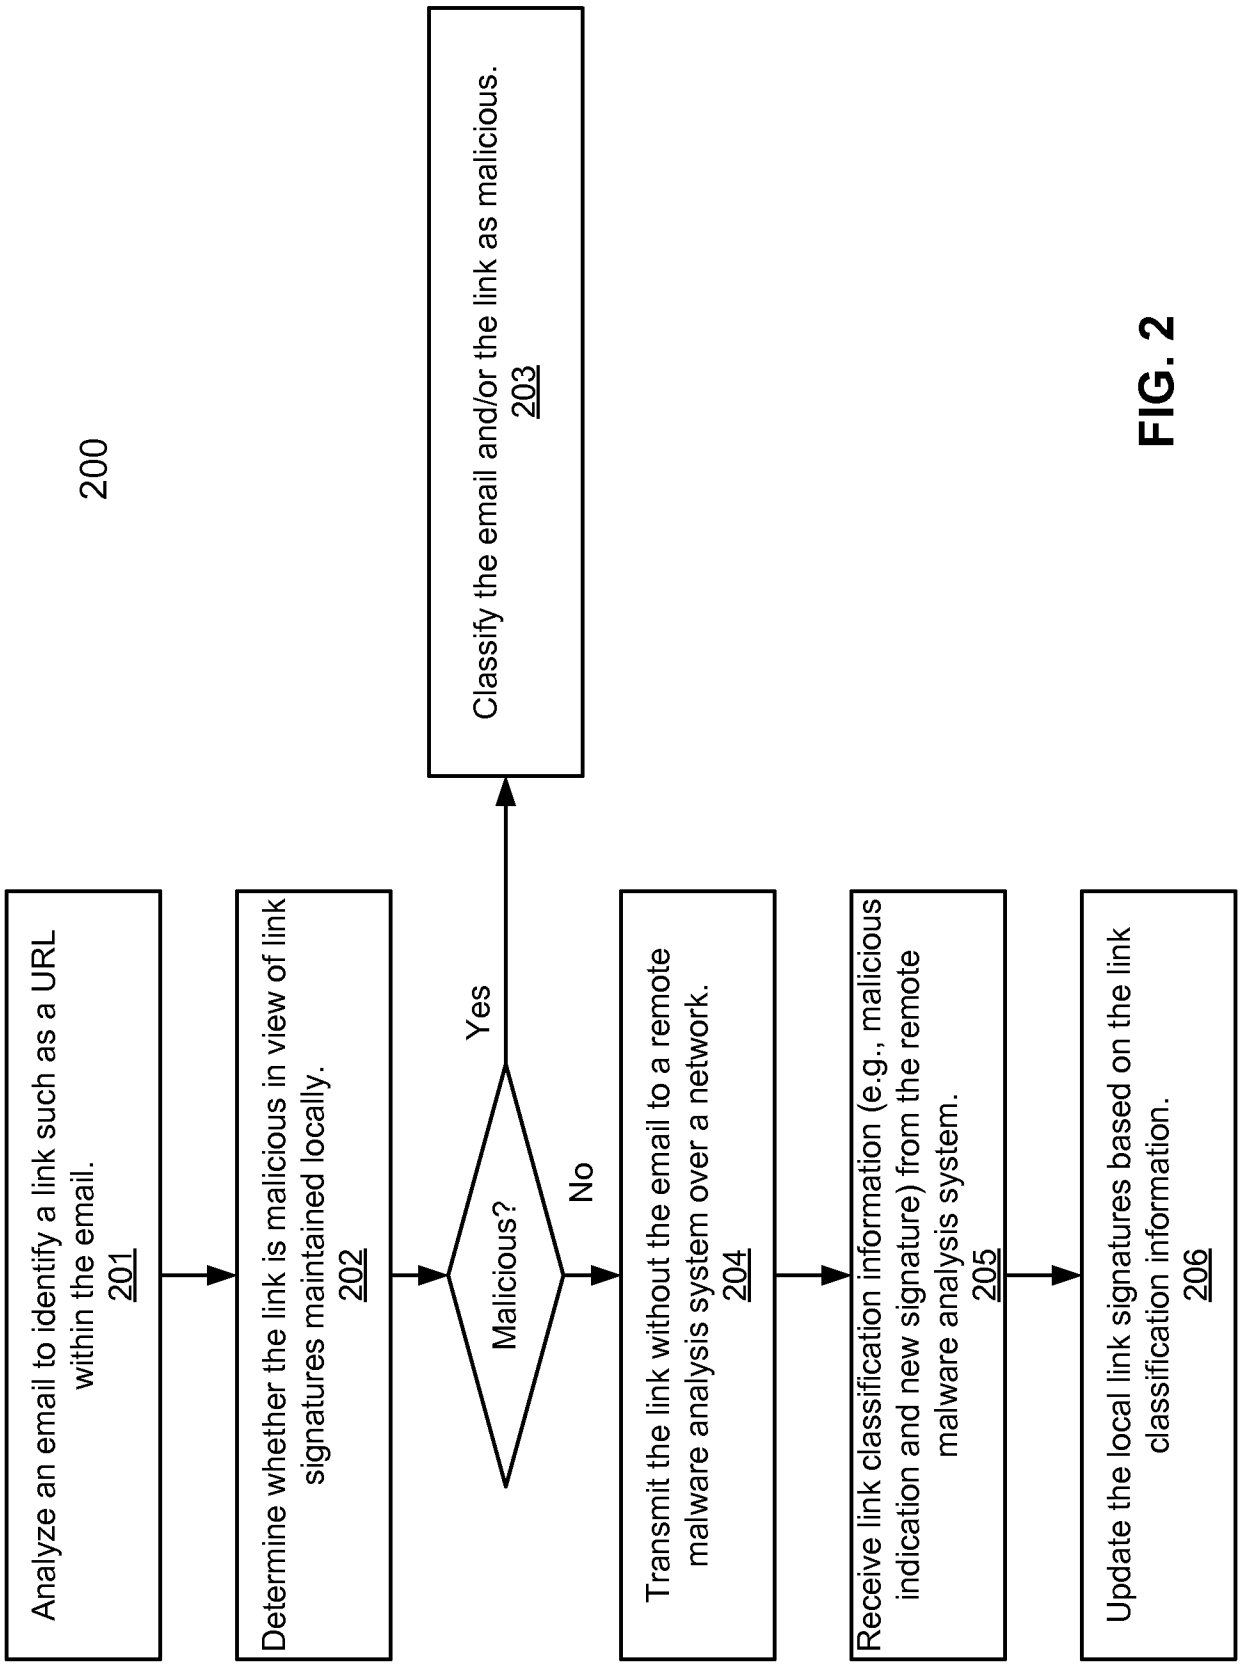
200
Analyze an email to identify a link such as a URL within the email.
201
Determine whether the link is malicious in view of link signatures maintained locally.
202
Malicious?
Yes
No
Classify the email and/or the link as malicious.
203
Transmit the link without the email to a remote malware analysis system over a network.
204
Receive link classification information (e.g., malicious indication and new signature) from the remote malware analysis system.
205
Update the local link signatures based on the link classification information.
206

FIG. 2

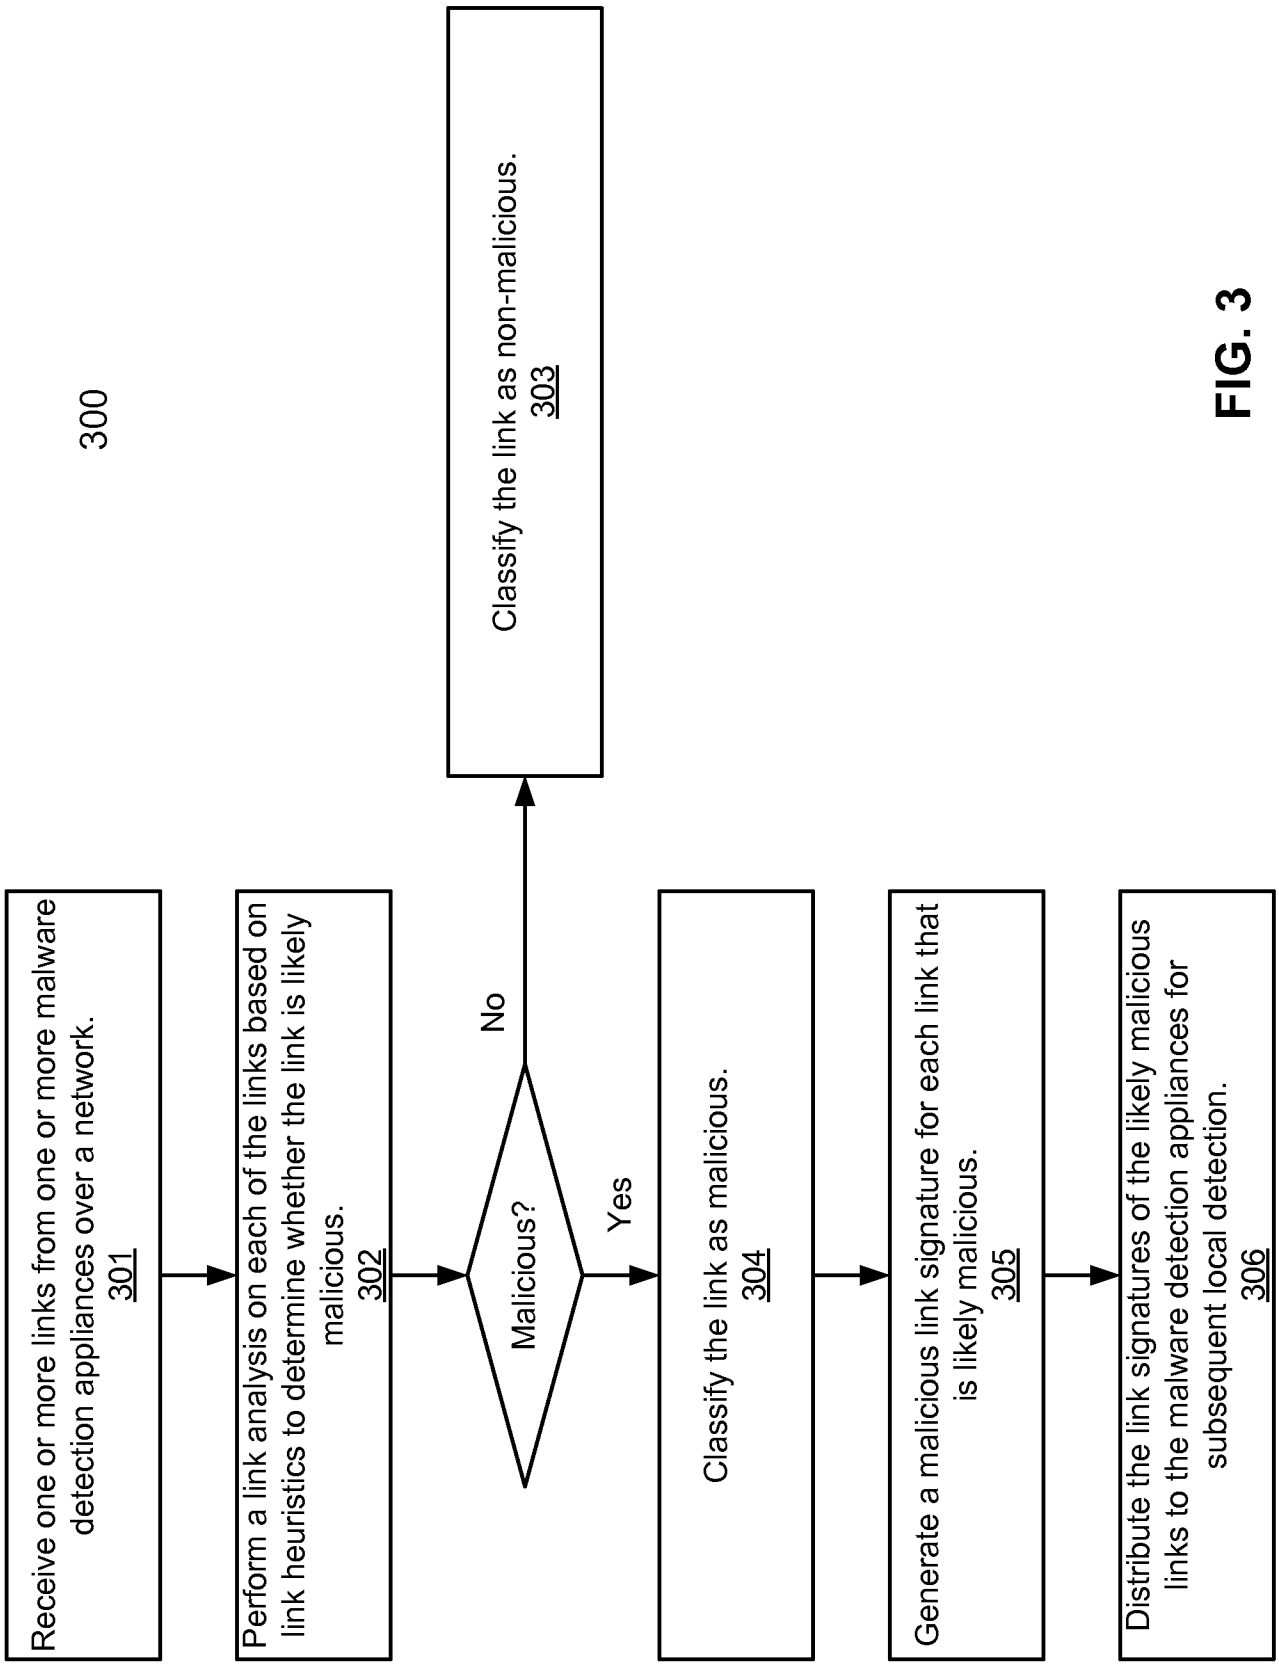
300
Receive one or more links from one or more malware detection appliances over a network.
301
Perform a link analysis on each of the links based on link heuristics to determine whether the link is likely malicious.
302
Malicious?
No
Classify the link as non-malicious.
303
Yes
Classify the link as malicious.
304
Generate a malicious link signature for each link that is likely malicious.
305
Distribute the link signatures of the likely malicious links to the malware detection appliances for subsequent local detection.
306

FIG. 3

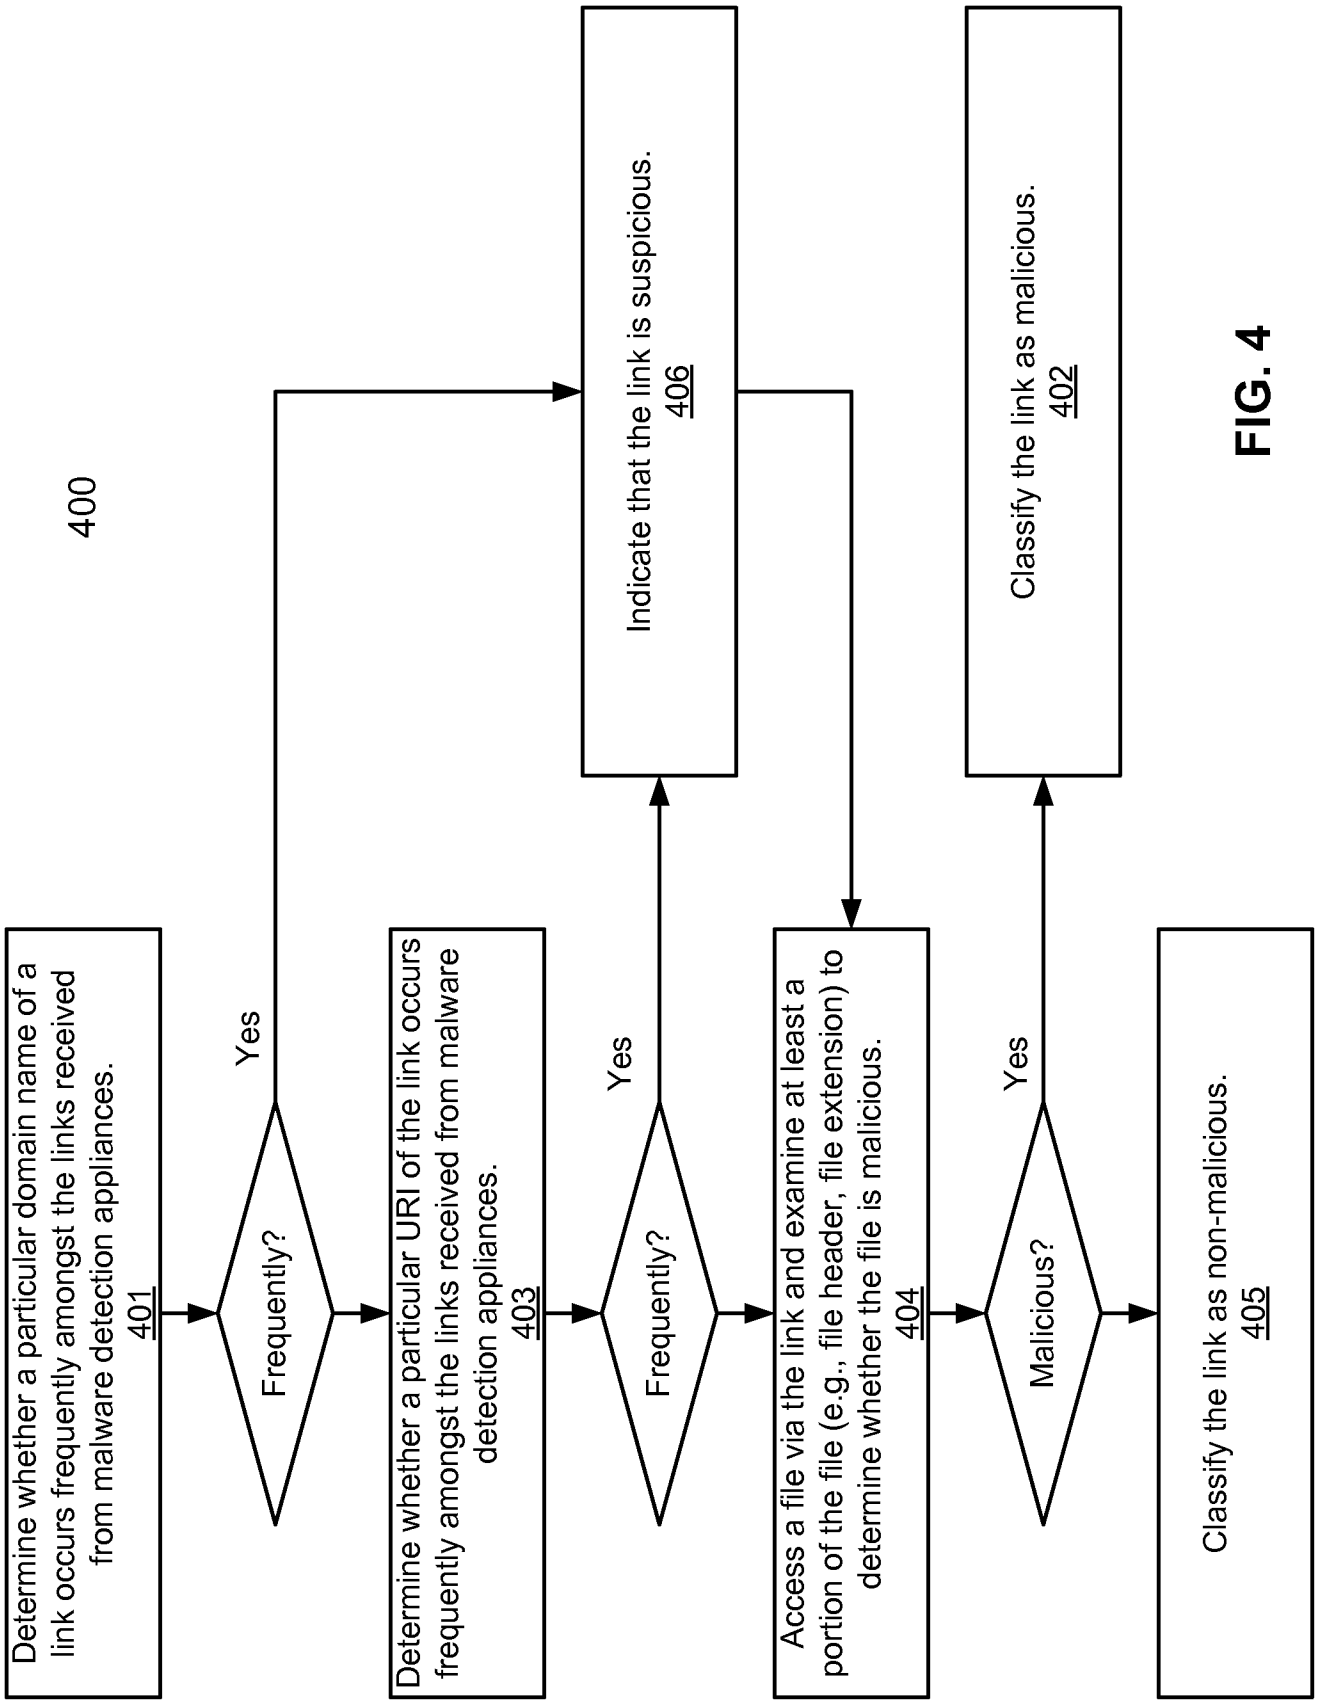
400
Determine whether a particular domain name of a link occurs frequently amongst the links received from malware detection appliances.
401
Frequently?
Yes
Determine whether a particular URI of the link occurs frequently amongst the links received from malware detection appliances.
403
Frequently?
Yes
Indicate that the link is suspicious.
406
Access a file via the link and examine at least a portion of the file (e.g., file header, file extension) to determine whether the file is malicious.
404
Malicious?
Yes
Classify the link as malicious.
402
Classify the link as non-malicious.
405

FIG. 4

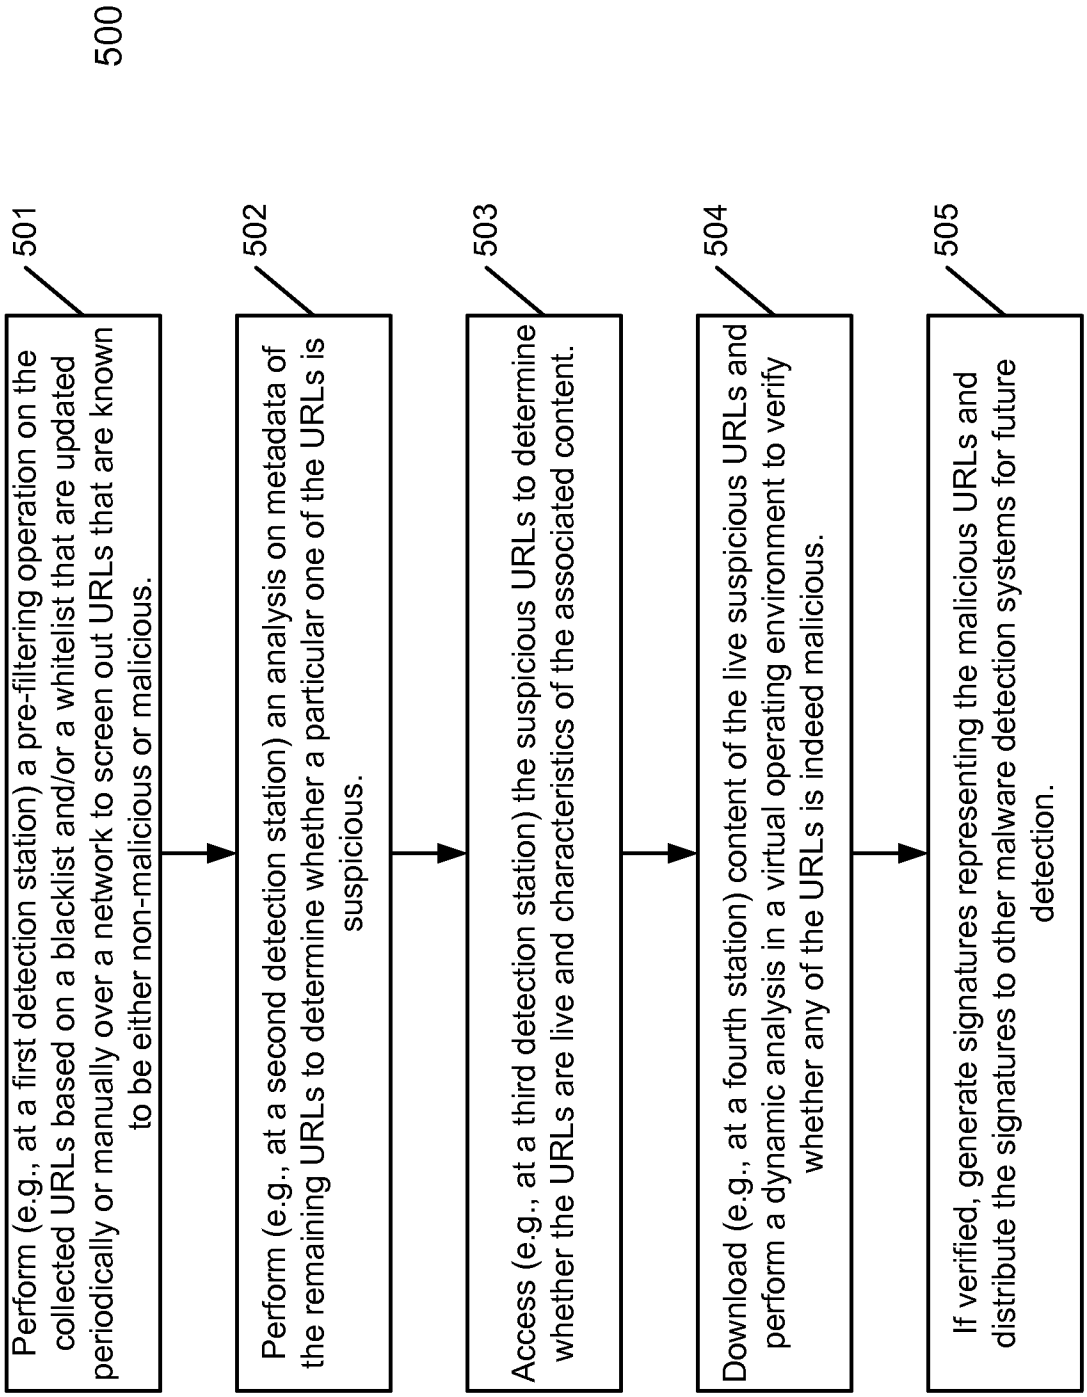
500
501
Perform (e.g., at a first detection station) a pre-filtering operation on the collected URLs based on a blacklist and/or a whitelist that are updated periodically or manually over a network to screen out URLs that are known to be either non-malicious or malicious.
502
Perform (e.g., at a second detection station) an analysis on metadata of the remaining URLs to determine whether a particular one of the URLs is suspicious.
503
Access (e.g., at a third detection station) the suspicious URLs to determine whether the URLs are live and characteristics of the associated content.
504
Download (e.g., at a fourth station) content of the live suspicious URLs and perform a dynamic analysis in a virtual operating environment to verify whether any of the URLs is indeed malicious.
505
If verified, generate signatures representing the malicious URLs and distribute the signatures to other malware detection systems for future detection.

FIG. 5

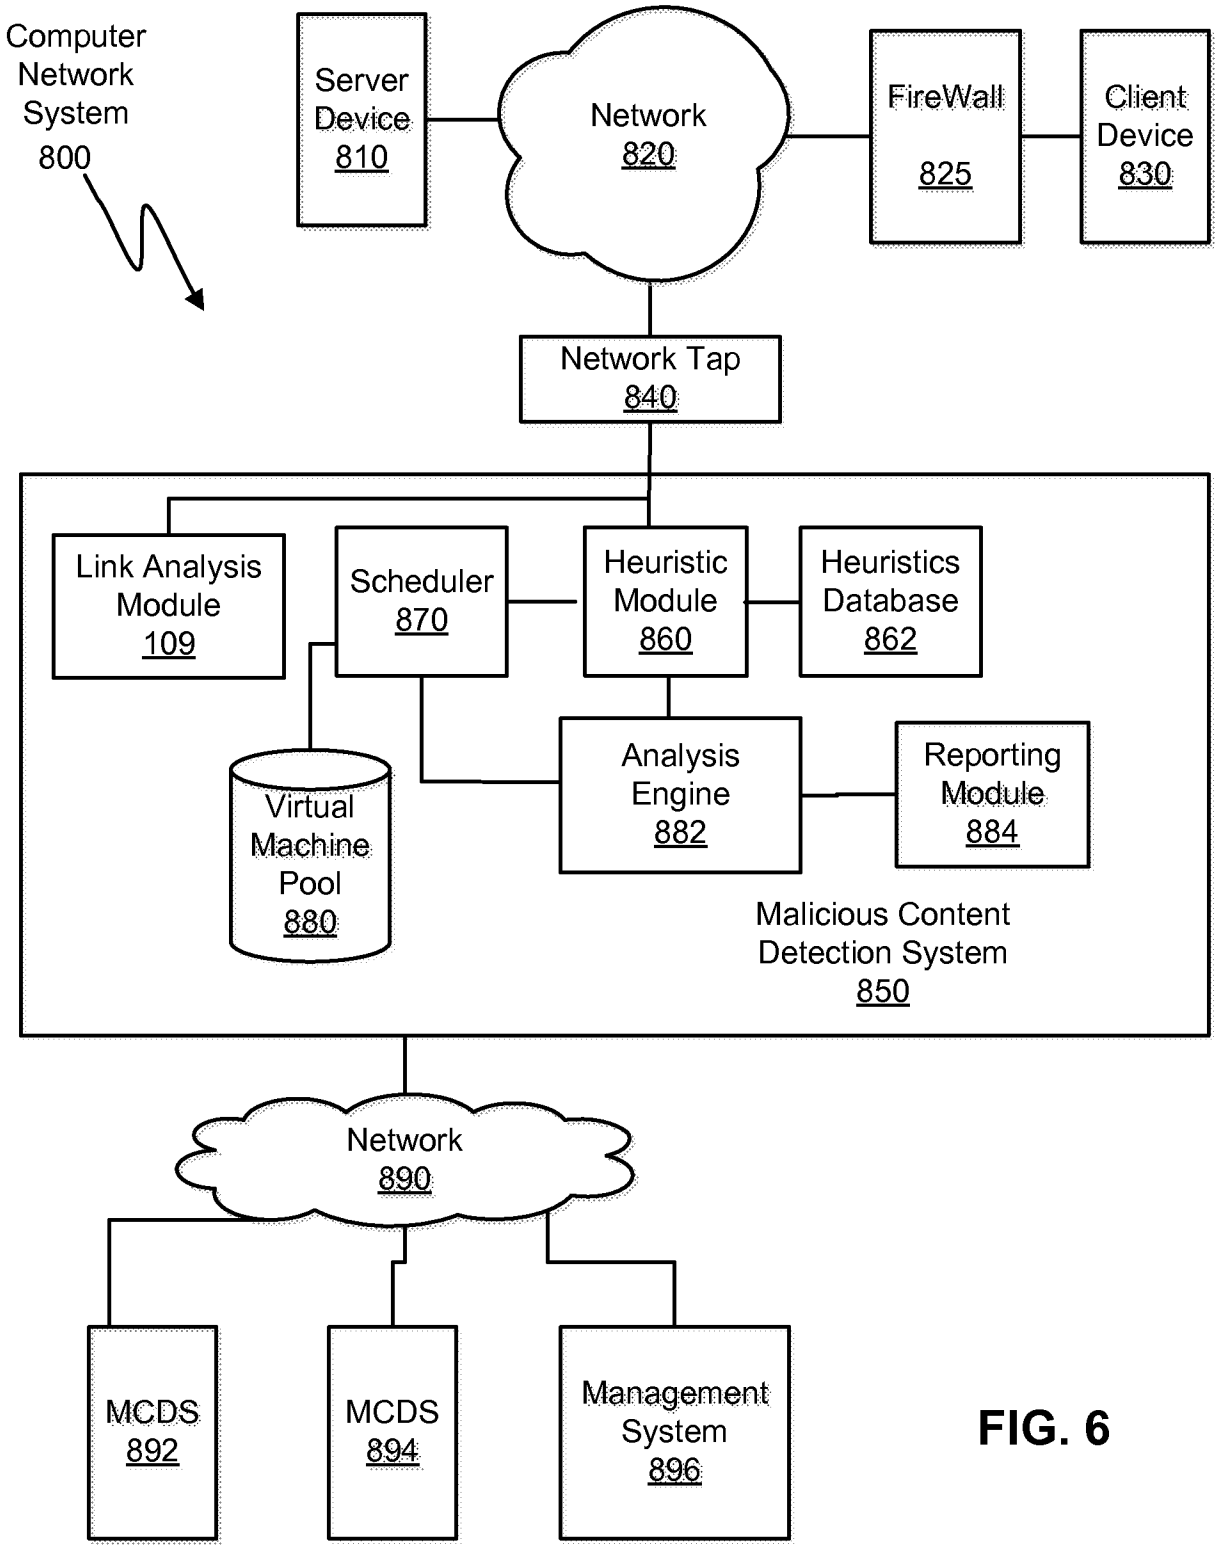
Computer Network System 800
Server Device 810
Network 820
FireWall 825
Client Device 830
Network Tap 840
Link Analysis Module 109
Scheduler 870
Heuristic Module 860
Heuristics Database 862
Virtual Machine Pool 880
Analysis Engine 882
Reporting Module 884
Malicious Content Detection System 850
Network 890
MCDS 892
MCDS 894
Management System 896

FIG. 6

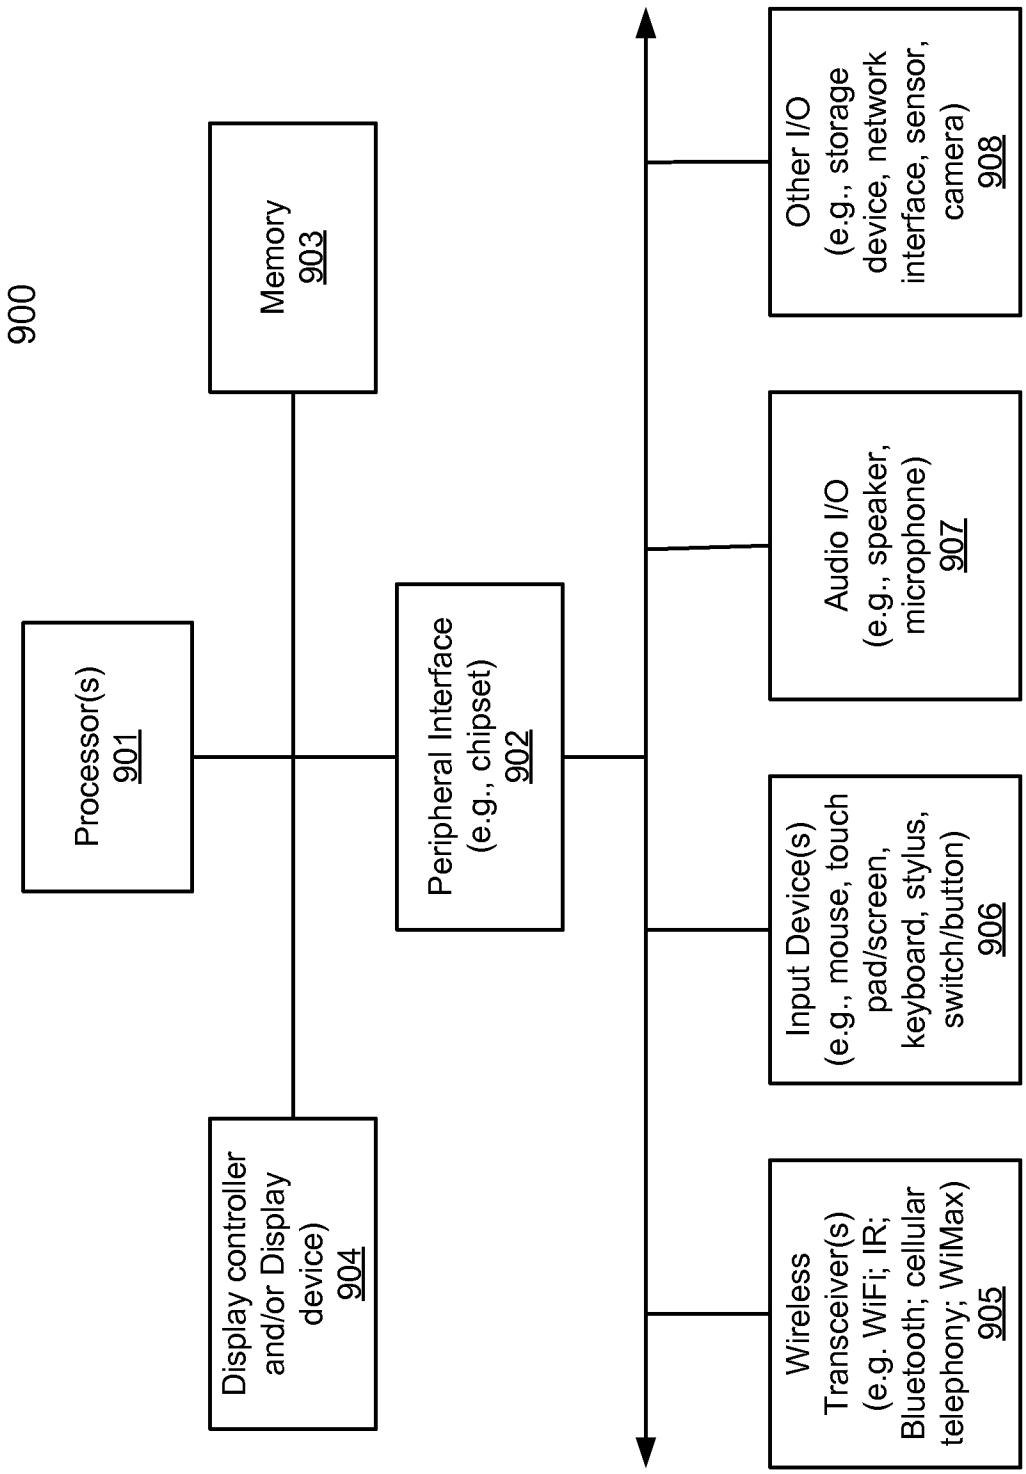
900
Processor(s)
901
Memory
903
Peripheral Interface
(e.g., chipset)
902
Display controller
and/or Display
device)
904
Other I/O
(e.g., storage
device, network
interface, sensor,
camera)
908
Audio I/O
(e.g., speaker,
microphone)
907
Input Device(s)
(e.g., mouse, touch
pad/screen,
keyboard, stylus,
switch/button)
906
Wireless
Transceiver(s)
(e.g. WiFi; IR;
Bluetooth; cellular
telephony; WiMax)
905

FIG. 7

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/043724

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F21/56 H04L29/06
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2011/314546 A1 (AZIZ ASHAR [US] ET AL) 22 December 2011 (2011-12-22) paragraphs [0004], [0005], [0018], [0024], [0030], [0032], [0043], [0044] figures 1,2,5	1-28
X	US 2011/247072 A1 (STANIFORD STUART GRESLEY [US] ET AL) 6 October 2011 (2011-10-06) paragraphs [0020], [0025], [0030], [0033], [0059], [0077], [0079] figures 1,3	1-28
A	US 2012/054869 A1 (YEN CHUI-TIN [US] ET AL) 1 March 2012 (2012-03-01) paragraphs [0048], [0053] table 4	1-28

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
31 October 2014	10/11/2014

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Pajatakis, Emmanouil

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2014/043724

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011314546 A1	22-12-2011	EP 2700009 A1	26-02-2014
		JP 2014513834 A	05-06-2014
		US 2011314546 A1	22-12-2011
		WO 2012145066 A1	26-10-2012
US 2011247072 A1	06-10-2011	EP 2666093 A1	27-11-2013
		JP 2014504765 A	24-02-2014
		US 2011247072 A1	06-10-2011
		US 2012222121 A1	30-08-2012
		WO 2012100088 A1	26-07-2012
US 2012054869 A1	01-03-2012	EP 2612488 A1	10-07-2013
		US 2012054869 A1	01-03-2012
		WO 2012030530 A1	08-03-2012