



(51) International Patent Classification:
G06F 11/00 (2006.01)

(21) International Application Number:
PCT/US2009/056623

(22) International Filing Date:
11 September 2009 (11.09.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/096,795 13 September 2008 (13.09.2008) US

(71) Applicant (for all designated States except US): **THE TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK** [US/US]; 412 Low Memorial Library, 535 West 116th Street, New York, 10027 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **STOLFO, Salvatore, J.** [US/US]; 80 Kenilworth Road, Ridgewood, NJ 07450 (US). **LI, Wei-Jen** [—/US]; 6246 Ellwell Crescent, Rego Park, NY 11374 (US).

(74) Agent: **POH, Philip, R.**; Byrne Poh LLP, 11 Broadway, Suite 865, New York, NY 10004 (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, 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, IS, JP, KE, KG, KM, 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, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, 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: METHODS, MEDIA, AND SYSTEMS FOR DETECTING ATTACK ON A DIGITAL PROCESSING DEVICE




FIG. 5A

(57) Abstract: Methods, media, and systems for detecting attack are provided. In some embodiments, the methods include: selecting a data segment in at least one portion of an electronic document; determining whether the arbitrarily selected data segment can be altered without causing the electronic document to result in an error when processed by a corresponding program; in response to determining that the arbitrarily selected data segment can be altered, arbitrarily altering the data segment in the at least one portion of the electronic document to produce an altered electronic document; and determining whether the corresponding program produces an error state when the altered electronic document is processed by the corresponding program.

**METHODS, MEDIA, AND SYSTEMS FOR
DETECTING ATTACK ON A DIGITAL PROCESSING DEVICE**

Cross Reference to Related Applications

[0001] This application claims the benefit of United States Provisional Application No. 61/096,795, filed September 13, 2008, which is hereby incorporated by reference herein in its entirety.

Statement Regarding Federally Sponsored Research Or Development

[0002] The invention was made with government support under Grant No. DE-AC05-76RL01830 awarded by Department of Energy (DOE), Grant No. 60NANB1D0127 awarded by the U.S. Department of Homeland Security through the Institute for Information Infrastructure Protection (I3P), and Grant No. W911NF-06-1-0151-49626-CI awarded by the Army Research Office (ARO). The government has certain rights in the invention.

Technical Field

[0003] The disclosed subject matter relates to methods, media, and systems for detecting attack on a digital processing device.

Copyright Notice

[0004] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

Background

[0005] Documents created and executed by various applications, including, for example, document rendering applications such as Microsoft® Word® and Adobe® Acrobat® include not only simple binary content interpreted by the document rendering applications, but also can include, as part of the documents themselves, software necessary to interpret data in the documents. Because of their ability to contain and execute software,

such documents can be considered complex code injection platforms. The injected code can be of various types, such as, for example, Macros (e.g., scripts written in Microsoft® Visual Basic®) and Javascript® (e.g., embedded in Adobe PDF® files).

[0006] While the ability to embed software into documents provides various advantages to users, it can also be used by attackers to launch attacks on digital data processing devices. In some cases, malicious code may attack upon execution. In other cases, embedded malicious code can lie dormant for use in a future multi-partite attack. For example, one type of attack embeds malicious code in the padding areas of the binary file format of documents or to replace normal textual data with malicious code.

[0007] One issue in inhibiting such attacks is that it can be difficult for a user or a system to determine whether code embedded in a document is, for example, useful and friendly or harmful and malicious. For example, software can be injected into a document as obfuscated encoded code (e.g., code represented as image data that, when decoded and rendered at runtime, can be executed to perform malicious activities). In some cases, attackers may even entice a user to launch embedded malicious code. For example, as illustrated in FIG. 20, embedded malicious object 2010 has the message "CLICK HERE" displayed below it. If the user follows these instructions, the user will have launched an attack on the user's own system. In some cases, a parsed document in the Object Linking and Embedding (OLE) structured storage format, which contains nodes and directories, can harbor various exploits, such as buffer overflows or vulnerabilities to other applications. For example, FIG. 22 illustrates an example of the internal structure of a parsed document in OLE format, where attackers may craft data that exploit the vulnerabilities which redirect the execution of Microsoft® Word® to a particular location to execute arbitrary embedded malicious code, such as in the "Table" sector.

[0008] In some cases, attackers may obfuscate or shape the attacking code so that it appears to be the same as, for example, benign code surrounding it. Code, including benign code, tends to have a high entropy statistical distribution, so, some attackers, for example, may inject malicious code into benign code in an attempt to avoid detection. FIG. 21 illustrates an uninfected Microsoft® Word® document 2111 and the same document 2112 embedded with a known malicious code sample (in this case, the malicious code is known as "Slammer"). A Symantec® anti-virus scanner has been installed and is running on this system, however, it does not detect the embedded malicious code even though Slammer is a

known attacker. In addition, there is no discernable change to the appearance of the document that may, for example, make a user suspicious that malicious code is present.

[0009] Generally speaking, embedding malicious code within documents is a convenient approach to attack a digital processing device. Such attacks can be targeted and difficult to stop due to the number of document-exchange vectors and particular vulnerabilities in word processing programs. Moreover, detecting malicious code embedded in a document is increasingly difficult due to the complexity of modern document formats.

[0010] Accordingly, it is desirable to provide methods, media, and systems that overcome these and other deficiencies of the prior art.

Summary

[0011] Methods, media, and systems for detecting attack are provided. In some embodiments, methods for detecting are provided. The methods include comparing at least part of a document to a static detection model; determining whether attacking code is included in the document based on the comparison of the document to the static detection model; executing at least part of the document; determining whether attacking code is included in the document based on the execution of the at least part of the document; and if attacking code is determined to be included in the document based on at least one of the comparison of the document to the static detection model and the execution of the at least part of the document, reporting the presence of an attack.

[0012] In some embodiments, computer-readable media containing computer-executable instructions that, when executed by a processor, cause the processor to perform methods for detecting an attack on an application are provided. The methods include comparing at least part of a document to a static detection model; determining whether attacking code is included in the document based on the comparison of the document to the static detection model; executing at least part of the document; determining whether attacking code is included in the document based on the execution of the at least part of the document; and if attacking code is determined to be included in the document based on at least one of the comparison of the document to the static detection model and the execution of the at least part of the document, reporting the presence of an attack.

[0013] In some embodiments, systems for detecting attack including, an interface in communication with a network; a memory; and a processor in communication with the memory and the interface are provided, wherein the processor compares at least part of a

document to a static detection model; determines whether attacking code is included in the document based on the comparison of the document to the static detection model; executes at least part of the document; determines whether attacking code is included in the document based on the execution of the at least part of the document; and if attacking code is determined to be included in the document based on at least one of the comparison of the document to the static detection model and the execution of the at least part of the document, reports the presence of an attack.

[0014] In some embodiments, methods for detecting are provided. The methods include executing at least part of a document, wherein a load order of shared objects is controlled; detecting for an indication of a crash; and reporting the presence of an attack, based on the detection for an indication of a crash.

[0015] In some embodiments, methods for detecting are provided. The methods include receiving a first document in a first format; converting the document to a second format to create a second document; converting the second document to the first format to create a third document; comparing the first document to the third document; determining whether attacking code is included in the first document based on the comparison; and if attacking code is determined to be included, reporting the presence of an attack.

[0016] In some embodiments, methods for detecting malicious code in electronic documents are provided, the methods include: selecting a data segment in at least one portion of an electronic document; determining whether the arbitrarily selected data segment can be altered without causing the electronic document to result in an error when processed by a corresponding program; in response to determining that the arbitrarily selected data segment can be altered, arbitrarily altering the data segment in the at least one portion of the electronic document to produce an altered electronic document; and determining whether the corresponding program produces an error state when the altered electronic document is processed by the corresponding program.

[0017] In some of these methods, the electronic document is a word processing document. In some of these methods, the corresponding program is a word processor. In some of these methods, the at least one portion of the electronic documents is altered by changing the values of data in the at least one portion by a given value or by changing the values of data in the at least one portion by an arbitrarily selected displacement.

[0018] In some embodiments, systems for detecting malicious code in electronic documents are provided, the systems comprising: at least one digital processing device that:

selects a data segment in at least one portion of an electronic document; determines whether the arbitrarily selected data segment can be altered without causing the electronic document to result in an error when processed by a corresponding program; in response to determining that the arbitrarily selected data segment can be altered, arbitrarily alters the data segment in the at least one portion of the electronic document to produce an altered electronic document; and determines whether the corresponding program produces an error state when the altered electronic document is processed by the corresponding program.

[0019] In some of these systems, the electronic document is a word processing document. In some of these systems, the corresponding program is a word processor. In some of these systems, the at least one portion of the electronic documents is altered by changing the values of data in the at least one portion by a given value or by changing the values of data in the at least one portion by an arbitrarily selected displacement.

[0020] In some embodiments, computer-readable media containing computer-executable instructions that, when executed by a processor, cause the processor to perform a method for detecting malicious code in electronic documents, are provided, the method comprising: selecting a data segment in at least one portion of an electronic document; determining whether the arbitrarily selected data segment can be altered without causing the electronic document to result in an error when processed by a corresponding program; in response to determining that the arbitrarily selected data segment can be altered, arbitrarily altering the data segment in the at least one portion of the electronic document to produce an altered electronic document; and determining whether the corresponding program produces an error state when the altered electronic document is processed by the corresponding program.

[0021] In some of these media, the electronic document is a word processing document. In some of these media, the corresponding program is a word processor. In some of these media, the at least one portion of the electronic documents is altered by changing the values of data in the at least one portion by a given value or by changing the values of data in the at least one portion by an arbitrarily selected displacement.

Brief Description of the Drawings

[0022] FIG. 1 is a simplified illustration of a method for detecting attack on a digital processing device in accordance with some embodiments of the disclosed subject matter.

[0023] FIG. 2 is table listing sections that a Microsoft® Word® document can be parsed into in accordance with some embodiments of the disclosed subject matter.

[0024] FIG. 3 is a simplified illustration of a method for detecting content anomalies in accordance with some embodiments of the disclosed subject matter.

[0025] FIG. 4 is a simplified illustration of a method for generating, training, and sharing a binary-based content anomaly detection model and for using the content anomaly model to detect content anomalies in accordance with some embodiments of the disclosed subject matter.

[0026] FIG. 5 is an illustration of two environments in accordance with some embodiments of the disclosed subject matter.

[0027] FIG. 5A is a schematic diagram of a system suitable for generating at least part of a dynamic detection model using the environments illustrated in FIG. 5 in accordance with some embodiments of the disclosed subject matter.

[0028] FIG. 6 is a simplified illustration of a method for detecting attack on a digital processing device in accordance with some embodiments of the disclosed subject matter.

[0029] FIG. 7 is an illustration of a screenshot of a message indicating an application has crashed that can be used to detect attack in accordance with some embodiments of the disclosed subject matter.

[0030] FIG. 8 is an illustration of a screenshot of a message that can be interpreted as indicating an application has crashed in accordance with some embodiments of the disclosed subject matter.

[0031] FIG. 9 is a simplified illustration of a method for generating at least part of a dynamic detection model in accordance with some embodiments of the disclosed subject matter.

[0032] FIG. 10A is an illustration showing changes which can be interpreted as malicious in accordance with some embodiments of the disclosed subject matter.

[0033] FIG. 10B is an illustration showing changes which can be interpreted as benign in accordance with some embodiments of the disclosed subject matter.

[0034] FIG. 11 is an illustration of a screenshot of a message that can be interpreted as indicating attack in accordance with some embodiments of the disclosed subject matter.

[0035] FIG. 12 is another an illustration of a screenshot of a message that can be interpreted as indicating attack in accordance with some embodiments of the disclosed subject matter.

[0036] FIG. 13 is a simplified illustration of a method for detecting attack on a digital processing device in accordance with some embodiments of the disclosed subject matter.

[0037] FIGS. 14-18 are illustrations of screenshots of a graphical user interface (GUI) for a toolkit in accordance with some embodiments of the disclosed subject matter.

[0038] FIG. 19 is an illustration of a system that can be used in accordance with some embodiments of the disclosed subject matter.

[0039] FIG. 20 is an illustration of a screenshot of embedded malicious code that can be detected in accordance with some embodiments of the disclosed subject matter.

[0040] FIG. 21 is another illustration of a screenshot of embedded malicious code that can be detected in accordance with some embodiments of the disclosed subject matter.

[0041] FIG. 22 is an illustration showing an example of the internal structure of a parsed document in Object Linking and Embedding (OLE) structured file format in accordance with some embodiments of the disclosed subject matter.

[0042] FIG. 23 is a simplified illustration of a method for detecting malicious code in electronic documents in accordance with some embodiments of the disclosed subject matter.

[0043] FIG. 24 is an illustration of a screenshot of two documents before and after applying the arbitrary data transformation mechanism in accordance with some embodiments of the disclosed subject matter.

[0044] FIGS. 25 and 26 are illustrations showing examples of the detection of malicious code embedded in electronic documents in accordance with some embodiments of the disclosed subject matter.

[0045] FIG. 27 is an illustration showing examples of messages indicating that the document may be damaged in accordance with some embodiments of the disclosed subject matter.

Detailed Description

[0046] Methods, systems, and media for detecting and/or inhibiting attack on a digital processing device are provided. In some embodiments, the methods, systems, and media of the disclosed subject matter can detect attack by comparing received documents to models created by analyzing known benign and malicious applications and/or by executing received documents and comparing the behavior of the execution to models of behavior of execution of known benign and/or malicious applications. For example, as illustrated in FIG. 1, in some embodiments, a document 10 can be received at, for example, a digital processing device connected to a network. Document 10 can be of various file types and can include, for example, embedded code. Document 10 can be parsed, at 20, into various sections. In some embodiments, static analysis can be performed, at 30, on document 10 and/or sections of documents 10, to determine whether malicious code is detected within document 10. In performing static analysis, at 30, document 10 and/or sections of document 10 can be compared to a detection model 35 of, for example, known malicious and/or known benign code. Dynamic analysis can be performed, at 40, on document 10 and/or sections of document 10 to determine whether malicious code is detected within document 10. In performing dynamic analysis, at 40, behavior exhibited by execution of document 10 and/or sections of document 10 can be compared to a detection model 45 of, for example, the behavior of known malicious and/or known benign code. The presence or absence of malicious code in document 10 can be responded to, at 50, by, for example, reporting the presence of malicious code to a user and/or a system administrator, cleaning document 10 of the detected malicious code, and/or deleting and/or quarantining document 10. In some embodiments, if malicious code is detected by, for example, dynamic analysis, at 40, but not by static analysis, at 30, detection model 35 can be updated, at 60, so that, for example, the same or similar malicious code detected, at 40, can be detected, at 30, if seen again. Document 10 can include documents and/or embedded objects of various formats, such as, for example, Microsoft® Word® formats, Microsoft® Excel® formats, Microsoft® Power Point® formats, Adobe® Portable Document Format, PostScript, Corel® WordPerfect® formats, Joint Photographic Experts Group (JPG), Tagged Image File Format (TIFF), Moving Picture Experts Group (MPEG) (e.g., MP3, MP4, etc.), and/or Resource Interchange File Format (RIFF), etc.

[0047] In some embodiments, document 10 may not be parsed at 20, and instead, for example, static analysis, at 30, and/or dynamic analysis, at 40, can be performed directly on

document 10. However, in some embodiments, static analysis, at 30, and/or dynamic analysis, at 40, can be performed on sections of document 10. In such embodiments, document 10 can be parsed, at 20, into various sections. The sections can include, for example, embedded objects, tables, images, plain text, injected code, etc. Various parsing systems and/or methods can be used, at 20, to parse a document 10 of various formats into various sections. For example, Apache® POI, available at <http://poi.apache.org/>, can be used to parse Microsoft® Word® documents into their constituent structures. These structures can be referenced by, for example, structure reference names using Object Linking and Embedding (OLE). As listed in table 210 of FIG. 2, the parsed sections of a Microsoft® Word® document can include various sections. Sections can be arbitrarily named and can be defined by an author of a document. In addition, some sections (e.g., sections 211) can be included by default in a Microsoft® Word® Document. For example, referring to default sections 211 of table 210 of FIG. 2, VBA_PROJECT identifies a Visual Basic Application, Data identifies binary data, OTable and ITable identify tabular structures, PROJECT can identify a macro and/or a Visual Basic project, SummaryInformation identifies summary information of a document, and WordDocument identifies the text portion of a document.

[0048] Various systems and methods can be used to perform static analysis, at 30 of FIG. 1, such as, for example, Mahalanobis distance calculation, n-gram analysis, n-perm analysis, and/or common n-gram analysis. For example, FIG. 3 is a simplified illustration of a method for detecting content anomalies that can be performed, at 30 of FIG. 1. As shown, at 302, a content anomaly detection model can be generated and trained. In some embodiments, this model can be part of, or all of, model 35 of FIG. 1. At 304, the content anomaly detection model can be used to detect anomalies. Various embodiments of the method of FIG. 3 are described in, for example, International Publication No. WO 2007/100916, entitled "Systems, Methods, and Media for Outputting a Dataset based upon Anomaly Detection," which is hereby incorporated by reference herein in its entirety.

[0049] Other methods which can be performed at 30 of FIG. 1 include, for example, methods for generating, training, and sharing a binary-based content anomaly model and for using the content anomaly model to detect content anomalies, as illustrated in, for example, FIG. 4. As shown, at 402, a binary-based content anomaly detection model is generated and trained. In some embodiments, this model can be part of or all of model 35 of FIG. 1. In some embodiments, a binary-based content anomaly model can be generated and trained using known benign training datasets. In some embodiments, a binary-based content

anomaly model can be generated and trained using known anomaly signatures. In some embodiments, a binary-based anomaly detection model can be trained in cooperation with a host-based detector. At 404, the binary-based content anomaly detection is shared. A shared model can be received from one or more remote sites to compare with and update a local model (e.g., model 35). At 406, the binary based content anomaly detection model is used to detect content anomalies. For example, an anomaly score of an input dataset can be used to determine whether the input dataset is anomalous. Various embodiments of the method of FIG. 4 are described in, for example, International Publication No. WO 2007/100915, entitled "Systems, Methods, and Media for Outputting Data based on Anomaly Detection," which is hereby incorporated by reference herein in its entirety. Other examples of systems and methods that can be used to detect attack, for example, at 30 of FIG. 1, are described in, for example, International Publication No. WO 2005/047862, entitled "Apparatus Method and Medium for Identifying Files using N-gram Distribution of Data," which is hereby incorporated by reference herein in its entirety.

[0050] FIG. 5 illustrates two environments which can be used in accordance with some embodiments of dynamic analysis, at 40, of FIG. 1. Environment 500 illustrates an emulated environment and can run, for example, Microsoft® Windows® 501 on a digital processing device. A virtual machine, such as, for example, VMware WorkStation can be installed and Linux® 502 can be installed and run in it. A Microsoft® Windows® binary translator and emulator 503, such as, for example, CrossOver® Office Standard 5.0.1 can be installed in Linux® 502. An application 504, such as, for example, Microsoft® Word®, can be run inside emulator 503 and used to open and/or execute a document 10. Environment 550 illustrates a non-emulated environment and can run, for example, Microsoft® Windows® 551 on a digital processing device. Though environment 550 is considered non-emulated, in some embodiments, a virtual machine, acting as a sandbox to protect the underlying system, can be used to run an instance of Microsoft® Windows® 552 inside Microsoft® Windows® 551. Inside Microsoft® Windows® 551, an application 553, can be executed. In some embodiments, application 504 and application 553 can be the same application with the same set up. For example, both application 504 and application 553 can be Microsoft® Office Word 2007® with the same updates (e.g., service packs, security updates, bug fixes, etc.).

[0051] Some embodiments can manipulate the order in which shared libraries are loaded by, for example, application 504 and/or 553. Various applications in various systems

employ shared libraries that can include various shared code and/or data. In some systems, the order in which some shared libraries are loaded to memory typically does not cause benign code to crash. However, attacks launched by malicious code can be designed with a particular shared library load order in mind and changing the order can cause a malicious application and/or the application and/or operating system running it, to crash. Some embodiments can use an emulated environment, such as, for example, environment 500, to reveal code exceptions which depend on the order in which shared objects are loaded to memory by changing the order in which shared objects would normally be loaded. For example, Microsoft® Windows® implements shared objects as DLLs (Dynamic-link libraries) and changing the load order of DLLs can cause some malicious programs to crash or to cause, for example, the application and/or operating system executing them to crash. As illustrated in FIG. 6, in some embodiments, as part of dynamic analysis, at 40 of FIG. 1, a document 10 can be opened, at 610, in environment 500, with the load order of DLLs altered, and it can be determined, at 620, whether a system crash has occurred. If a crash has occurred, the document can be reported, at 50, as including malicious code.

[0052] In some embodiments, the determination of whether an application crashes, at 620, can be detected and/or reported by the application and/or the operating system. For example, FIG. 7 illustrates a user being informed by pop-up window 710 that Microsoft® Word® has crashed. Crashes can include, for example, entering an infinite loop, an unexpected termination of the application, and/or the inability to open document 10. Another example of behavior that can be considered a crash, at 620, is the presence of pop-up window 810 of FIG. 8, which alerts a user that an application was not able to read a document 10. In some embodiments, in addition to, or instead of, monitoring messages from an application or operating system of detection of a crash, an application and/or operating system can be monitored directly, for example, by monitoring the status and/or presence of one or more process identifiers.

[0053] If a crash is not detected at 620, it can be determined, at 630 and 640, whether executing document 10 results in unexpected behavior. In some embodiments, detecting unexpected behavior can include, for example, comparing the behavior observed when executing a document 10 in a first environment to behavior seen when executing known benign and/or malicious documents in the first environment. In some embodiments, detecting unexpected behavior can include executing a document 10 in various environments and comparing the differences in behavior observed when executing a document 10 in

various environments to behavior seen when executing known benign and/or malicious various environments. For example, differences between the files generated and/or altered and/or differences between various events posted by an application opening document 10. The various environments can include, for example, environment 500 of FIG. 5 and environment 550 of FIG. 5. In some embodiments, malicious behavior resulting from executing a document 10 can be detected by comparing behavior observed when opening document 10 to a model 45, which can include, for example, a list or lists of expected and/or unexpected behavior and/or a list or lists of expected and/or unexpected differences in behavior among execution in various environments.

[0054] To generate a model 45 of expected and/or unexpected behavior and/or expected and/or unexpected differences in behavior, benign and malicious documents can be executed and changes to the systems in which they are executed can be monitored and recorded. FIG. 5A illustrates a system that can be used to generate a model 45. For example, method 910 of FIG. 9 can be performed on the system of FIG. 5A. In some embodiments, training document 580 can be a set of benign documents and/or sections of benign documents and can be opened and/or executed, at 911, in environment 550. The same set of benign documents 580 and/or sections of benign documents 580 can be opened and/or executed, at 912, in environment 500. The behavior observed by execution, at 911 and 912, can be recorded to create records 913 and 914. Records 913 and 914 can be compared, at 915, and the differences can be recorded and considered a record of expected differences 920. Record 920 can be saved as part of model 45.

[0055] To identify unexpected differences that indicate malicious code, some embodiments can execute a collection of malicious documents. For example, training documents 580 can include known malicious documents and/or sections of malicious documents (possibly including malicious documents and/or sections of documents identified using static analysis, at 30, of FIG. 1). From these differences, for example, known differences 920 can include a list of unexpected differences (e.g., malicious changes) and be included as part of model 45. In some embodiments, expected differences can be used to avoid false positives and unexpected differences can be used to improve malicious code detection. Changes which occur both when executing a document in both environment 550 and environment 500 can still be included in model 45 if, for example, they provide an indication of malicious behavior. For example, if a malicious application exhibits the same

behavior regardless of its environment, but this behavior is different than the behavior of a benign application, it can be included in model 45.

[0056] Returning to FIG. 6, at 640, model 45 can be compared to changes observed, at 630, considering document 10. If the changes are considered to be malicious, document 10 can be reported, at 50, as including malicious code.

[0057] FIG. 10A, box 1010, illustrates an example of changes caused by the execution of a malicious document. In this case, lines 1011 have been added to the file index.dat as a result of executing a malicious document. However, on this system, when opening benign documents, either no change in the index.dat file is observed or lines 1012 are added. Accordingly, in some embodiments, if the execution of a document causes lines 1011 to be added to index.dat, the document can be considered malicious, at 640. Box 1020, line 1021, identifies a portion of a user registry section that, in some systems, is only modified when malicious documents are opened. Accordingly, in some embodiments, a document that adds line 1021 to user.reg can be considered malicious, at 640.

[0058] FIG. 10B, box 1030, illustrates example changes made, in some systems, by execution of a benign document. The contents of the system section of the registry in an emulated environment before running the benign document are shown in lines 1031. The contents of the system section of the registry in an emulated environment after running the benign document are shown in lines 1032. Box 1040, line 1041, illustrates a key entry that is expected to be added to the user registry section in some systems. In such systems, when changes 1032 or 1041 are observed they will not be considered to provide an indication of attack.

[0059] In considering some other changes, for example, observed using method 910, assume the following. There is a system where all benign files modify system.reg and user.reg files and most of the time also modify index.dat. However, there are some cases of malicious files that do not change user.reg. This constitutes a considerable change from normal behavior and can be used to categorize a document that exhibits this behavior as malicious. Hence, in such a system, the absence of these changes can be used to indicate a file is malicious, at 640. In this example system, it can also be that both benign and malicious documents that use network connections modify the system.reg registry section. However, some of the malicious documents modify system.reg considerably more than is usual for benign files. Such behavior can be used to reveal the presence of malicious code in an unknown file, at 640.

[0060] Returning to FIG. 6, in some embodiments, if the changes are not considered malicious, at 640, it can be determined, at 650, whether an abnormal message was produced, and if so, document 10 can be reported, at 50, as including malicious code. For example, some embodiments can examine system messages that do not cause an operating system or emulated environment to fail, but can be indicators of malicious code causing part of an application to gracefully terminate. For example, a user can be presented with a pop-up window requesting his input or intervention before the document can be viewed. Such a popup message can be used for detection of malicious documents. For example, FIG. 11 and FIG. 12 illustrate pop-up messages 1110 and 1120 that can be considered an indication of malicious code. Which pop-up messages to consider malicious can be determined by comparing pop-up messages observed at 650 to a list of known messages which indicate the presence of malicious code. This list can include messages detected, for example, when running a known malicious document in an environment 500, but not in environment 550 or detected only when running malicious code in any environment. In some embodiments, the changes that are considered malicious, at 630, or the messages, at 650, can be controlled by user and/or administrator settings. This can be used, for example, to set a level of risk a user of the system is willing to accept.

[0061] Again, referring to FIG. 6, if an abnormal message was not detected, at 650, the document can be reported as benign, at 50. In different embodiments, various tests can be performed, at 40, in various orders and whether a specific test is run can depend on the success or failure of another test. For example, in some embodiments, test 640 may only be run if test 620 was successful, while in other embodiments each test may be run every time. Whether to run, and in what order to run, various tests (e.g., 620, 640, and 650) can be controlled by settings set by, for example, a user and/or system administrator. The response to the detection of attack can also be controlled by such settings. In addition, dynamic analysis, at 40, can include various other tests.

[0062] As illustrated in FIG. 13, some embodiments can convert, at 1310, document 10 from a first file format to a document 1320 of a second file format. Document 1320 can then be converted, at 1330, to a third document 1340 of the first file format. Document 1340 and document 10 can be compared, at 1350, to determine differences. At 1360, it can be determined whether the differences indicate the presence of malicious code. Comparing, at 1350, can include, for example, a comparison of the data which makes up document 10 and document 1340. In such embodiments, the method of FIG. 13 can be performed for example,

at 30 of FIG. 1. Comparing, at 1350, can also include executing document 10 and document 1340, in at least one environment, and comparing their behavior and/or differences in their behavior to, for example, known differences between benign and malicious code. In such embodiments, the method of FIG. 13 can be performed, for example, at 40 of FIG. 1.

[0063] In some embodiments, converting, at 1310 can include for example, using a virtual PDF printer. If, for example, a Microsoft® Word® document includes a table, which includes malicious code, converting the document (including the table), at 1310, to a PDF and then converting the PDF, at 1330, back to a Word® document can produce a Word® document which appears the same as the original, but no longer contains the malicious code. In addition, the original Word® document and the new Word® document can be compared and it can be determined that differences between them appear to be due to malicious code. In some embodiments, detection model 35 and/or detection model 45 can be updated to include, for example, signatures of, or behavior observed by, the malicious code detected at 1360.

[0064] Some embodiments provide a toolkit that can implement static analysis techniques and/or parsing, such as, for example, those illustrated at 20 and 30 of FIG. 1. Some embodiments of such a toolkit can include a graphical user interface (GUI). The toolkit can be implemented using various methods. For example, it can be implemented using the Java 2 Platform Standard Edition 5.0 Development Kit (JDK 5.0) and can operate on a platform using J2SE Runtime Environment 5.0 (JRE 5.0). The toolkit can implement static analysis techniques using, for example, a database application on top of a MySQL database.

[0065] An embodiment of a GUI to a toolkit implementing static analysis methods is illustrated in FIG. 14. The GUI includes five tabs. Tabs 1410 and 1420 are directed to parsing and displaying documents. Tabs 1430, 1440, and 1450 are directed to performing and displaying the results of training and testing documents. Tab 1410, labeled, "Parse Single File", displays the filename 1411 of document 10 and names of the parsed sections 1412. The information about the selected section, in this case section 1413, is displayed in panel 1414. The information can include, for example, the file length, section length, section offset 1415 in this file, and the exact byte sequence 1416. The byte values of sequence 1416 are represented in decimal format, ranging from -128 to 127, and the offsets 1415 can be displayed using both decimal and hexadecimal format. A bar 1417 is located at the bottom of the display that shows the location of the selected section in the file. The bar 1417 represents

the entire file 1411. Section 1418 of the bar represents the selected section 1413. In some cases, a section can have more than one subsection and can be distributed in different areas in a file. In some embodiments, various colors and/or patterns can be used to identify the various sections and/or subsections. For example, section 1418 of the bar can be red, and the rest of the bar 1417 can be blue.

[0066] In some embodiments, a user can save the byte sequence of a selected section 1413. This can be done, by using the "Save Section to File" button 1419. A section 1413 can be loaded and displayed in the tab "Load & Display Section" 1420 using button 1510 as illustrated in FIG. 15. In the "Load Experimental Files" tab 1430, illustrated in FIG. 16 a user can specify and load benign training files 1611, malicious training files 1612, and testing files 1613.

[0067] Tab 1440, "Run Experiment," illustrated in FIG. 17, includes panel 1710, which displays configuration information, and panel 1720, which displays results. The run test button 1711, starts an experiment. Button 1712 can save the current training model and button 1713 can load a previously saved training model. Configuration panel 1710 can include various configurations options. For example, option 1715 can be used to set the size of n-grams extracted from the files. Option 1716 can be set to remove malicious n-grams from benign files. For example, when training the benign files, n-grams that appear in a malicious model can be removed. Option 1717 can be set to remove benign n-grams from malicious files. For example, when training the malicious files, n-grams that appear in a benign model can be removed. Option 1718 can be set to ignore zero bytes. For example, when training the model, zero byte sequences (e.g., 0, 00, 000, 0000) can be skipped.

[0068] When option 1719 is set, the entire file is trained and tested. When option 1720 is set, only uncommon sections are trained and tested. Common sections include, for example, WordDocument, SummaryInformation, 1 Table, CompObj, Data, and DocumentSummaryInformation. When option 1721 is set, only the text section (e.g. a WordDocument section) is trained and tested in the files. When option 1721 is set, sections that have a "macro" keyword are trained and tested. Using option 1722 and field 1723, a user can select a specific section to test. Using option 1724 and field 1726, a user can specify a keyword or keywords of a section to be trained and tested.

[0069] As illustrated in results panel 1720, in some embodiments, the score of a document that is calculated from the benign and malicious models is displayed here. The higher the similarity score, the closer the section of the document is to the model that was

loaded. The score of a section and the average score of the displayed sections can be displayed when more than one section is selected.

[0070] Tab 1450, "Extra Experimental Result", illustrated in FIG. 18, can graphically display a suspect malicious code location within a file. The suspicious regions are represented in bold 1821 in panel 1820 and are also shown at 1831 in bar 1830. In some embodiments, byte sequences of selected testing sections are checked using a Bloom filter in the benign training model. A byte sequence with size n (n -gram) is deemed to be suspicious if it does not appear in the benign model.

[0071] FIG. 19 illustrates a system 1900 that can be used in accordance with some embodiments of the disclosed subject matter. System 1900 can include, among other things, digital processing device 1902, display 1904, input device 1906, and memory 1908, which can be interconnected. Digital processing device 1902 can be, for example, a processor, a digital signal processor, a controller, etc. In some embodiments, memory 1908 contains a program for controlling digital processing device 1902. Memory 1908 can also contain applications for performing the methods described herein, for example, the method illustrated in FIG. 1. In some embodiments, various other applications can be resident in the memory 1908 and executed by digital processing device 1902. System 1900 can include additional components, for example, additional input and/or output devices. Various systems 1900 can be interconnected and/or connected with other devices through, for example, networks such as the Internet, cellular networks, local area networks, and/or wide area networks. In some embodiments, the method illustrated in, for example, FIG. 1, can be implemented across various systems 1900. For example, with reference to FIG. 1, static analysis, at 30, can be performed on a first system 1900 and dynamic analysis, at 40, can be performed on a second system 1900. A document 10, model 35, and/or model 45 can be sent to and from various systems 1900 through various networks.

[0072] International Publication No. WO 2008/036665, filed September 18, 2007 and entitled "Methods, Media, and Systems for Detecting Attack on a Digital Processing Device," is hereby incorporated by reference herein in its entirety.

[0073] It should be noted, however, that additional detection approaches can be provided. For example, there may be situations where malicious code can evade detection approaches that use static and/or dynamic detection and analysis.

[0074] For example, the ITable sector shown in FIG. 22, which illustrates an example of the internal structure of a parsed document in Object Linking and Embedding (OLE)

format, is currently one of the most vulnerable portions of Word documents. The data in the ITable sector can include text such as "Times New Roman," identifying fonts, or numbers used to define sizes, locations, flags, and indices, or code blocks that can be a script or encrypted code. An attacker would likely find it convenient and safe to embed their malicious code in the data section or padding areas of Microsoft Word or any other suitable word processing program.

[0075] In addition to the pointers and data, the ITable sector can include a particular sequence of characters which, in some cases, are keywords or special values that cannot be changed without crashing the word processing program (e.g., numbers that indicate the beginning of a document or a sector, streams of Microsoft® Office® reference schemas, and/or numbers that indicate the end of structures). It should be noted that these values are not arbitrarily modified.

[0076] In accordance with some embodiments of the disclosed subject matter, arbitrary data transformation mechanisms can be provided that are designed to detect stealthy embedded attacks that may use mimicry strategies to evade the above-mentioned static and dynamic detection and analysis approaches. More particularly, some embodiments can modify and damage potentially embedded malicious code or any other suitable data segments. For example, in some embodiments, the arbitrary data transformation mechanisms can arbitrarily change the value of one or more data segments in a document in such a way as to purposely damage any hidden malicious code that may be embedded in one or more sections of the document. Consequently, the embedded malicious code not only fails, but also introduces a system exception or error (e.g., an application crash, a recognizable error state, etc.) that can be detected with little difficulty.

[0077] In addition, it should be noted that altering normal data used to describe the content of the document would not crash the application (e.g., a word processing program) in cases where no malicious code existed. The display of the document can be different or distorted, but the application does not crash. For example, the characters, font size, and type are changed in the document, but the word processing program can display the document without causing exceptions.

[0078] It should also be noted that these mechanisms do not require training. For example, these arbitrary data transformation mechanisms may not require the use of benign training files, malicious training files, testing files, etc. Instead, the arbitrary data

transformation mechanisms disables potentially embedded malicious code by altering vulnerable data values.

[0079] It should further be noted that, in some embodiments, these arbitrary data transformation mechanisms can be applied in a safe sandbox environment or virtual machine, where such a transformation can be reversed after testing the document with the altered or transformed data.

[0080] These mechanisms can be used in a variety of applications. For example, these mechanisms can be incorporated into the functionality of an application (e.g., Microsoft® Word® or another suitable word processing program), where the testing is handled within the application. In another example, these mechanisms can be added as a security feature to enhance the safety of Microsoft® Word® documents. The arbitrary data transformation mechanisms can be used as a safety check when opening Microsoft® Word® documents, where the mechanisms can be applied within Word in a sandbox environment or virtual machine or by a separate analysis application, such as the detection mechanisms described above or a third party antivirus scanner. If the document is deemed benign, the document can then be opened by the word processing program in its normal fashion after reversing the arbitrary data transformation.

[0081] Methods, systems, and media for detecting malicious code in electronic documents are provided. In some embodiments, the methods, systems, and media of the disclosed subject matter can detect malicious code in electronic documents by arbitrarily changing data in at least a portion of the document to produce an altered electronic document and then determine whether the altered electronic document produces an error (e.g., an application crash, a recognizable error state, etc.) when processed. For example, as illustrated in FIG. 23, a data segment in a portion of a document can be selected at 2310. The data segment can be selected from a data section or padding area of the document (e.g., the ITable sector shown in FIG. 22). In some embodiments, the data segment can be arbitrarily selected from multiple data segments in the data section and/or padding area of the document.

[0082] At 2320, in some embodiments, it can be determined whether the selected data segment in the portion of the document is capable of being changed. For example, in addition to pointers and data, the ITable sector and/or other sectors of the document can include a particular sequence of characters which, in some cases, are keywords or special values that cannot be changed without crashing the word processing program (e.g., numbers that indicate the beginning of a document or a sector, streams of Microsoft® Office®

reference schemas, and/or numbers that indicate the end of structures). These values are not arbitrarily modified.

[0083] In some embodiments, it can be determined whether the selected data segment in the portion of the document is a pointer. It should be noted that embedding malicious code in pointers can be difficult. For example, if the malicious code is created as a pointer value, the attacker has to ensure that the document does not cause an error in the word processing program before the malicious code is launched. That is, arbitrarily replacing a pointer value with another would likely crash the word processing program when it processes that pointer. In another example, pointers are generally short (e.g., from one bit to a few bytes) and are generally not contiguous. Because of this, there may not be sufficient space to place the malicious code in a pointer area without overwriting the data values or padding values adjacent to the pointer. Accordingly, in some embodiments, pointer values can be excluded from being arbitrarily changed.

[0084] At 2330, the values of the selected data segment in the portion of the document can be arbitrarily changed. For example, the values of the data segment can be arbitrarily changed to arbitrarily chosen different values (for non-zero data values that can be safely changed). In response, the binary content of illegally embedded code is changed to arbitrary values and, when processed, forces it to crash or to be disabled. That is, the execution of buffer overflow or other vulnerabilities is not avoided. Instead, the malicious code is damaged such that it purposely crashes or creates an error state.

[0085] Alternatively, normal data (data without malicious code) used to describe the documents can be changed to other values without serious error. It should be noted that the display of the altered document can be changed and can, in some cases, even appear incomprehensible. However, the word processing program or any other suitable program processing the document does not crash when malicious code is not present.

[0086] In some embodiments, for the byte values that can be changed, those data values can be increased or decreased by an arbitrarily chosen displacement x (e.g., changing the character "A" to "Q," increasing numerical values by 2, etc.). For example, the value of x can range from 1 to 3 (or -3 to -1), so a y -byte long data section has $|2x^y|$ possible permutations, where x can be positive or negative. The range of x is 256, which is the total possible values of a byte.

[0087] Alternatively, the value displacements can be changed to specific values that are defined according to the type of the structure using that data. In doing this, the display of documents may not be damaged and can be transformed into another style or type.

[0088] It should be noted that the arbitrary data transformation mechanisms described herein can transform a different data segment each time and, because each single byte has 256 possible values and an n-byte data sequence has 256^n possible permutations, an attacker cannot predict or guess what transformation is applied to a document's data sections. By changing the byte values in a random non-predictable fashion, the attacker cannot predict or guess what transformation is being applied to the data sections of a particular document. For example, if the arbitrarily data transformation mechanism is performed whenever a document is opened, the changed byte values can be different each time a document is opened. Accordingly, it is difficult to continuously guess the transformation using a brute force approach. Each attempted probe requiring the opening of a document produces a newly transformed instance of the document, thereby creating moving targets that thwart brute force guessing approaches.

[0089] In some embodiments, in response to arbitrarily changing the values of the data segment, the display of the document can be distorted. For example, the font display for ASCII data used in the ITable sector can appear with the "Times New Roman" font name, where the corresponding byte values are "54 69 6D 65 73 20 4E 65 77 20 52 6F 6D 61 6E." These data values are the data that describe the text type font. Arbitrarily changing one of these values to another value, including the extended ASCII characters, does not cause a word processing program, such as Microsoft® Word®, to crash. Instead, the word processing program generally chooses a default font or style if the transformed values were deemed incomprehensible.

[0090] Referring back to FIG. 23, at 2340, it is determined whether the corresponding program (e.g., the word processing program) produces an error when the altered electronic document is processed by the corresponding program. For example, these arbitrary data transformation mechanisms can be applied in a safe sandbox environment or a virtual machine.

[0091] FIG. 24 is an illustration of a screenshot of two documents before and after applying the arbitrary data transformation mechanism in accordance with some embodiments of the disclosed subject matter. As shown, the data representing the text type 2415 (i.e., FFN or Font Family Name) is arbitrarily changed. It should be noted that the altered document

can be opened and displayed without any error. Alternatively, in some embodiments, the word processing program displays a blank page or displays some strange sequence of characters when malicious code is not embedded in the document. If the document is deemed benign, the document can be opened by the word processing program in its normal fashion after reversing the arbitrary data transformation.

[0092] On the other hand, stealthy embedded malicious code residing in the data portion is also changed, and, subsequently, either the word processing program crashes or the malicious code is disabled when an attempt is made to execute the malicious code. For example, the hexadecimal Opcode values "6A" and "EB" represent the push and jmp x86 instructions, respectively. If the byte values are increased by 1, they become "6B" and "EC," which are not correct Opcodes. While there can be some instances in which the changed code is valid, it can become another completely unintended instruction. Accordingly, the corresponding program (e.g., a word processing program, an operating system, etc.) is not able to correctly execute the attackers' shellcode and either crashes or terminates the process.

[0093] FIGS. 25 and 26 are illustrations showing examples of the detection of malicious code embedded in electronic documents in accordance with some embodiments of the disclosed subject matter. In particular, the binary content of two buffer overflow attacks is shown. In FIG. 25, the list of 4-byte values is a list of pointers (the fcPlfLfo structure), and the zeros are padding values. The highlighted values 2510 are embedded shellcode, which is an infinite loop causing memory corruption. In FIG. 26, the structure is used to describe the Font Family Name (FFN). While the pointer values are not changed, the byte values of the malicious code embedded in the padding area is changed using the arbitrary data transformation mechanisms and, accordingly, the attack is disabled. Whether the shellcode exhibits obvious or mimicry behavior, the arbitrary data transformation mechanisms are capable of detecting the malicious code.

[0094] It should be noted that the word processing program generally does not crash. In some embodiments, the word processing program or any other suitable program displays a message indicating that the document may be damaged. FIG. 27 is an illustration showing two examples of such messages. For example, pop-up window 2710 indicates that the document name or path is not valid. In another example, pop-up window 2720 indicates that the document may be corrupt as the word processing program was unable to read and/or open the document. These error states can include, for example, the inability to open the

document, the inability to locate a document, entering an infinite loop causing memory corruption, etc.

[0095] These arbitrary data transformation mechanisms can be used in a variety of applications. For example, these mechanisms can be incorporated into the functionality of a word processing program (e.g., Microsoft® Word® or another suitable word processing program), where the testing is handled within the application. It should be noted that, in some embodiments, sections parsed in a proprietary format can be directly tested within the application program without developing an emulator. For example, the Microsoft Office Binary File Format Documentation describes the document format to parse the binary into the correct sectors. By performing the arbitrary data transformation mechanisms within a word processing program or any other suitable application program, the malicious code detection can be performed (e.g., using an internal emulator or sandbox environment) without providing knowledge of the specific proprietary application.

[0096] Alternatively, these mechanisms can be added a security feature to enhance the safety of Microsoft® Word® documents. The arbitrary data transformation mechanisms can be used as a safety check when opening Microsoft® Word® documents, where the mechanisms can be applied within Word in a sandbox environment or virtual machine or by a separate analysis application, such as the detection mechanisms described above or a suitable third party antivirus scanner.

[0097] Accordingly, methods, systems, and media for detecting malicious code in electronic documents are provided.

[0098] Although the invention has been described and illustrated in the foregoing illustrative embodiments, it is understood that the present disclosure has been made only by way of example, and that numerous changes in the details of implementation of the invention can be made without departing from the spirit and scope of the invention, which is limited only by the claims that follow. Features of the disclosed embodiments can be combined and rearranged in various ways within the scope and spirit of the invention. For example, with reference to FIG. 6, in some embodiments, detection model 45 can store information regarding messages (e.g., pop-up windows) and can be used at, for example, test 650.

What is claimed is:

1. A method for detecting malicious code in electronic documents, the method comprising:
 - selecting a data segment in at least one portion of an electronic document;
 - determining whether the arbitrarily selected data segment can be altered without causing the electronic document to result in an error when processed by a corresponding program;
 - in response to determining that the arbitrarily selected data segment can be altered, arbitrarily altering the data segment in the at least one portion of the electronic document to produce an altered electronic document; and
 - determining whether the corresponding program produces an error state when the altered electronic document is processed by the corresponding program.
2. The method of claim 1, wherein the electronic document is a word processing document.
3. The method of claim 1, wherein the corresponding program is a word processing program.
4. The method of claim 1, wherein the at least one portion of the electronic document is altered by changing values of the data segment in the at least one portion by a given value.
5. The method of claim 1, wherein the at least one portion of the electronic document is altered by changing values of the data segment in the at least one portion by an arbitrarily selected displacement.
6. The method of claim 1, wherein the error is at least one of: a crash of the corresponding program and a recognizable error state.
7. The method of claim 1, wherein the altered electronic document is processed in a sandbox environment.

8. The method of claim 1, further comprising disabling the malicious code in response to the produced error state.

9. A system for detecting malicious code in electronic documents, the system comprising:

at least one digital processing device that:

selects a data segment in at least one portion of an electronic document;

determines whether the arbitrarily selected data segment can be altered without causing the electronic document to result in an error when processed by a corresponding program;

in response to determining that the arbitrarily selected data segment can be altered, arbitrarily alters the data segment in the at least one portion of the electronic document to produce an altered electronic document; and

determines whether the corresponding program produces an error state when the altered electronic document is processed by the corresponding program.

10. The system of claim 9, wherein electronic document is a word processing document.

11. The system of claim 9, wherein the corresponding program is a word processing program.

12. The system of claim 9, wherein the at least one digital processing device is further configured to alter the at least one portion of the electronic document by changing values of the data segment in the at least one portion by a given value.

13. The system of claim 9, wherein the at least one digital processing device is further configured to alter the at least one portion of the electronic document by an arbitrarily selected displacement.

14. The system of claim 9, wherein the error is at least one of: a crash of the corresponding program and a recognizable error state.

15. The system of claim 9, wherein the at least one digital processing device is further configured to process the altered electronic document in a sandbox environment.

16. The system of claim 9, wherein the at least one digital processing device is further configured to disable the malicious code in response to the produced error state.

17. A computer-readable medium containing computer-executable instructions that, when executed by a processor, cause the processor to perform a method for detecting malicious code in electronic documents, the method comprising:

selecting a data segment in at least one portion of an electronic document;

determining whether the arbitrarily selected data segment can be altered without causing the electronic document to result in an error when processed by a corresponding program;

in response to determining that the arbitrarily selected data segment can be altered, arbitrarily altering the data segment in the at least one portion of the electronic document to produce an altered electronic document; and

determining whether the corresponding program produces an error state when the altered electronic document is processed by the corresponding program.

18. The computer-readable medium of claim 17, wherein the electronic document is a word processing document.

19. The computer-readable medium of claim 17, wherein the corresponding program is a word processing program.

20. The computer-readable medium of claim 17, wherein the at least one portion of the electronic document is altered by changing values of the data segment in the at least one portion by a given value.

21. The computer-readable medium of claim 17, wherein the at least one portion of the electronic document is altered by changing values of the data segment in the at least one portion by an arbitrarily selected displacement.

22. The computer-readable medium of claim 17, wherein the error is at least one of: a crash of the corresponding program and a recognizable error state.

23. The computer-readable medium of claim 17, wherein the altered electronic document is processed in a sandbox environment.

24. The computer-readable medium of claim 17, wherein the method further comprises disabling the malicious code in response to the produced error state.






FIG. 1

211		210		211	
_VBA_PROJECT	Data	Module1	Pictures		
OTable	Equation Native	MyMap	PRINT		
1Table	Events	NewMacros	PROJECT		
AttachContents	GeneralMacroStream	ObjInfo	PSMclsIstTable		
AttachDesc	Header	OCXNAME	SummaryInformation		
AttachPres	Image Contents	Ole	ThisDocument		
autoOpen	Internal	Ole10ItemName	ThisWorkbook		
CDraw	Macros	Ole10Native	VisioDocument		
CompObj	MailAttachment	OlePres000	VisioInformation		
CONTENTS	MailEnvelope	Payload	WordDocument		
Contents	MailStream	PIC	WordProData		
CoreIDRAW	META	PICT	Workbook		

FIG. 2







FIG. 5





FIG. 5A

6/28






FIG. 7




FIG. 8

9/28




FIG. 9




FIG. 10A




FIG. 10B

12/28




FIG. 11






FIG. 12

14/28








FIG. 15




FIG. 16




FIG. 17




FIG. 18




FIG. 19





FIG. 20







FIG. 22

**FIG. 23**

 %7Eanhabib%2F104.doc - Microsoft Word File Edit View Insert Format Tools Table Window Help       Heading 1 Times New Roman 12 B 1 2 3	Traditions and Cultures 104- Lecture 851 Winter 2000/2001 Justice and Virtue Description This course will introduce you to some of the central and political philosophy through the works of some of the western tradition. We will discuss <u>The Republic</u> by Plato, <u>Niccolo Machiavelli</u> (1469-1527), <u>Leviathan</u> by Thomas Hobbes, <u>Second Treatise on Civil Government</u> and <u>A Letter Concerning Human Faculties</u> (1632-1704), <u>The Social Contract</u> and <u>A Discourse on the Origin and Foundation of the Inequality of Mankind</u> by Jean-Jacques Rousseau (1712-1778), <u>Contract</u> by David Hume (1711-1776). There are three aims of this course. The first aim is to provide philosophical perspectives on moral and political topics.	Traditions and Cultures 104- Lecture 851 Winter 2000/2001 Justice and Virtue Description This course will introduce you to some of the central and political philosophy through the works of some of the western tradition. We will discuss <u>The Republic</u> by Plato, <u>Niccolo Machiavelli</u> (1469-1527), <u>Leviathan</u> by Thomas Hobbes, <u>Second Treatise on Civil Government</u> and <u>A Letter Concerning Human Faculties</u> (1632-1704), <u>The Social Contract</u> and <u>A Discourse on the Origin and Foundation of the Inequality of Mankind</u> by Jean-Jacques Rousseau (1712-1778), <u>Contract</u> by David Hume (1711-1776). There are three aims of this course. The first aim is to provide philosophical perspectives on moral and political topics.	Traditions and Cultures 104- Lecture 851 Winter 2000/2001 Justice and Virtue Description This course will introduce you to some of the central and political philosophy through the works of some of the western tradition. We will discuss <u>The Republic</u> by Plato, <u>Niccolo Machiavelli</u> (1469-1527), <u>Leviathan</u> by Thomas Hobbes, <u>Second Treatise on Civil Government</u> and <u>A Letter Concerning Human Faculties</u> (1632-1704), <u>The Social Contract</u> and <u>A Discourse on the Origin and Foundation of the Inequality of Mankind</u> by Jean-Jacques Rousseau (1712-1778), <u>Contract</u> by David Hume (1711-1776). There are three aims of this course. The first aim is to provide philosophical perspectives on moral and political topics.	Traditions and Cultures 104- Lecture 851 Winter 2000/2001 Justice and Virtue Description This course will introduce you to some of the central and political philosophy through the works of some of the western tradition. We will discuss <u>The Republic</u> by Plato, <u>Niccolo Machiavelli</u> (1469-1527), <u>Leviathan</u> by Thomas Hobbes, <u>Second Treatise on Civil Government</u> and <u>A Letter Concerning Human Faculties</u> (1632-1704), <u>The Social Contract</u> and <u>A Discourse on the Origin and Foundation of the Inequality of Mankind</u> by Jean-Jacques Rousseau (1712-1778), <u>Contract</u> by David Hume (1711-1776). There are three aims of this course. The first aim is to provide philosophical perspectives on moral and political topics.
---	---	---	---	---

FIG. 24

2500

xxxglobal.txt - WordPad			
File	Edit	View	Insert
Format	Help		
0000C990	11 00 00 00 00 00 00 00	00 00 00 9C 69 1A 5A 00	11Z.
0000C9A0	11 00 00 00 00 00 00 00	00 00 00 5A 6E 29 54 00	11ZN]T.
0000C9B0	11 00 00 00 00 00 00 00	00 00 00 4D 2D BC 52 00	11M-R.
0000C9C0	11 00 00 00 00 00 00 00	00 00 00 96 OD F9 2F 00	11I.
0000C9D0	11 00 00 00 00 00 00 00	00 00 00 93 6A 00 62 00	11j.b.
0000C9E0	11 00 00 00 00 00 00 00	00 00 00 50 56 A1 63 00	11PV.C.
0000C9F0	11 00 00 00 00 00 00 00	00 00 00 8D 08 F3 76 00	11V.
0000CA00	11 00 00 00 00 00 00 00	00 00 00 89 2C 90 31 00	11I.
0000CA10	11 00 00 00 00 00 00 00	00 00 00 CE 06 9C 2C 00	11Z.
0000CA20	11 00 00 00 00 00 00 00	00 00 00 18 79 9D 41 00	11y.A.
0000CA30	11 00 00 00 00 00 00 00	00 00 00 D4 11 B4 26 00	11&.
0000CA40	11 00 00 00 00 00 00 00	00 00 00 53 5D 68 24 00	11S)H\$.
0000CA50	11 00 00 00 00 00 00 00	00 00 00 AC 12 36 78 00	116x.
0000CA60	11 00 00 00 00 00 00 00	00 00 00 BD 1A 3A 71 00	11q.
0000CA70	11 00 00 00 00 00 00 00	00 00 00 05 38 7E 6E 00	118-n.
0000CA80	11 00 00 00 00 00 00 00	00 00 00 51 35 02 66 00	11Q5.f.
0000CA90	11 00 00 00 00 00 00 00	00 00 00 1E 68 8E 13 00	11h...
0000CAA0	11 00 00 00 00 00 00 00	00 00 00 4D 01 48 90 6E	11MMH.
0000CAB0	11 00 00 00 00 00 00 00	00 00 00 00 00 00 00 00	11h...
0000CAC0	11 00 00 00 00 00 00 00	00 00 00 FE 59 F7 15 00	11Y...
0000CAD0	11 00 00 00 00 00 00 00	00 00 00 1C 19 86 45 00	11E.
0000CAE0	11 00 00 00 00 00 00 00	00 00 00 6F 57 2A 01 00	11OW*.
0000CAF0	11 00 00 00 00 00 00 00	00 00 00 6E 4E 90 16 00	11nN...
0000CBB0	11 00 00 00 00 00 00 00	00 00 00 04 10 C6 71 00	11q.
For Help, press F1		NUM	

2510

FIG. 25




FIG. 26

2700





FIG. 27

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 09/56623

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 11/00 (2010.01)

USPC - 726/22

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)

IPC(8): G06F 11/00 (2010.01)

USPC: 726/22

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

USPC: 726/22-25; 717/124, 126, 127, 131, 141; 714/100, 1, 25, 38; 713/187-189 (keyword limited; terms below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

PubWEST (PGPB,USPT,EPAB,JPAB); USPTO AppFT, Google Scholar, Google Patents

electronic digital document media text detect discover find identif expos illegal malicious virus worm intru anomal code program software instruction select choos data string segment word byte determin check test alter transform chang modif application process run execut

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2008/0215891 A1 (HORNE et al.) 04 September 2008 (04.09.2008), entire document, especially Abstract; Fig. 1A, para [0010], [0052], [0054], [0056], [0231]	1-24
Y	US 2007/0234070 A1 (HORNING et al.) 04 October 2007 (04.10.2007), entire document, especially Abstract; Fig. 1, 31A, para [0076], [0501], [0536], [0623]	1-24
Y	US 2008/0010538 A1 (SATISH et al.) 10 January 2008 (10.01.2008), entire document, especially Abstract; Fig. 3, para [0040]	8, 16, 24
A	US 2008/0178299 A1 (MERKLE et al.) 24 July 2008 (24.07.2008), entire document	1-24
A	US 7,236,610 B1 (LUO et al.) 26 June 2007 (26.06.2007), entire document	1-24

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


* 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

15 April 2010 (15.04.2010)

Date of mailing of the international search report

22 APR 2010

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774