



- (51) **International Patent Classification:**
G06F 11/00 (2006.01)
- (21) **International Application Number:**
PCT/US2013/068451
- (22) **International Filing Date:**
5 November 2013 (05.11.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/803,029 18 March 2013 (18.03.2013) US
- (71) **Applicant: THE TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK** [US/US]; MC 4308, 535 West 116th Street, New York, New York 10027 (US).
- (72) **Inventors; and**
- (73) **Applicants (for US only): SETHUMADHAVAN, Lakshminarasimhan** [IN/US]; Apt. 10C, 560 Riverside Drive, New York, New York 10027 (US). **DEMME, John** [US/US]; MC0401, 1214 Amsterdam Avenue, New York, New York 10027 (US). **SCHMITZ, Jared** [US/US]; MC0401, 1214 Amsterdam Avenue, New York, New York 10027 (US). **TANG, Adrian** [US/US]; MC0401, 1214 Amsterdam Avenue, New York, New York 10027 (US). **STOLFO, Sal** [US/US]; MC0401, 1214 Amsterdam Avenue, New York, New York 10027 (US). **MAYCOCK, Matthew** [US/US]; MC0401, 1214 Amsterdam Avenue, New York, New York 10027 (US).
- (74) **Agents: RABINOVITCH, Ido et al.;** Hunter Clark PLLC, Suite 213T, 900 Cummings Center, Beverly, Massachusetts 01915 (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))

(54) **Title:** DETECTION OF ANOMALOUS PROGRAM EXECUTION USING HARDWARE-BASED MICRO-ARCHITECTURAL DATA




FIG. 2

(57) **Abstract:** Disclosed are devices, systems, apparatus, methods, products, media and other implementations, including a method that includes obtaining hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for a hardware device executing one or more processes, and determining based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process. In some embodiments, determining based on the hardware-based micro-architectural data whether the at least one of the one or more processes corresponds to a malicious process may include applying one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process.

DETECTION OF ANOMALOUS PROGRAM EXECUTION USING HARDWARE-BASED MICRO-ARCHITECTURAL DATA

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of, and priority to, provisional U.S. application Serial No. 61/803,029, entitled “SYSTEMS AND METHODS TO DETECT ANOMALOUS PROGRAM EXECUTION USING PROCESSOR MICROARCHITECTURAL EVENTS,” and filed March 18, 2013, the content of which is incorporated herein by reference in its entirety.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

[0002] This invention was made with government support under FA 8750-10-2-0253 awarded by the Defense Advanced Research Projects Agency (DARPA). The government has certain rights in the invention.

BACKGROUND

[0003] The proliferation of computers in a particular domain is generally followed by the proliferation of malicious processes (e.g., malware) in that domain. For example, systems that include the latest Android devices are laden with viruses, rootkits spyware, adware and other classes of malicious processes. Despite the existence of anti-virus software, malware threats (as well as threats from other types of malicious processes) persist and are growing. Unfortunately, there exist myriad ways to subvert commercial anti-virus software, including simply disabling the anti-virus. Furthermore, malware can mutate into new variants, which makes static detection of malware difficult.

[0004] Examples of some common malware processes are provided below:

Malware	Brief Description
Worm	Malware that propagates itself from one infected host to other hosts via exploits available on the surface (system call interfaces)

	of the operating system.
Virus	Malware that attaches itself to running programs and spreads itself through users' interactions with various systems.
Polymorphic Virus	A virus that, when replicating to attach to a new target, alters its payload to evade detection, i.e., takes on a different shape but performs the same function.
Metamorphic Virus	A virus that, when replicating to attach to a new target, alters both the payload and functionality, including the framework for generating future changes.
Trojan	Malware that masquerades as non-malware and acts maliciously once installed (opening backdoors, interfering with system behavior, etc.)
AdWare	Malware that forces the user to deal with unwanted advertisements.
SpyWare	Malware that secretly observes and reports on users computer usage and personal information accessible therein.
Botnet	Malware that employs a user's computer as a member of a network of infected computers controlled by a central malicious agency.
Rootkit	A malware package that exploits security holes in the operating system to gain superuser access. Usually, a rootkit attempts to hide its existence while performing malicious superuser activities by tampering with the file system.

[0005] Malicious processes, such as malware, were originally created to attain notoriety or for fun, but today malware deployment is mostly motivated by economic gains. There are reports of active underground markets for personal information, credit cards, logins into sensitive machines in the United States, etc. Also, malicious processes such as malware have been developed to target specific computers for industrial espionage purposes and/or for sabotage.

SUMMARY

[0006] The devices, systems, apparatus, methods, products, media and other implementations disclosed herein include a method including obtaining hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for a hardware device executing one or more processes, and determining based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process.

[0007] Embodiments of the method may include at least some of the features described in the present disclosure, including one or more of the following features.

[0008] Obtaining the hardware-based micro-architectural data may include obtaining the hardware-based micro-architectural data at various time instances.

[0009] Obtaining the hardware-based micro-architectural data at the various time instances may include performing one or more of, for example, a data push operation initiated by the hardware device to send the micro-architectural data, and/or a data pull operation, initiated by an antivirus engine, to send the micro-architectural data.

[0010] Obtaining the hardware-based micro-architectural data may include obtaining multi-core hardware-based micro-architectural data resulting from execution of the one or more processes on a processor device with multiple processor cores, and correlating the respective hardware-based micro-architectural data obtained from each of the multiple processor cores to the one or more processes.

[0011] Determining based on the hardware-based micro-architectural data whether the at least one of the one or more processes corresponds to a malicious process may include applying one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process.

[0012] Applying the one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process may include matching the obtained hardware-based micro-architectural

data to previously identified patterns of hardware-based micro-architectural data associated with one or more malicious processes.

[0013] The method may further include obtaining updates for the previously identified patterns of hardware-based micro-architectural data associated with the one or more malicious processes.

[0014] Obtaining the updates may include downloading encrypted data for the previously identified patterns of hardware-based micro-architectural data associated with the one or more malicious processes to an antivirus engine in communication with the hardware device providing the hardware-based micro-architectural data, decrypting at the antivirus engine the downloaded encrypted data for the previously identified patterns of hardware-based micro-architectural data associated with the one or more malicious processes, and updating a revision counter maintained by the antivirus engine indicating a revision number of a most recent update of the previously identified patterns of hardware-based micro-architectural data.

[0015] The one or more machine learning procedures may include one or more of, for example, a k-nearest neighbor procedure, a decision tree procedure, a random forest procedure, an artificial neural network procedure, a tensor density procedure, and/or a hidden Markov model procedure.

[0016] The malicious process may include one or more of, for example, a malware process, and/or a side-channel attack process.

[0017] The hardware-based micro-architectural data may include one or more of, for example, processor load density data, branch prediction performance data, and/or data regarding instruction cache misses.

[0018] In some variations, a system is provided that includes a hardware device executing one or more processes, and an antivirus engine in communication with the hardware device. The antivirus engine is configured obtain hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for the hardware device executing the one or more processes, and determine based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process.

[0019] Embodiments of the system may include at least some of the features described in the present disclosure, including at least some of the features described above in relation to the method, as well as one or more of the following features.

[0020] The antivirus engine configured to obtain the hardware-based micro-architectural data may be configured to obtain the hardware-based micro-architectural data at various time instances.

[0021] The antivirus engine configured to obtain the hardware-based micro-architectural data at the various time instances may be configured to receive the micro-architectural data in response to one or more of, for example, a data push operation initiated by the hardware device, and/or a data pull operation initiated by the antivirus engine.

[0022] The antivirus engine configured to determine based on the hardware-based micro-architectural data whether the at least one of the one or more processes corresponds to a malicious process may be configured to apply one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process.

[0023] The antivirus engine configured to apply the one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process may be configured to match the obtained hardware-based micro-architectural data to previously identified patterns of hardware-based micro-architectural data associated with one or more malicious processes.

[0024] The antivirus engine may further be configured to obtain updates for the previously identified patterns of hardware-based micro-architectural data associated with the one or more malicious processes.

[0025] In some variations, a computer readable media storing a set of instructions executable on at least one programmable device is provided. The set of instructions, when executed, causes operations including obtaining hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for a hardware device executing one or more processes, and determining based, at least in part, on the hardware-based micro-architectural data whether

at least one of the one or more processes executing on the hardware device corresponds to a malicious process.

[0026] Embodiments of the computer readable media may include at least some of the features described in the present disclosure, including at least some of the features described above in relation to the method and the system.

[0027] In some variations, an apparatus is provided. The apparatus includes means for obtaining hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for a hardware device executing one or more processes, and means for determining based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process.

[0028] Embodiments of the apparatus may include at least some of the features described in the present disclosure, including at least some of the features described above in relation to the method, the system, and the computer readable media.

[0029] Unless defined otherwise, all technical and scientific terms used herein have the same meaning as commonly or conventionally understood. As used herein, the articles “a” and “an” refer to one or to more than one (i.e., to at least one) of the grammatical object of the article. By way of example, “an element” means one element or more than one element. “About” and/or “approximately” as used herein when referring to a measurable value such as an amount, a temporal duration, and the like, is meant to encompass variations of $\pm 20\%$ or $\pm 10\%$, $\pm 5\%$, or $+0.1\%$ from the specified value, as such variations are appropriate to in the context of the systems, devices, circuits, methods, and other implementations described herein. “Substantially” as used herein when referring to a measurable value such as an amount, a temporal duration, a physical attribute (such as frequency), and the like, is also meant to encompass variations of $\pm 20\%$ or $\pm 10\%$, $\pm 5\%$, or $+0.1\%$ from the specified value, as such variations are appropriate to in the context of the systems, devices, circuits, methods, and other implementations described herein.

[0030] As used herein, including in the claims, “or” or “and” as used in a list of items prefaced by “at least one of” or “one or more of” indicates that any combination of the listed items may be used. For example, a list of “at least one of A, B, or C” includes any of the combinations A or B or C or AB or AC or BC and/or ABC (i.e., A and B and C). Furthermore, to the extent more

than one occurrence or use of the items A, B, or C is possible, multiple uses of A, B, and/or C may form part of the contemplated combinations. For example, a list of “at least one of A, B, or C” may also include AA, AAB, AAA, BB, etc.

[0031] As used herein, including in the claims, unless otherwise stated, a statement that a function, operation, or feature, is “based on” an item and/or condition means that the function, operation, function is based on the stated item and/or condition and may be based on one or more items and/or conditions in addition to the stated item and/or condition.

[0032] Details of one or more implementations are set forth in the accompanying drawings and in the description below. Further features, aspects, and advantages will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0033] These and other aspects will now be described in detail with reference to the following drawings.

[0034] FIG. 1 includes example illustrations of graphs of hardware micro-architectural activity for several different processes.

[0035] FIG. 2 is a schematic diagram of an example system to detect malicious processes

[0036] FIG. 3 is a flow chart of an example procedure to detect malicious processes.

[0037] FIG. 4 is a schematic diagram of an example system in which an AV engine is implemented.

[0038] FIG. 5 is a table of Android malware families, and detection results therefor, tested using, for example, the systems and procedures of FIGS 2-4.

[0039] FIG. 6 is a graph illustrating accuracy of various binary classifiers when applied to micro-architectural data produced, in part, by Android malware.

[0040] FIG. 7 includes graphs showing the accuracy of classifiers when applied to rootkits processes.

[0041] FIG. 8A is a schematic diagram of an example security update payload.

[0042] FIG. 8B is a flowchart of an example procedure to receive a security update payload and update the configuration of an AV engine.

[0043] Like reference symbols in the various drawings indicate like elements.

DESCRIPTION

[0044] Described herein are systems, devices, apparatus, computer program products, and other implementations for detection of anomalous program execution processes, such as malware. In some implementations, hardware-based micro-architectural data, including hardware-based micro-architectural counter data (e.g., from hardware-based performance counters) is obtained from a hardware device (such as a processor / computing-based device), and analyzed (e.g., to analyze the temporal behavior of executing processes that resulted in the micro-architectural data) using machine-learning procedures (e.g., classification procedures), to identify malicious processes from one or more processes executing on the hardware device being monitored

[0045] Generally, processes executing on a hardware-implemented controller device (be it a general-purpose processor, an application-specific controller, etc.) exhibit phase behavior. A process (be it a malicious or non-malicious process) that is configured to achieve a particular functionality may perform activity A for a while, then switch to activity B, then to activity C. Although such a process may alternate in the exact order of performance of the activities, typically the process would need to perform activities A, B, and C to accomplish its particular functionality. The activity phases that correspond to a particular process typically correspond to patterns in architectural and micro-architectural events. Additionally, different processes (e.g., different programs configured to perform different functionalities) result in different hardware-based micro-architectural behavior. For example, FIG. 1, providing example illustrations of micro-architectural activity graphs for several different processes, show that the micro-architectural behavior for different processes (in the example of FIG. 1, the processes were from the SPEC benchmark suite) tend to be different, resulting in different hardware micro-architectural traces or patterns. For example, as shown in FIG. 1, the behavior for the L1 exclusive hits and the executed branch instructions monitored for the 'bzip2' process (as illustrated in graphs 102 and 104) is different from the behavior for the L1 exclusive hits and the

executed branch instructions for the ‘mcf’ process (illustrated in graphs 112 and 114), which in turn is different from the behavior for the L1 exclusive hits and the executed branch instructions for the ‘sjeng’ process (illustrated in graphs 122 and 124). It is to be noted that monitoring the micro-architectural behavior may be facilitated by specific built-in counters, and/or may be achieved by measuring / monitoring occurrences of event at specific points on the circuits of the hardware device being monitored.

[0046] Accordingly, processes executing on a hardware device (e.g., hardware-based controller device) may be distinguished (and thus identified) based on such time-varying micro-architectural signatures / traces. Generally, minor variations in the exact implementation of a particular process do not significantly affect the generated hardware-based micro-architectural traces resulting from the process, and, therefore, identifying the process and/or a determining whether the process is malicious or not (e.g., via machine learning classification procedures, heuristic and non-heuristic procedures analyzing the micro-architectural data, etc.) can still be performed. This is because regardless of how a malicious process (e.g., malware) writers change the underlying implementation (e.g., the software program), the semantics of the process do not change significantly. For instance, if a piece of malware is designed to collect and log GPS data, then no matter how its writer re-arranges the code, the process will still have to collect and log GPS data. In other words, the activity phases characterizing the process will generally remain regardless of the specific implementation of the process. Additionally, a particular task that needs to be accomplished will include various sub-tasks that cannot be significantly modified. For example, a GPS logger will always have to warm up the GPS, wait for signals, decode the data, log it, and, at some future point, exfiltrate the data back to the rogue user (privacy thief) seeking to obtain the data. As a result of these generally invariant operations required to accomplish particular tasks or processes, particular phases of the malicious process’ execution remain relatively invariant for different implementation variations.

[0047] Thus, hardware-based micro-architectural data (e.g., data from hardware performance counters) such as processor load density data, branch prediction performance data, data regarding instruction cache misses, etc., can be used to identify malware and/or other types of malicious processes. Experimental results (more particularly discusses below) show that the detection techniques/procedures described herein tend to be robust to variations in malware programs (or

other types of malicious processes). Thus, after examining a small set of variations within a family of malware on a processing platform (e.g., Android ARM and Intel Linux platforms), many variations within that family may be substantially accurately detected. Further, various implementations described herein enable malicious process detectors, such as the detectors described herein, to run securely beneath the system software, thus reducing, or all together avoiding, the danger of being turned off.

[0048] Accordingly, in some embodiments, methods, systems, devices, products, and other implementations are disclosed that include a method including obtaining hardware-based micro-architectural data, including, for example, hardware-based micro-architectural counter data, for a hardware device executing one or more processes, and determining based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the processor-based system corresponds to a malicious process. The malicious process being identified / detected may include one or more of, for example, a malware process, and/or a side-channel attack process.

[0049] With reference to FIG. 2, a schematic diagram of an example system 200 to detect and/or resolve malicious processes is shown. The system 200 includes an antivirus (AV) engine 210 that comprises, in some embodiments, a performance counter sampling unit (also referred to as a “sampler”) 212, a performance counter database 214 that stores / maintains representative micro-architectural profiles or signatures (including performance counter profiles or signatures) corresponding to various processes (including malware processes), and micro-architectural data collected by the sampling unit 212, and a classifier 216 configured to analyze the collected hardware micro-architectural data to determine if the one or more processes running on the hardware device being observed / monitored includes at least one malicious process (in some embodiments, the classifier 216 may also be configured to more particularly identify such a malicious process). The AV engine 210 is generally in communication with one or more hardware devices such as processor devices 220 and/or 222 shown in FIG. 2.

[0050] The sampling unit 212 is configured to obtain hardware-based micro-architectural data, including, for example, hardware-based micro-architectural performance counter data from the one or more hardware-devices, which may include devices such as controller devices, e.g., processor devices such as the devices 220 and 222, or any other type of controller devices

including controller devices implemented using modules such as an FPGA (field programmable gate array), an ASIC (application-specific integrated circuit), a DSP processor, etc. Generally, hardware-based controller devices include hardware-related performance counters that may be configured to count a variety of events such as cycles, instructions, cache misses, etc. In some implementations, these performance counters are used to assist in software performance optimization. For example, the Intel For x86 processor device implements four (4) configurable performance counters, and the OMAP4460 processor with dual ARM Cortex-A9 cores includes six (6) configurable performance counters. The AV engine 210 is implemented to obtain micro-architectural data (e.g., performance counter data) from known controller designs, and as such the AV engine 210 may be configured to obtain micro-architectural data from specific known performance counters particular to the hardware that is being monitored by the AV engine. That is, knowledge of the specific architecture of the hardware to be monitored may be required in order to obtain performance counter data and other micro-architectural data from the performance counters corresponding to the specific architecture. Examples of micro-architectural counters used on an Intel x86 processor architecture include:

- 0x0440 -- L1D_CACHE_LD.E_STATE;
- 0x0324 -- L2_RQSTS.LOADS;
- 0x03b1 -- UOPS_EXECUTED.PORT (1 or 2); and
- 0x7f88 -- BR_INST_EXEC.ANY.

[0051] Examples of common counters (feature event number assignments) on the ARM Cortex-A9 cores architecture, through which micro-architectural data can be obtained, include event numbers:

- 0x06 – Memory-reading instruction architecturally executed (counter increments for every instruction that explicitly read data);
- 0x07 – Memory-writing instruction architecturally executed (counter increments for every instruction that explicitly wrote data);
- 0x0C – Software change of PC, except by an exception, architecturally executed (counter does not increment for a conditional instruction that fails its condition code);

- 0x0D – Immediate branch architecturally executed (counter counts for all immediate branch instructions that are architecturally executed);
- 0x0F – Unaligned access architecturally executed (counter counts each instruction that is an access to an unaligned address); and
- 0x12 – Counter counts branch or other change in program flow that could have been predicted by the branch prediction resources of the processor.

[0052] Additional information on micro-architectural counters that may be implemented on the ARM Cortex-A9 cores architecture is provided, for example, at “ARM[®] Architecture Reference Manual, Arm[®] v7-A and ARM[®] v7-R edition, Errata markup,” the content of which is incorporated herein by reference in its entirety.

[0053] In some embodiments, the sampling unit 212 may be configured to obtain hardware micro-architectural data (including micro-architectural performance counter data) from the counters of the hardware monitored through data push procedures and/or through data pull procedures. For example, when pulling data, the AV engine 210 initiates the data collection, causing hardware targets (e.g., specific hardware performance counters implemented in the hardware being monitored) to be accessed by, for example, interrupting execution of the counters and/or querying the counters without interruptions. In some embodiments, the AV engine 210 may be configured, e.g., via the sampling module 212, to interrupt the hardware once every N cycles (where N may be a constant pre-determined number, or may be a varying number, e.g., based on a random or pseudo-random generator), and sample the various performance / event counters, as well as other values (e.g., the currently executing process' PID). When performing sampling operations using an interrupt-based procedure, the sampling unit 212 may be configured to send control signals or otherwise cause the executing hardware to be interrupted, access the performance counters and/or other storage hardware, and retrieve the values stored on the counters of the interrupted hardware for further processing by the AV engine 210. In some embodiments, upon interruption of the hardware and/or the counters, the interrupted hardware may first store data held by its various performance counters in a central storage location (e.g., in a state stack), and the data stored at the central storage location may then be accessed and retrieved by the sampling unit 212. When implementing a data-push sampling mode, data held by the performance counters (and/or other sampling points on the hardware being monitored)

may be configured to be communicated to the AV engine 210 (e.g., to the sampling unit 212) at regular or irregular intervals, with or without interrupting the execution of the hardware being monitored or of its performance counters. Thus, in such embodiments, the hardware device to be monitored is configured to initiate sending the micro-architectural data to the AV engine 210. For example, in a data push mode, the hardware device being monitored may be configured to send micro-architectural data without needing to receive a request (e.g., from the sampling unit 212).

[0054] The sampling operations implemented by the sampling unit 212 of the AV engine 210 thus obtain time-based data of the output of the various hardware performance counters (and/or other output points) monitored for one or more processes executing on the hardware being monitored. As noted, in addition to micro-architectural data, information such as a process' ID (e.g., PID) is also recorded to enable associating / correlating the micro-architectural data with the process whose execution resulted in the obtained micro-architectural data. By also recording processes' IDs and associating / correlating them with the obtained micro-architectural data, the implementations described herein can track micro-architectural data resulting from execution of a process across different hardware devices. For example, in situations where a system being monitored includes multiple processor cores (each with its own set of performance counters), where processes / threads may be suspend and resume execution on different cores, maintaining processes' PID's along with obtained micro-architectural data may enable tracking the behavior of processes as they switch execution to different hardware devices.

[0055] In some embodiments, the sampling unit 212 may be realized, at least in part, on the hardware device being monitored. For example, the sampling unit 212 may be implemented as a hardware realization on a specialized hardware-based controller such as an FPGA, an ASIC, etc.) In some embodiments, the micro-architectural database 212 may be realized, at least in part, as a software implementation executing on a machine that includes a processor-based device that is being monitored by the AV engine 210 to detect malicious processes that are executing on the machine. For example, one of a processor-device's multiple general-purpose cores may be allocated to execute a software realization of at least part of the AV engine.

[0056] As noted the AV engine 210 also includes a micro-architectural database 214 configured to store the micro-architectural data obtained from the hardware being monitored / observed, as

well as pre-determined data sets, obtained from remote nodes (e.g., servers), that include data representative of micro-architectural signatures / traces of known malicious processes (e.g., time-series traces for various micro-architectural events or performance counters) and training data that includes micro-architectural data (e.g., time-based data) for non-malicious / benign processes. As will be described below in greater details, in some embodiments, the AV engine 210 is periodically (at regular or irregular intervals) updated to include new or modified micro-architectural signature data defining the behavior of new or existing malicious processes by receiving from a remote node micro-architectural signature data.

[0057] In some embodiments, the database 214 may be realized, at least in part on the hardware device being monitored. In some embodiments, the micro-architectural database 214 may be realized, at least in part, as a software implementation executing on a machine that includes a processor-based device being monitored by the AV engine 210 (e.g., allocating one of a processor-device's multiple general-purpose cores to execute a software realization of the database 214).

[0058] Collection of micro-architectural data (including micro-architectural performance counter data) using, for example, the sampling unit 212, and/or storage of the collected data using, for example, the database 214, provides a relatively large amount of labeled data that includes micro-architectural data resulting from execution of malicious processes (e.g., malware) and non-malicious processes. Thus, in some embodiments, the classifier 216 (also referred to as a machine-learning engine) is configured to determine whether at least one of the processes with respect to which the micro-architectural data was collected corresponds to a malicious process (e.g., whether some of the micro-architectural data traces collected potentially resulted from execution of the at least one malicious process) and/or identify the at least one malicious process.

[0059] In some implementations, a classifier, such as the classifier 216 of the AV engine 210, may be configured to iteratively analyze training input data and the input data's corresponding output (e.g., a determination of a process type and/or identification of a process corresponding to the input data), and derive functions or models that cause subsequent micro-architectural inputs, collected from the hardware being monitored, to produce outputs consistent with the classifier's learned behavior. Such a classifier should be configured to distinguish malicious processes from non-malicious processes.

[0060] Generally, machine learning classifiers are configured to examine data items and determine to which of N groups (classes) each data item belongs. Classification procedures can produce a vector of probabilities, e.g., the likelihoods of the data item belonging to each class. In the case of malicious process detection, two classes may be defined: malicious process (e.g., malware) and non-malicious process (e.g., non-malware). As a result, the output from classifiers may include probabilities representing the likelihood of a data item being malicious. In situations where a particular classifier is not adapted to process / classify time-series data (like the time-series micro-architectural data collected by the AV engine 210) this difficulty can be overcome by arranging input data (e.g., corresponding to micro-architectural events occurring at a particular location of the hardware, such as at a particular counter) that occurred at different time instances into a single vector of features that is presented as input to the classifier. Under this approach, time-based data may be consolidated into a vector of data, where each vector point corresponds to a micro-architectural sample for a certain counter or location that occurred at a different time instance. Additionally and/or alternatively, another approach for processing time-dependent data (micro-architectural data) using classifiers that are generally not configured to handle sequences of time-dependent data is to separately process with such a classifier data points taken for a particular process at different time instances, and aggregate the classifier's results in order to classify the entire process. In some embodiments, different aggregation operations may be applied to a classifier's results, and the aggregation operation that is determined (e.g., through testing and experimentation) to yield the best classification results may be used to perform future aggregation operations. For example, one aggregation operation that may be used is a simple average operation. Another aggregation operation that may be used is a weighted average operation in which, for example, data points which are equally probable to belong to each of the various available classes are given zero weight, whereas data points with high probabilities are given relatively large weights.

[0061] The types of classifiers that may be used to process / analyze the collected micro-architectural data points corresponding to the executing processes belong to two main classifier categories: linear classifiers, and non-linear classifiers. Linear classification procedures are configured to attempt to separate n -dimensional data points by a hyperplane – points on one side of the plane are points of class X and points on the other side are of class Y . Non-linear

classifiers generally do not rely on this type of linear separation. Thus, any operation to derive a classification may be applied.

[0062] In some of the implementations described herein, non-linear classifiers were used to perform the data processing / analysis operations to reflect the fact that the data (e.g., micro-architectural data) that was used to determine whether at least one executing process may be malicious, or to identify a malicious process, may not necessarily be linearly-separable. Some examples of classifiers, configured to determine if a particular process (for which micro-architectural time-based data was collected) is malicious or non-malicious, that may be used with implementations of the AV engine 210 include:

- K-Nearest Neighbors (KNN) – A KNN classifier is trained by inserting the training data points along with their labels into a spatial data structure, like a k-dimensional tree (referred to as a “k-d-tree”) used for organizing points/data in a k-dimensional space. In order to classify a data point, that point’s k nearest neighbors (in Euclidean space) are found using the spatial data structure. The probability that the data point is of a particular class is determined by how many of the data point’s neighbors are of that class and how far they are from each other.
- Decision Tree – Another way to classify data points is to use a non-spatial tree called a decision tree. This tree is built by recursively splitting training data into groups on a particular dimension. The dimension and split points are chosen to minimize the entropy with each group. These decisions can also integrate some randomness, decreasing the quality of the tree but helping to prevent overtraining. After some minimum entropy is met, or a maximum depth hit, a branch terminates, storing in it the mix of labels in its group, e.g., 30% malware vs. 70% non-malware. To classify a new data point, the decision tree traverses the tree to find the new point’s group (leaf node), and returns the stored mix.
- Random Forest – One way to increase the accuracy of a classifier is to use a lot of different classifiers and combine the results. In a random forest, multiple decision trees are built using some randomness. When classifying a new data point, the results of all trees in the forest are weighted equally to produce a result.

- Artificial Neural Network (ANN) – A neural network machine attempts to model biological brains by including neurons which are connected to each other with various weights. The weight values between connections can be varied, thus enabling the neural network to adapt (or learn) in response to training data it receives. In feed-forward neural nets, input values are supplied at one edge and propagate through a cycle-less network to the output nodes. In some embodiments, one input neuron for each dimension, and two output nodes (e.g., one indicating the probability that malware is running, one indicating the probability that non-malware is running) are defined.
- Tensor Density – this classifier discretizes the input space into different buckets. Each bucket contains the mix of classes in the training data set. A data point is classified by finding its bin and returning the stored mix. Generally, a tensor density classifier uses $O(1)$ lookup time, and is thus considered to be time-efficient.

[0063] Other classifiers that may be used also include, in some embodiments, a support vector machine configured to generate, for example, classification functions or general regression functions. In some embodiments, the classifiers may be implemented using regression techniques to derive best-fit curves, a classification procedure based on hidden Markov model, and/or other types of machine learning techniques. In embodiments in which a hidden Markov model-based classifier is used, patterns in the data (e.g., micro-architectural data) being processed may be identified using self-similarity analysis, and the transitions in patterns may be used to build the hidden Markov model with which malware/goodware can be predicted / classified. Additionally, linear classification techniques like kernel methods which are capable of accurately classifying data but with reduced computational requirements may also be used.

[0064] To train the classifiers to identify suspected malicious processes based on micro-architectural data collected from a hardware-based device that is to be monitored, in some implementations, a remote system whose hardware configuration may be the same or similar to the hardware configuration of the hardware device with respect to which the procedures described herein are performed, may execute variants of a known malicious process (e.g., a malware discovered and/or tracked by some third party). Micro-architectural data resulting from execution of the variants of the particular malware (e.g., represented in a form that may be similar to the data used to generate graphs similar to those illustrated in FIG. 1) is collected.

Periodically, data representative of the micro-architectural data captured by the remote system may be communicated to the AV engine 210, and stored on the database 214. The remote system may also provide micro-architectural data corresponding to known non-malicious processes. The example micro-architectural data communicated by the remote system may be used to train the classifier 216 by providing that micro-architectural data and the respective identities and/or type (e.g., malicious or non-malicious) of the processes that caused that micro-architectural data to be produced to at least some of the one or more of the classifiers 216a-n. The training data will cause the classifiers 216a-n to be configured (e.g., dynamically configured) so that upon presenting similar micro-architectural data (collected from the hardware device to be monitored) to the classifiers, output consistent with the processes types/identities of the training data will be produced.

[0065] As noted, an AV engine, such as the AV engine 210 of FIG. 2, may be realized entirely in hardware (e.g., implemented as a module on the hardware device that is to be monitored), entirely in software (e.g., as a multi-module application executing on a computing system that includes the hardware to be monitored), or as a hardware-software combination implementation in which one component (e.g., the sampling unit 212 of FIG. 2) is implemented in hardware, while the database and classifier units 214 and 216 are implemented via software). If implemented at least partly by software, the software components may be configured to communicate with the hardware component (e.g., using an interfacing procedure) to receive data (e.g., micro-architectural data obtained by the sampling unit) and/or to transmit data or control signals to the hardware-based component.

[0066] In addition to being configured to collect and store micro-architectural data and analyze collected micro-architectural data to determine whether or not malicious behavior is occurring (and possibly more particularly identify the malicious process(es)), the AV engine 210 is also configured to take certain actions if a threat is detected (e.g., shut down the hardware or report the malicious behavior), and update the AV engine with malicious processes definitions and micro-architectural signatures. More particularly, there are a wide variety of security policies that can be implemented by an AV engine such as the AV engine 210. Some viable security policies include:

- Using the AV engine as a first-stage malware predictor – When the AV engine suspects a program to be malicious it can run more sophisticated behavioral analysis on the program. Hardware analysis happens ‘at speed’ and is significantly faster than behavioral analysis used by malicious process analysts to create signatures. Such pre-filtering can avoid costly behavioral processing for ‘goodware.’
- Migrating sensitive computation – In multi-tenant settings such as public clouds, when the AV engine suspects that an active thread on the system is being attacked (e.g., through a side-channel) the AV engine can move the sensitive computation. In some scenarios it may be acceptable for the AV system to simply kill a suspect process.
- Using the AV engine for forensics – Logging data for forensics is expensive as it often involves logging all interactions between the suspect process and the environment. To mitigate these overheads, the information necessary for forensics can be logged only when the AV engine suspects that a process is malicious.
- Screening for goodware – In some embodiments, the hardware-based micro-architectural data collected can be used to identify non-malicious processes, and to corroborate that those processes are in fact non-malicious. For example, in some implementations, underlying code samples of processes identified by the AV engine as non-malicious can be analyzed by, for example, comparing the code sample to known code samples (that were previously obtained) corresponding to the processes analyzed. If the examined underlying code of the executing processes matches the known code listing previously obtained, the executing process is confirmed as being non-malicious.

[0067] Thus, there is a broad spectrum of actions that can be taken based on the AV engine’s output. The systems and procedures described herein to implement an AV engine should be flexible enough to implement the above-described security policies. Conceptually, this means that, in some embodiments, the AV engine should be able to interrupt computation on any given core and run the policy payload on that machine. This requires the AV engine to be able to issue a non-maskable inter-processor interrupt. Optionally, in some embodiments, the AV engine can communicate to the OS or supervisory software that it has detected a suspect process so that the system can start migrating other co-resident sensitive computation. In some embodiments, the AV engine may also be configured to run in the highest privilege mode.

[0068] Additionally, as noted, in some embodiments, the AV engine 210 may be configured to be updated with new malware signatures as they become available, or when new classification techniques are implemented. The AV update should be implemented in a way to prevent attackers from compromising the AV. For instance, a rogue user should not be able to mute the AV engine or subvert the AV engine to create a persistent, high-privilege rootkit.

[0069] Generally, security updates may include one or more of, for example, a classifier, an action program that specifies security policies, a configuration file that determines which performance features are to be used with what classifiers, micro-architectural data for malicious and/or non-malicious processes, and/or an update revision number. This data can be delivered to the AV engine securely using techniques / procedures adapted for a hardware setting. A schematic diagram of an example security update payload 800 that is to be sent from a system security vendor, including the various encryption levels applied to the payload, is depicted in FIG. 8A. An example procedure 850, generally performed by an AV engine, to receive a security update payload (such as the encrypted payload 800) and update the configuration of the AV engine, is depicted FIG. 8B. As shown in the figure, the procedure 850 includes receiving 855 the payload, and decrypting 860 the payload with a “verif” key embedded in the hardware (on which the AV engine is implemented). A determination is then made 865 of whether a resulting hash of the “verif” matches the expected hash of the verif key embedded in the hardware. If it doesn’t, the procedure 850 terminates 870. If there is a match of the hash of the “verif” key, a determination is made 875 of the integrity of the payload with a SHA-2 hash function. If the integrity is confirmed, the payload is decrypted 885 with an AES key (otherwise, the procedure terminates 880), and upon a determination that the update revision number indicated in the payload is in agreement with a revision number indicator maintained in the hardware device (at 890), the updates in the payload are applied 895.

[0070] As indicated in relation to the operation 890 of the procedure 850, in some embodiments, the hardware device on which the AV engine is, at least partly, implemented, maintains the revision number of the last update, and that revision number is incremented on every update. This is to prevent / inhibit an attacker from rolling back the AV system, which an attacker might do to prevent the system from discovering new malicious processes. The AV engine may offer

this protection by rejecting updates with a revision number that is older than the revision number maintained in the hardware counter.

[0071] With reference now to FIG. 3, a flowchart of an example procedure 300 to detect malicious processes is shown. The procedure 300 includes obtaining 310 hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for a hardware device executing one or more processes. As noted, in some embodiments, obtaining the micro-architectural data may be performed by a sampling unit, which may be implemented, at least partly, in hardware as part of the hardware-device that is to be monitored (i.e., the hardware device executing the one or more processes with respect to which the micro-architectural data is to be collected). In some embodiments, the micro-architectural data may be obtained periodically at regular or irregular intervals (e.g., at intervals of length determined by a pseudo random process), and may be obtained through a data-pull (e.g., by the sampling unit initiating the collection of the micro-architectural data, with or without interruption the hardware device being observed) or through a data push process (e.g., the hardware device initiating periodic communication of micro-architectural data to an AV engine).

[0072] Based, at least in part, on the obtained hardware-based micro-architectural data, a determination is made 320 whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process. In some embodiments, a more specific determination may be made of the type or identity of the at least one of the one or more processes executing on the hardware device. As noted, determination of whether at least one process is malicious and/or the type and/or identity of the at least one of the one or more processes executing on the hardware device may be performed using a machine learning system that may include one or more classifiers (such as the one or more classifiers 216a-n) that were trained with training data including micro-architectural data for variants of known malicious and non-malicious processes.

[0073] Thus, when a variant of a known malicious process executes on the hardware device to be monitored, even in situations where the exact implementation of the malicious process has been modified, the malicious process will generally perform operations that are characteristic of that process (e.g., accessing particular modules, retrieving specific types of data, etc.). These operations that are characteristic of the known malicious process will result in a micro-

architectural data signature (that may be represented as a time series), which may then be identified (or at least identified as being malicious or non-malicious) through at least one of the one or more classifiers of the machine learning system of the AV engine.

[0074] With reference to FIG. 4, an example system 400 in which an AV engine (such as the AV engine 210 of FIG. 2) is implemented, is shown. The system 400 includes a hardware device such as controller device 410, which may be a processor-based personal computer, a specialized computing device, and so forth, and which includes, in some implementations, a processor-based unit such as central processor unit (CPU) 412. In some embodiments, the controller device 410 may be realized, at least in part, using modules such as an FPGA (field programmable gate array), an ASIC (application-specific integrated circuit), a DSP processor, etc.

[0075] As noted, in some embodiments, at least part of the AV engine may be implemented in hardware directly on the hardware device that is to be monitored, and/or may be implemented in software executing on a dedicated and secure controller device. For example, as depicted in FIG. 4, the CPU 412 may be a multi-core processor, and the hardware portion of the AV engine may thus be realized on one or more of the cores 413 of the CPU 412, and be configured (e.g., through pre- or post-manufacturing programming) to perform one or more of the functions of the AV engine (e.g., collect micro-architectural data). If the hardware device to be monitored is an application-specific controller device (e.g., implemented as an application-specific integrated circuit), the hardware-portion of the AV may be realized at the time of manufacturing of the controller, e.g., as a special-purpose malware detection units that sit on a network-on-chip, on-chip/off-chip FPGA, or off-chip ASIC co-processor. These choices represent different trade-offs in terms of flexibility and area- and energy-efficiency. Moving security protection to the hardware level solves several problems and provides some interesting opportunities. For example, it ensures that the security system cannot be disabled by software, even if the kernel is compromised. Second, because the security system runs beneath the operating system, the security system might be able to protect against kernel exploits and other attacks against the kernel. Third, because the hardware itself is being modified (to accommodate at least some portions of the AV engine), arbitrary static and dynamic monitoring capabilities can be added. This gives the security system extensive viewing capabilities into software behavior.

[0076] As further shown in FIG. 4, in addition to the CPU 412 and/or other application-specific hardware to implement controller functionality, the system 400 includes main memory, cache memory and bus interface circuits (not shown in FIG. 4). For example, the controller device 410 may include a mass storage element 414, such as a hard drive or flash drive associated with the system. The computing system 400 may further include a keyboard, or keypad, or some other user input interface 416, and a monitor 420, e.g., a CRT (cathode ray tube), LCD (liquid crystal display) monitor, etc., that may be placed where a user can access them.

[0077] The controller device 410 is configured to facilitate, for example, the implementation of operations to obtain hardware-based micro-architectural data resulting from execution of one or more processes on the CPU 412 and/or on some other application-specific device on which processes are executing (or can be executed) and determine, based on the micro-architectural data obtained, whether at least one of one or more of the processes executing on the controller device 410 of the system 400 is a potentially malicious process (e.g., malware). In some embodiments, identities of the one or more processes executing on the hardware of the controller device 410 may be determined based on the micro-architectural data collected. The storage device 414 may thus include a computer program product that when executed on, for example, a processor-based implementation of the controller device 410 causes the device to perform operations to facilitate the implementation of procedures described, including the procedures to obtain micro-architectural data and determine based on that data whether at least one of the one or more executing processes is potentially malicious.

[0078] The controller device 410 may further include peripheral devices to enable input/output functionality. Such peripheral devices may include, for example, a CD-ROM drive and/or flash drive (e.g., a removable flash drive), or a network connection (e.g., implemented using a USB port and/or a wireless transceiver), for downloading related content to the connected system. Such peripheral devices may also be used for downloading software containing computer instructions to enable general operation of the respective system/device. As noted, alternatively and/or additionally, in some embodiments, special purpose logic circuitry, e.g., an FPGA (field programmable gate array), an ASIC (application-specific integrated circuit), a DSP processor, etc., may be used in the implementation of the system 400. Other modules that may be included with the controller device 410 are speakers, a sound card, a pointing device, e.g., a mouse or a

trackball, by which the user can provide input to the system 400. The controller device 410 may include an operating system, e.g., Windows XP® Microsoft Corporation operating system, Ubuntu operating system, etc.

[0079] Computer programs (also known as programs, software, software applications or code) include machine instructions for a programmable processor, and may be implemented in a high-level procedural and/or object-oriented programming language, and/or in assembly/machine language. As used herein, the term “machine-readable medium” refers to any non-transitory computer program product, apparatus and/or device (e.g., magnetic discs, optical disks, memory, Programmable Logic Devices (PLDs)) used to provide machine instructions and/or data to a programmable processor, including a non-transitory machine-readable medium that receives machine instructions as a machine-readable signal. Non-transitory computer readable media can include media such as magnetic media (such as hard disks, floppy disks, etc.), optical media (such as compact discs, digital video discs, Blu-ray discs, etc.), semiconductor media (such as flash memory, electrically programmable read only memory (EPROM), electrically erasable programmable read only Memory (EEPROM), etc.), any suitable media that is not fleeting or not devoid of any semblance of permanence during transmission, and/or any suitable tangible media.

[0080] Some or all of the subject matter described herein may be implemented in a computing system that includes a back-end component (e.g., as a data server), or that includes a middleware component (e.g., an application server), or that includes a front-end component (e.g., a client computer having a graphical user interface or a Web browser through which a user may interact with an embodiment of the subject matter described herein), or any combination of such back-end, middleware, or front-end components. The components of the system may be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network (“LAN”), a wide area network (“WAN”), and the Internet.

[0081] The computing system may include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server generally arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0082] To evaluate the efficacy of the systems and procedures described herein to detect / identify the existence of different types of malicious processes, testing for different hardware device (e.g., different processor architecture) and for different malicious processes (malware, side channel attack processes, etc.) was conducted.

[0083] The evaluations and testing performed included testing to determine the efficacy of the systems and procedures described herein to detect Android malware. Examples of Android malware include malware to create advertisements and install unwanted links, cluttering up the user's device, etc. More advanced malware may take advantage of a phone's features, e.g., making phone calls or sending text messages to premium services, resulting in charges on the user's cell phone bill. Other type of Android malware may compromise a user's privacy in various ways, including accessing information like phone numbers, contact information, IMEI numbers and other sensitive data. Moreover, many Android-based mobile devices have GPS capability, and therefore malware may be capable of physically tracking victims.

[0084] The systems and procedures described herein were applied to Android malware obtained from various sources that catalog or study malwares. The malware data sets acquired were divided into families of variants. In families with only one variant, different execution cycles were used to acquire micro-architecture data for the malware specimen. For families with more than one variant, some of the variants were used for training purposes (e.g., generally about a 1/3 of the variants were used for training), while the remaining variants were used for testing (e.g., to determine if the system and procedures described herein would detect those variants). FIG. 5 is a table 500 of some of the Android malware families that were tested. Column 502 in the table 500, identified as APKs, indicates the number of variants that were available for the respective malware families.

[0085] To train the classifiers that were used to process micro-architectural data, micro-architectural performance data was collected on all malware samples. In the evaluations and testing performed, the collection infrastructure operated at the thread level. In addition to data on malware, data for 86 non-malware applications was also collected, resulting in data on 92,475 non-malware threads. These data sets were used for both training and testing, with generally a 1/3 of the data sets used for training, and the rest of the data sets used for testing. The performance of the various classifiers that may be used to process the micro-architectural data

can be adjusted using their respective configuration parameters. For instance, for a k -Nearest Neighbors (KNN) classifier, k is an adjustable configuration parameter. To identify an optimal set of parameters, the classifier (or, in some situations, several classifiers) that is chosen is the one that identifies the most malware correctly. However, the amount of malware identified varies with false positive rate. As a classifier is configured to make it more sensitive, more malware is identified, but non-malicious, legitimate processes are then also identified as malware. To determine which classifier (or classifiers) to use, in some embodiments, the classifier(s) that performs best (on the training data) for a given false positive percentage may be selected.

[0086] FIG. 6 contains a graph 600 showing the accuracy of various binary classifiers in detecting Android malware. As illustrated in the graph 600, as the false positives rate is increased, the classifiers find more malware (column 504 of table 500 in FIG. 5 provides the rate for correctly identifying the executed processes for the various malware families on which the classifiers, such the decision tree classifier, were applied for a false positive rate of 10% or better). The results obtained for the Android malware testing indicate that the classifiers tested work properly and that micro-architectural data (including micro-architectural performance counter data) can, with simple analysis, be used to detect Android malware with relatively good accuracy. For example, the “AnserverBot” malware had 187 known variants (which, as noted, were obtained from third parties that study and categorize malicious processes such as Android malware). Of those 187 known variants, 61 variants were used to train the classifiers of the systems and procedures described herein. After being trained with those 61 variants, the classifiers tested were able to identify 96.6% of the threads of the remaining 126 variants.

[0087] Evaluations and testing to determine the efficacy of the systems and procedures described herein was also performed on known Linux rootkits. Rootkits are malicious software that attackers install on compromised systems to evade detection and maximize their period of access on the systems. Once installed, rootkits hide their presence in the systems, typically by modifying portions of the operating systems to obscure specific processes, network ports, files, directories and session log-on traces. With their stealth capabilities, they can pose a significant threat to systems security due to the difficulty in detecting such infections.

[0088] In evaluating and testing the efficacy of the systems and procedures to detect Linux rootkit processes, two publicly available Linux rootkits were used that, once loaded, gave an attacker the ability to hide log-on session traces, network ports, processes, files and directories. The two rootkits used were:

1. Average Coder Rootkit – This rootkit works as a loadable kernel module that hide traces via hooking the kernel file system function calls. It is loaded into the kernel via the Linux command *insmod*. It allows the attacker to modify the system information to hide at runtime by writing via the echo command to a pre-defined file */proc/buddyinfo*.
2. Jynx2 Rootkit – This rootkit functions as a shared library and is installed by configuring the LDPRELOAD environment variable to reference this rootkit. When this is done, the rootkit is executed as a shared library whenever any program runs. The information it hides is pre-configured at compile-time and cannot be modified once it is loaded.

[0089] The Linux operating system has native utility programs that produce listings of the system current state (such as current process listing and network ports). To evade detection, the rootkits are designed to obscure portions of the output of these programs. Therefore, it is likely that micro-architecture performance counter data for these programs (produced on the processor device on which the rootkit processes are executing) will show some degree of deviation after a rootkit infection. To examine the presence of such deviation, collection of per-process performance counter data focused on the following processes

<u>Program</u>	<u>Relevant function</u>
ps	List active running processes
ls	List files and directories
who	List active log-on sessions
netstat	List active network connections

[0090] Micro-architectural performance counter data was collected for various arbitrarily selected event types (such as number of branch misprediction, number of data TLB misses, number of L1 instruction cache reads) for multiple execution runs of all the programs. Two sets of data were collected – one set was collected before the rootkits were installed (that collected set was referred to as the “clean set”), and the second set was collected after the system was infected with the rootkits (that set was referred to as the “dirty set”). To introduce variation to the

execution flows of the programs, each run of the programs was executed with a random combination of their associated parameters. To do this, a list of command-lines comprising the program names combined with a random set of valid parameters was generated. Each command-line was then randomly tagged as either clean or dirty to indicate the set of data it would be used in. An example subset list of the command-lines is provided below:

```
(clean) netstat -n
(clean) netstat -nt
(dirty) netstat -ntu
...
(dirty) ls -l /usr/include
(clean) ls -ld /home
(dirty) ls -lar /home/user
(clean) ls -lart ../
...
```

[0091] With the random list of command-lines generated, per-process per-run performance data was collected. Additionally, to reduce input bias and to make the collected data more realistic, the action of various users logging into the server and doing a series of tasks (like creating new files and running new processes) was simulated. Because both the rootkits that were used have different stealth capabilities and target the outputs of different programs, dirty data was collected separately for each rootkit. The collection of the data for each rootkit was performed with the following programs it was designed against:

<u>Program</u>	<u>Average Coder</u>	<u>Jynx2</u>
ps	✓	
ls		✓
who	✓	
netstat	✓	

[0092] While the dirty data was being collected, the information hidden by the rootkits was also varied. This included adding to, and removing from, the list of network ports, files, processes and log-on session logs that were hidden by the rootkits. As with the testing performed with respect to the Android malware, the micro-architectural data collected about the rootkits was divided into testing and training sets, with 1/3 of the data being used for training a large number of classifiers, and the remaining data used for testing the trained classifiers. The classifiers were trained to determine if the processes / programs for which micro-architectural data was being collected were running with or without rootkits (i.e., whether or not there was a rootkit contamination). FIG. 7 includes graphs 700 showing the accuracy (in terms of the number of correctly identified malicious threads as a function of false-positive rate) of the classifiers used as part of AV engine implemented herein. Although the accuracy achieved by the systems and procedures described herein for rootkit detection is generally lower than that achieved when the systems and procedures were applied to the Android malware, it is to be noted that because rootkits do not operate as separate programs, but rather are configured to dynamically intercept programs' normal control flow, the training data used is affected by the rootkits to a relatively small degree. As a result, identification of rootkits is generally more difficult than identification of other types of malicious processes.

[0093] Evaluation and testing of the systems and procedures described herein was also performed in relation to side-channel attacks. The term side-channel refers to unintended information leakage present in real implementations of systems. Because specific implementations cannot adhere to the idealized axioms of theoretical models, side-channels can be used to steal information from theoretically secure systems. For example, RSA cryptographic keys can be stolen by observing the performance of the branch predictor, or of the caches, for most existing implementations. Common side-channel mediums include acoustic or electrical signals, power draw, application-level timing channels, architectural or micro-architectural affects, or, in general, any shared resources. Although side-channel attacks are not generally considered malware, they render security of a hardware-based device vulnerable, and furthermore, have characteristic micro-architectural behavior that may be detected by the systems and procedures described herein.

[0094] A side-channel “attacker” process is a process that gets placed within the system in such a way that it shares a resource and uses that resource to learn information. Micro-architectural examples include sharing a network card, a core pipeline, memory bandwidth and caches. In embodiments involving side-channel attacks on a cache, shared on-chip caches can leak tremendous amounts of data that can be readily used to, for example, steal cryptographic keys and/or other types of private data. Intuitively, attacker programs that exploit micro-architectural side-channels should have clear signatures in terms of performance. For example, side-channel attack processes repeatedly thrash a particular shared resource so as to gauge all the activity of the victim process with respect to that shared resource. Micro-architectural events and performance counters are therefore likely to take on extreme values during such attacks, and thus indicate that the occurrence of attacker programs / processes (and possibly identify those attacker programs / processes).

[0095] To test the efficacy of the systems and procedures described herein to determine the occurrence of side-channel attacks (and/or identify the specific side-channel attack processes), an array of cache side-channel attacks was implemented. Variants of the standard prime-and-probe technique were implemented, in which an attacker program / process wrote to every line in the L1 data cache, and then scanned the cache repeatedly (using a pattern chosen at compile time) to read every line. Whenever a miss occurred, it meant there was a conflict miss caused by the victim process sharing the cache. The result data of a successful prime-and-probe attack includes data about the cache lines used by the victim process over time. The prime-and-probe variants were implemented and executed against an OpenSSL victim process. The cache side-channel attack processes were compared against a wide array of normal processes, which included programs of SPEC2006 int, SPEC2006 fp, PARSEC, web browsers, games, graphics editors and other common desktop applications, as well as generic system-level processes.

[0096] As with testing performed for the Android malware and Linux rootkits, 1/3 of the micro-architectural data collected was used to train the classifiers of the AV engine (namely, the KNN, DecisionTree, Tensor, RandomForest, and FANN classifiers). In this case, the training data included 3872 normal program threads and 12 attack threads. The trained classifiers were used to analyze the remaining two thirds of the collected data. The classifiers achieved perfect results when analyzing the 7744 normal threads and 24 attacks threads of this example testing, detecting

all 24 attack threads without producing any false positives. The results also indicated that in processing side-channel attack micro-architectural data, it did not matter which particular classifier was used.

[0097] Although particular embodiments have been disclosed herein in detail, this has been done by way of example for purposes of illustration only, and is not intended to be limiting with respect to the scope of the appended claims, which follow. Some other aspects, advantages, and modifications are considered to be within the scope of the claims provided below. The claims presented are representative of at least some of the embodiments and features disclosed herein. Other unclaimed embodiments and features are also contemplated.

WHAT IS CLAIMED IS:

1. A method comprising:
obtaining hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for a hardware device executing one or more processes; and
determining based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process.
2. The method of claim 1, wherein obtaining the hardware-based micro-architectural data comprises:
obtaining the hardware-based micro-architectural data at various time instances.
3. The method of claim 2, wherein obtaining the hardware-based micro-architectural data at the various time instances comprises:
performing one or more of a data push operation initiated by the hardware device to send the micro-architectural data, or a data pull operation, initiated by an antivirus engine, to send the micro-architectural data.
4. The method of claim 1, wherein obtaining the hardware-based micro-architectural data comprises:
obtaining multi-core hardware-based micro-architectural data resulting from execution of the one or more processes on a processor device with multiple processor cores; and
correlating the respective hardware-based micro-architectural data obtained from each of the multiple processor cores to the one or more processes.
5. The method of claim 1, wherein determining based on the hardware-based micro-architectural data whether the at least one of the one or more processes corresponds to a malicious process comprises:
applying one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process.

6. The method of claim 5, wherein applying the one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process comprises:

matching the obtained hardware-based micro-architectural data to previously identified patterns of hardware-based micro-architectural data associated with one or more malicious processes.

7. The method of claim 6, further comprising:

obtaining updates for the previously identified patterns of hardware-based micro-architectural data associated with the one or more malicious processes.

8. The method of claim 7, wherein obtaining the updates comprises:

downloading encrypted data for the previously identified patterns of hardware-based micro-architectural data associated with the one or more malicious processes to an antivirus engine in communication with the hardware device providing the hardware-based micro-architectural data;

decrypting at the antivirus engine the downloaded encrypted data for the previously identified patterns of hardware-based micro-architectural data associated with the one or more malicious processes; and

updating a revision counter maintained by the antivirus engine indicating a revision number of a most recent update of the previously identified patterns of hardware-based micro-architectural data.

9. The method of claim 5, wherein the one or more machine learning procedures comprise one or more of: a k-nearest neighbor procedure, a decision tree procedure, a random forest procedure, an artificial neural network procedure, a tensor density procedure, or a hidden Markov model procedure.

10. The method of claim 1, wherein the malicious process comprises one or more of: a malware process, or a side-channel attack process.

11. The method of claim 1, wherein the hardware-based micro-architectural data comprise one or more of: processor load density data, branch prediction performance data, or data regarding instruction cache misses.

12. A system comprising:
a hardware device executing one or more processes; and
an antivirus engine in communication with the hardware device, the antivirus engine configured to:
obtain hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for the hardware device executing the one or more processes; and
determine based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process.

13. The system of claim 12, wherein the antivirus engine configured to obtain the hardware-based micro-architectural data is configured to:
obtain the hardware-based micro-architectural data at various time instances.

14. The system of claim 13, wherein the antivirus engine configured to obtain the hardware-based micro-architectural data at the various time instances is configured to:
receive the micro-architectural data in response to one or more of: a data push operation initiated by the hardware device, or a data pull operation initiated by the antivirus engine.

15. The system of claim 12, wherein the antivirus engine configured to determine based on the hardware-based micro-architectural data whether the at least one of the one or more processes corresponds to a malicious process is configured to:
apply one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process.

16. The system of claim 15, wherein the antivirus engine configured to apply the one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process is configured to:

match the obtained hardware-based micro-architectural data to previously identified patterns of hardware-based micro-architectural data associated with one or more malicious processes.

17. The system of claim 16, wherein the antivirus engine is further configured to:
obtain updates for the previously identified patterns of hardware-based micro-architectural data associated with the one or more malicious processes.

18. A computer readable media storing a set of instructions executable on at least one programmable device that, when executed, causes operations comprising:

obtaining hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for a hardware device executing one or more processes; and

determining based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process.

19. The computer readable media of claim 18, wherein determining based on the hardware-based micro-architectural data whether the at least one of the one or more processes corresponds to a malicious process comprises:

applying one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process.

20. The computer readable media of claim 18, wherein applying the one or more machine-learning procedures to the hardware-based micro-architectural data to determine whether the at least one of the one or more processes corresponds to the malicious process comprises:

matching the obtained hardware-based micro-architectural data to previously identified patterns of hardware-based micro-architectural data associated with one or more malicious processes.

21. An apparatus comprising:

means for obtaining hardware-based micro-architectural data, including hardware-based micro-architectural counter data, for a hardware device executing one or more processes; and

means for determining based, at least in part, on the hardware-based micro-architectural data whether at least one of the one or more processes executing on the hardware device corresponds to a malicious process.




FIG. 1




FIG. 2

3 / 9




FIG. 3




FIG. 4

5 / 9

500

502

Malware Family	APKs	Training Threads	Testing Threads	Correctly Identified	Rate
Tapsnake	1	31	3	3	100%
Zitmo	1	5	1	1	100%
Loozfon-android	1	25	7	7	100%
Android.Steek	3	9	9	9	100%
Android.Trojan.Qicsomos	1	12	12	12	100%
CruseWin	1	2	4	4	100%
Jifake	1	7	5	5	100%
AnserverBot	187	9716	11904	11505	96.6%
Gone60	9	33	67	59	88.1%
YZHC	1	9	8	7	87.5%
FakePlayer	6	7	15	13	86.7%
LoveTrap	1	5	7	6	85.7%
Bgserv	9	119	177	151	85.3%
KMIN	40	43	30	25	83.3%
DroidDreamLight	46	181	101	83	82.2%
HippoSMS	4	127	28	23	82.1%
Dropdialerab	1	18*	16*	13	81.3%
Zsone	12	44	78	63	80.8%
Endofday	1	11	10	8	80.0%
AngryBirds-LeNa.C	1	40*	24*	19	79.2%
jSMShider	16	101	89	70	78.7%
Plankton	25	231	551	432	78.4%
PJAPPS	16	124	174	136	78.2%
Android.Sumzand	1	8	9	7	77.8%
RogueSPPush	9	236	237	184	77.6%
FakeNetflix	1	27	8	6	75.0%
GEINIMI	28	189	203	154	75.9%
SndApps	10	110	77	56	72.7%
GoldDream	47	1160	237	169	71.3%
CoinPirate	1	8	10	7	70.0%
BASEBRIDGE	1	14*	72	46	63.8%
DougaLeaker.A	6	12*	35*	22	62.9%
NewZitmo	1	5	8	5	62.5%
BeanBot	8	122	93	56	60.2%
GGTracker	1	16	15	9	60.0%
FakeAngry	1	7	10	5	50.0%
DogWars	1	14	8	2	25.0%

FIG. 5

6 / 9




FIG. 6




FIG. 7




FIG. 8A

9 / 9




FIG. 8B

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 13/68451

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 11/00 (2014.01)

USPC - 726/23

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 (2014.01)

USPC - 726/23

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

USPC - 713/320, 713/300, 713/502, 726/34, 380/203, 713/500, 726/23, 713/601, 713/400, 713/501, 726/25, 713/600, 726/22 (See keywords Below)

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

Thomsoninnovation.com; Patbase; Google Scholar; Google Patents; Gogole.com; Freepatentsonline; ProQuest Dialog

Search Terms: hardware based, micro architectural, counter, performance counter, malware, side channel, virus, botnet, DoS, learn, train, pattern, behavior, time, sampling, frequency, processor, core, process, thread, load, performance, c

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2008/0184371 A1 (MOSKOVITCH et al.), 31 July 2008 (31.07.2008), entire document, especially Abstract; para [0016], [0018], [0021], [0034], [0198]	1-21
Y	US 2012/0254641 A1 (ROTEM et al.), 04 October 2012 (04.10.2012), entire document, especially Abstract; para [0003], [0025], [0029], [0050]	1-21
Y	US 2003/0070087 A1 (GRYAZNOV et al.), 10 April 2003 (10.04.2003), entire document, especially Abstract; para [0009], [0010], [0029]	8
A	US 2011/0228926 A1 (SHUMOW et al.), 22 September 2011 (22.09.2011), entire document	1-21
A	US 2008/0189530 A1 (MCINTOSH et al.), 07 August 2008 (07.08.2008), entire document	1-21
A	US 7,590,880 B1 (HERSHMAN), 15 September 2009 (15.09.2009), entire document	1-21



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

13 April 2014 (13.04.2014)

Date of mailing of the international search report

02 MAY 2014

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