



(51) International Patent Classification:
G06F 21/56 (2013.01)

(21) International Application Number:
PCT/US2016/053889

(22) International Filing Date:
27 September 2016 (27.09.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
14/923,184 26 October 2015 (26.10.2015) US

(71) Applicant: **MCCAFEE, INC.** [US/US]; 2821 Mission College Blvd., Santa Clara, CA 95054 (US).

(72) Inventors: **SHANMUGAVELAYUTHAM, Palaniverran, R.**; 1615 Clarkspur Lane, San Jose, CA 95129 (US). **YAMADA, Koichi**; 982 Cherrystone Drive, Los Gatos, CA 95032 (US). **SUKHOMLINOV, Vadim**; 2870 Kaiser Dr., Apt. 444, Santa Clara, CA 95051 (US). **MUTTIK, Igor**; Target House, Gatehouse Road, Aylesbury, bkm GB HP19 8ED (GB). **BAZHANIUK, Oleksandr**; 1470 NE Alex Way, #359, Hillsboro, OR 97124 (US). **BULYGIN, Yuriy**; 16475 SW Black Bird Drive, Beaverton, OR 97007 (US). **RUBAKHA, Dmitri, D.**; 2821 Mission College

Blvd, Santa Clara, CA 95054 (US). **MANKIN, Jennifer**; 1318 The Alameda, Apt. 337, San Jose, CA 95126 (US). **WOODWARD, Carl**; 2821 Mission College Blvd, Santa Clara, CA 95054 (US). **VAROGLU, Sevin**; 3600 Juliette Ln, Ms: Sc12-301, Santa Clara, CA 95054 (US). **MIRKIN, Dima**; Opal 19, Mevaseret Tzion, JM, IL, 9078619 (IL). **NAYSHTUT, Alex**; Hadagan 10/3, Gan Yavne, D, IL 80700 (IL).

(74) Agent: **SCHAFER, Richard, A.**; Blank Rome, LLP, 717 Texas Avenue, Suite 1400, Houston, TX 77002 (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, DJ, 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, KW, KZ, LA, LC, LK, LR, LS, 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,

[Continued on next page]

(54) Title: **HARDWARE HEURISTIC-DRIVEN BINARY TRANSLATION-BASED EXECUTION ANALYSIS FOR RETURN-ORIENTED PROGRAMMING MALWARE DETECTION**

(57) Abstract: A combination of hardware monitoring and binary translation software allow detection of return-oriented programming (ROP) exploits with low overhead and low false positive rates. Embodiments may use various forms of hardware to detect ROP exploits and indicate the presence of an anomaly to a device driver, which may collect data and pass the indication of the anomaly to the binary translation software to instrument the application code and determine whether an ROP exploit has been detected. Upon detection of the ROP exploit, the binary translation software may indicate the ROP exploit to an anti-malware software, which may take further remedial action as desired.




FIG. 1



GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, 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).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

HARDWARE HEURISTIC-DRIVEN BINARY TRANSLATION-BASED EXECUTION ANALYSIS FOR RETURN-ORIENTED PROGRAMMING MALWARE DETECTION

TECHNICAL FIELD

[0001] Embodiments described herein generally relate to techniques for detecting malware, and more particularly to techniques for detecting return oriented programming exploits.

BACKGROUND ART

[0002] Return and jump oriented programming (ROP/JOP) exploits are a growing threat for software applications. These techniques allow an attacker to execute code even if security measures such as non-executable memory and code signing are used. In ROP, an attacker gains control of the call stack and then executes carefully chosen machine instruction sequences, called “gadgets.” Each gadget typically ends in a return instruction and is code within an existing program (or library). Chained together via a sequence of carefully crafted return addresses, these gadgets allow an attacker to perform arbitrary operations. JOP attacks do not depend upon the stack for control flow, but use a dispatcher gadget to take the role of executing functional gadgets that perform primitive operations.

[0003] Detection of ROP exploits is complicated due to the nature of the attack. A number of techniques have been proposed to subvert attacks based on return-oriented programming.

[0004] The first approach is randomizing the location of program and library code, so that an attacker cannot accurately predict the location of usable gadgets. Address space layout randomization (ASLR) is an example of this approach. Unfortunately, ASLR is vulnerable to information leakage attacks and once the code location is inferred, a return-oriented programming attack can still be constructed. Randomization approach can be taken further by employing relocation at runtime. This complicates the process of finding gadgets but incurs significant overhead.

[0005] A second approach modifies the operating system with compiler support to track that return instructions actually divert control flow back to a location immediately following a call instruction. This prevents gadget chaining, but carries a heavy performance penalty. In addition, it is possible to mount JOP attacks without using return instructions at all, by using JMP instructions.

[0006] In a third approach, some Intrusion Protection Systems (IPSeS) invalidate memory pages of a process except one currently executed page. Most regular jumps land within the same page. Passing control flow to a different page causes an exception that allows the IPS to check the control flow. This technique may also introduce a noticeable overhead.

[0007] Better approaches to both ROP and JOP attacks that do not incur large performance penalties would be desirable.

BRIEF DESCRIPTION OF DRAWINGS

[0008] Figure 1 is a block diagram illustrating an example of existing performance monitor hardware that may be used for the exploit detection techniques described below.

[0009] Figure 2 is a block diagram illustrating an example of existing last branch recording hardware that may be used for the exploit detection techniques described below.

[0010] Figure 3 is a block diagram illustrating a system for detecting return-oriented and jump oriented programming exploits according to one embodiment.

[0011] Figure 4 is a flowchart illustrating a hardware-assisted technique for detecting return-oriented and jump-oriented programming exploits according to one embodiment.

[0012] Figures 5–6 are block diagrams illustrating two embodiments of programmable devices in which the techniques described herein may be implemented.

DESCRIPTION OF EMBODIMENTS

[0013] In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the invention. It will be apparent, however, to one skilled in the art that the invention may be practiced without these specific details. In other instances, structure and devices are shown in block diagram form in order to avoid obscuring the invention. References to numbers without subscripts or suffixes are understood to reference all instance of subscripts and suffixes corresponding to the referenced number. Moreover, the language used in this disclosure has been principally selected for readability and instructional purposes, and may not have been selected to delineate or circumscribe the inventive subject matter, resort to the claims being necessary to determine such inventive subject matter. Reference in the specification to “one embodiment” or to “an embodiment” means that a particular feature, structure, or characteristic described in

connection with the embodiments is included in at least one embodiment of the invention, and multiple references to “one embodiment” or “an embodiment” should not be understood as necessarily all referring to the same embodiment.

[0014] As used herein, the term “a computer system” can refer to a single computer or a plurality of computers working together to perform the function described as being performed on or by a computer system.

[0015] As used herein, the terms “anti-virus software” and “anti-malware software” mean computer software used to prevent, detect, and remove malicious software (malware). Although not all malware is a virus, anti-virus software has evolved to provide protection from both viruses and other types of malware, and the terms have become synonymous and may be used interchangeably herein.

[0016] As used herein, the terms “processor” mean any type of computational element, such as but not limited to, a microprocessor, a processor, central processing unit (CPU), a digital signal processing unit, a multiple core processor, a mobile device processor, a desktop processor, a single core processor, a system-on-chip (SoC) device, a complex instruction set computing (CISC) microprocessor, a reduced instruction set (RISC) microprocessor, a very long instruction word (VLIW) microprocessor, or any other type of processor or processing circuit on a single chip or integrated circuit.

[0017] Over the past few years, code reuse exploit attacks such as ROP have become popular as ROP allows bypassing mitigation techniques such as Data Execution Prevention (DEP). ROP and its variants JOP and Call oriented programming (COP) are code-reuse attacks, where existing code snippets called gadgets (ending with a RET/JUMP/CALL instruction) are chained together to construct malicious computations. In the discussion below, ROP is used as representative of ROP, JOP, COP, and other variants of ROP.

[0018] To prevent the attackers from knowing the addresses of gadgets needed for a ROP attack, modern software may employ an ASLR (Address Space Layout Randomization) mitigation technique. Unfortunately, ASLR can often be bypassed via various address discovery techniques. ROP exploits also have demonstrated various successful bypass techniques against API hooking-based caller validation, which are currently employed with Microsoft’s Enhanced Mitigation Experience Toolkit (EMET) and McAfee’s Host Intrusion Prevention System (HIPS). As of today, no reliable defense exists for ROP malware attacks.

[0019] Tools such as Binary Translation/Instrumentation (BT) provide fine grain instruction level behavioral monitoring and have proven to be a highly accurate method for detecting ROP malware. However, despite high detection efficacy with minimum false positives for ROP exploits, BT methods are prone to creating significant runtime overheads (10–15% performance impact) during monitoring.

[0020] Other anti-ROP methods may make use of one or more processor hardware features, such as Last Branch Record (LBR), Performance Monitoring Unit (PMU), processor trace (PT), and dedicated ROP heuristic detection hardware, such as hardware that attempts to detect ROP malware through detection of an RET–CALL imbalance. These approaches are hardware based and have low overhead. They may signal to anti-malware software about situations that look like ROP attacks: suspicious branching, RET–CALL biases, rate of branch or return mispredicts, etc. Their challenges lie in 1) the need for additional software analysis to make a final accurate decision (to reduce false positive rate), and 2) a possibility of malware escape and evasions when the threshold of hardware events is set too high.

[0021] A combined solution described herein utilizes events from the processor that trigger a binary translation (BT) based method. Binary Translation will provide forward execution analysis to eliminate false positives and detect an ROP attack with high accuracy. Techniques described below provide a low overhead and high accuracy detection/prevention solution against code reuse attacks involving a chain of ROP gadgets by using heuristic hardware-based events to trigger BT code analysis for forward execution to make a final exploit detection decision.

[0022] ROP detection hardware counters are used to detect branch discrepancies such as RET/CALL bias, return mispredicts, indirect branch mispredicts, etc. to trigger the binary translation based execution-monitoring method. BT provides opcode-level execution runtime analysis after the heuristic event trigger to determine the validity of a trigger to eliminate false positive detections. These may be hardware counters that are specific to ROP detection or hardware counters that may be configured for various purposes, including ROP detection, or both. Different models of processors may deploy different types of hardware counters.

[0023] No known good solutions exist today for providing complete coverage against sophisticated code-reuse attacks involving ROP gadgets. Existing solutions like Microsoft's EMET have fundamental deficiencies and have been proven to be bypassed many times. The

approach described herein is to use low-overhead hardware based heuristic events to detect execution anomalies and trigger highly effective BT analysis. This can provide high detection efficacy against ROP malware with a zero false positive rate and with low overhead. Heuristic hardware triggers allow system-wide monitoring and can detect anomalies on any software footprint, not only for single processes. Binary translation also provides low overhead false positive filtering on events, as false positive checks can be in-lined with a translated version of original code and eliminated for validated paths. In one implementation, the techniques described herein generated less than 1% application overhead with ~200 events per second to trigger BT based instruction level execution analysis. The BT may expose standardized application programming interfaces (APIs) from BT software for AV software integration, abstracting hardware differences between the hardware implementations (e.g., PMU based solutions, ROP heuristic hardware, etc.) and allowing seamless transition to new hardware enhancements.

[0024] The techniques described below work on the legacy binaries. No independent software vendor (ISV) or operating system vendor (OSV) enabling is needed. For example, some software-only solutions such as Microsoft Control Flow Guard technology require application recompiling to get a partial solution to control flow integrity checks, as well as support in both the compilers and OS.

[0025] The techniques described herein are applicable to any operating system and compiler. As the solution does not depend upon hooking and binary patching, it provides a unique and harder-to-bypass approach to protect against complex exploits.

[0026] A combination of a low-level device driver and a user-level library component allow easy security/AV software integration. The hardware implementation is abstracted, enabling solutions across different hardware implementations and allowing seamless transition to new hardware enhancements without impacting the AV software.

[0027] This invention provides a low overhead detection and prevention solution against code reuse attacks involving a chain of ROP gadgets by using heuristic hardware based events to trigger the binary translation code analysis, which provides novel instruction level forward execution checking to make a final exploit detection decision.

[0028] The following components are described below:

[0029] 1) ROP Heuristic Detection Hardware Counters—constant running hardware counters to detect anomalies of branch executions such as RET/CALL bias, return mispredicts, indirect branch mispredicts, far branches, etc. When dedicated ROP heuristic hardware does not exist on a processor, a set of model specific performance monitoring (Perfmon) counters may be used.

[0030] 2) ROP Heuristic (ROPh) Device Driver—configures the ROP heuristic detection hardware counters and interfaces with BT software to signal a branch anomaly event or any other kind of code flow anomaly (for example, a stack pointer modification often called stack pivoting). The discussion below is written in terms of branch anomalies as representative code flow anomalies, but the techniques described may be used for all types of code flow anomalies.

[0031] 3) Binary Translation (BT) software—provides instruction level execution runtime analysis after the heuristic event and applies a variety of instruction level execution analysis methods to make a final deposition (ROP malware versus false positive).

[0032] 4) AV (Anti-Virus) software—configures BT software through software interfaces (APIs). AV software receives a notification callback from BT software upon detection of malware activities and takes a security measure against the attack (ex. terminate the app).

[0033] FIG. 1 is a block diagram illustrating a system implementing a combined hardware-software ROP detection technique according to one embodiment. As illustrated, a workstation or client computer 100 incorporates a processor 110 that provides hardware counters that may be useful for ROP detection by the hardware. Although system 100 is illustrated as a desktop type computer, the hardware and software techniques described herein may be implemented in any type of programmable device, including mobile devices such as a smart phone. Although described herein generally as counters, the hardware features may be implemented in hardware in any desired way, including registers, special memory locations, etc.

[0034] A memory 105 coupled to the processor 110 may be used for storage of information related to the detection and analysis techniques described herein. The memory 105 may be connected to the processor in any desired way, including busses, point-to-point interconnects, etc. The memory may be also be used for storing instructions that when executed cause the computer 100 to execute the ROPh device driver 120, the BT software 130, and the anti-malware software 140.

[0035] One skilled in the art will recognize that other conventional elements of a computer system or other programmable device may be included in the system 100, such as a keyboard, pointing device, displays, etc.

[0036] Processor **110** may comprise, for example, a microprocessor, microcontroller, digital signal processor (DSP), application specific integrated circuit (ASIC), or any other digital or analog circuitry configured to interpret and/or execute program instructions and/or process data. In some embodiments, processor **110** may interpret and/or execute program instructions and/or process data stored in memory **105**. Memory **105** may be configured in part or whole as application memory, system memory, or both. Memory **105** may include any system, device, or apparatus configured to hold and/or house one or more memory modules. Each memory module may include any system, device, or apparatus configured to retain program instructions and/or data for a period of time (e.g., computer-readable storage media). Instructions, logic, or data for configuring the operation of system **100**, such as configurations of components such as the performance monitoring hardware 115, the ROPh device driver 120, the BT software 130, or anti-malware software 140 may reside in memory **105** for execution by processor **110**.

[0037] While a single processor **110** is illustrated in FIG. 1, the system 100 may include multiple processors. Furthermore, processor **110** may include multiple cores or central processing units.

[0038] Memory **105** may include one or more memory modules and comprise random access memory (RAM), read only memory (ROM), programmable read only memory (PROM), programmable read-write memory, and solid-state memory. Memory 105 may also include a 2of optical and magnetic, including solid-state storage elements, including removable media. The storage device may be a program storage device used for storage of software to control computer 100, data for use by the computer 100 (including performance monitoring configuration data), or both. The instructions for configuring the performance monitoring hardware as well as for processing PMU interrupts (PMIs) and analyzing the collected data may be provided on one or more machine readable media, used either as part of the memory 105 or for loading the instructions from the media into the memory 105. Although only a single memory 105 is illustrated in FIG. 1 for clarity, any number of memory devices, including any number of storage devices, may be provided as desired as part of the memory 105.

[0039] The computer system 100 may be any type of computing device, such as, for example, a smart phone, smart tablet, personal digital assistant (PDA), mobile Internet device, convertible tablet, notebook computer, desktop computer, server, or smart television.

[0040] For example, modern computer processors have a Performance Monitoring Unit (PMU) for monitoring selected events. The diagram in Figure 2 illustrates the core PMU and related registers 200 on Intel x86 processors. Processors from different manufacturers may have similar PMUs, although architectural details may differ. The PMU 210 has a plurality of fixed purpose counters 220. Each fixed purpose counter 220 can count only one architectural performance event, thus simplifying the configuration part. In addition to the fixed purpose counters 220, the core PMU also supports a plurality of general purpose counters 230 that are capable of counting any activity occurring in the core. Each core PMU 210 also has a set of control registers 240, 260, to assist with programming the fixed purpose counters 220 and general purpose counters 230. The PMU 210 also has Event Select registers 250 that correspond to each fixed purpose counter 220 and general purpose counter 230, which allows for specification of the exact event that should be counted. A global control register 260 allows enabling or disabling the counters 220, 230. A global status register 270 allows software to query counter overflow conditions on combinations of fixed purpose counters 220 and general purpose counters 230. A global overflow control register 280 allows software to clear counter overflow conditions on any combination of fixed-purpose counters 220 and general purpose counters 230. The elements illustrated in FIG. 1 are illustrative and by way of example only, and other elements and arrangements of elements may be provided as desired. In a multi-core processor, each core may have its own core PMU 200; alternately, a single PMU 200 may be shared among the multiple cores.

[0041] Modern processor architectures also provide a branch recording mechanism. Typically, the last branch recording mechanism tracks not only branch instructions (like JMP, Jcc, LOOP, and CALL instructions), but also other operations that cause a change in the instruction pointer, like external interrupts, traps, and faults. The branch recording mechanisms generally employ a set of processor model specific registers, referred to as a last branch record (LBR) stack, each entry of which stores a source address and a destination address of the last branch, thus the LBR stack provides a record of recent branches. Some embodiments of an LBR stack may also record an indication of whether the branch was mispredicted, i.e., one or more of the target of the branch and the direction (taken, not taken) was mispredicted. In

addition, control registers may allow the processor to filter which kinds of branches are to be captured in the LBR stack. FIG. 3 is a block diagram illustrating an LBR stack 300 with two sets of registers 310A and 310B in a conventional processor. Each LBR stack entry 310 includes one register with a from address field 320 and a mispredicted indicator 330, and another register with a to address field 340. Although only 2 LBR stack entries 310 are illustrated in the LBR stack 300 of FIG. 3 for clarity, implementations typically have more LBR stack entries 310. Although illustrated with the mispredict indicator as part of the register containing the from address 320, embodiments may place the mispredict indicator as part of the register containing the to address 340, or may place the mispredict indicator in a third register (not shown in FIG. 3). Other fields may be included in the LBR stack 300 as desired.

[0042] One of the ways the Event Select registers 250 may be configured is to cause the PMU 210 to count branch anomalies, such as branch mispredict events. These events may be caused by ROP exploits, as well as for other reasons. Where branch capture filtering is available, the filter may be employed to limit the captured branches to those of interest in ROP exploits. For JOP exploits, the branches of interest are typically near indirect jumps. For ROP exploits, the branches of interest are typically CALLs or RETs. However, embodiments may filter other types of branches or do no branch capture filtering, if desired. For example, COP exploits may use gadgets that end with indirect CALL instructions. In COP exploits, gadgets are chained together by pointing the memory-indirect locations to the next gadget in sequence. COP exploits may be detected using a similar approach to that used for detecting ROP exploits, with the branches of interest being CALLs.

[0043] When a mispredict event occurs (or, preferably, when a mispredict count exceeds a predetermined threshold) the reason for the misprediction may be analyzed by matching the expected program code flow with the real one extracted from the LBR stack 300. The analysis is fairly simple because the from and to addresses 320, 340 are readily available from the LBR stack 300 and the from and to addresses directly point to the code in question, allowing separating valid reasons (say, indirect CALL or deep recursion) from exploit behavior (by employing, for example, static code flow analysis of the program).

[0044] By using these facilities, embodiments disclosed herein can detect ROP exploits without significant processor overhead. Systems that include ROP-specific hardware features may similarly deploy those features for detecting ROP exploits.

[0045] Using the hardware PMU and related registers 200 and the LBR stack 300 to collect mispredicted branch data for analysis introduces the following advantages:

[0046] 1. Low overhead compared to all existing methods (all events are gathered by CPU via PMU 210 and LBR stack 300).

[0047] 2. Ease of analysis: LBR event data points exactly to the suspected code.

[0048] 3. High ROP detection rate with an ability to fine-tune the sensitivity and minimize the false positive rate.

[0049] 4. Generic to majority of processor platforms: most recent processor platforms already have all the hardware needed for implementing this invention.

[0050] 5. Operating system (OS) agnostic: events collection is fully hardware-based, with no OS interaction or enablement needed.

[0051] 6. Resilience to OS, Hypervisor, Basic Input/Output System (BIOS), and Unified Extensible Firmware Interface (UEFI) malware: even in the presence of an OS or firmware-based malware, events will be reliably collected and securely delivered to the monitoring agent.

[0052] 7. PMU logic allows counting mispredicted RET instructions and enabling a PMI once the counter reach a predetermined threshold. This provides additional hardware-supported sensitivity control to maximize the true positive rates (fine-tuning will allow catching the smallest observed ROP shell code sequences). Not every mispredicted branch indicates an exploit. In one embodiment, the threshold value may be empirically determined, based on analysis of detected ROP exploits. In some embodiments, the threshold value may be configured based on a policy.

[0053] Similar advantages may be achieved by using ROP-specific hardware in addition to or instead of general purpose PMUs 210, in which heuristic techniques may be used to flag a potential ROP event.

[0054] In one embodiment, PMU and related registers 200 are enhanced by additional hardware elements that keep a running sum of the number of return (RET) related instructions while subtracting out the number of call (CALL) related instructions dynamically during execution. The counting is computed over some sufficiently large window for the purpose of filtering out any short term excursions in this tabulation due to legitimate nested programming behavior. A threshold may be applied to detect if this accumulated sum exceeds a designated

value. In some embodiments, the designated value and the instruction window size may be controlled from by the device driver 120 under the direction of the BT software 130 or the anti-malware software 140. The controls allow for distinguishing legitimate code from abnormal behavior as possibly used by an attacker with ROP style code.

[0055] For ordinary non-ROP programs, over the long term the number of RETs generally approximately equals the number of CALLs, although the bias may be non-zero. In the case of a ROP-style program, detection would occur because the number of RETs over some window will greatly exceed the number of CALLs. When this threshold is exceeded, then the processor hardware can signal an exception that can be fielded by the device driver 120 and provided to the BT software 130 or anti-malware software 140.

[0056] Over the long term there can be a natural or legitimate imbalance between the number of RETs and number of CALLs. To address that aspect, the tabulated sum accumulated over multiple instruction windows may be biased toward zero by down-scaling the final accumulation value with each partial sum update. In one embodiment, this down-scaling can be achieved by a simple shifting of the accumulator of RETs-CALLs to the right by 1 or more bits every 'N' block of instructions. This has the effect of achieving visibility over a larger scope of execution while self-calibrating to the currently executing program context.

[0057] In addition, in one embodiment, the LBR logging capability described above may be enhanced with a freeze control. This freeze control could be activated when the ROP detect threshold mechanism is tripped, providing a backtrace of execution of where the program has been. The LBR mechanism could also be configured to filter or capture only call and return instructions. In this way the BT software 130 or AV software 140 can have a more precise picture of the last 16 (or whatever the LBR record depth support is) related CALL/RET transfers, which may more specifically identify the ROP related content. Embodiments may be configurable to enable the freeze control feature via the device driver 120.

[0058] A goal of this technology to work in a robust manner independent against future potential counter-measures that an attacker may attempt if aware of a RET instruction-based detection mechanism. Any ROP style attack must fundamentally be able to execute a program chain sequence as strung together from a series of address pointers held on the program stack.

[0059] One additional solution that may be monitored by processor hardware in one embodiment is to tabulate the number of stack related instructions and register-based control

transfers for a given instruction count window, looking for cases where the two are used at an unusually high rate. For example, instead of incorporating a RET instruction, an attacker may try to exploit the equivalent code fragment: “POP reg; JMP reg” instead. This would be more difficult to utilize given the lower probability of finding this 4 byte opcode sequence versus a 1 byte RET opcode at a useful location, but not impossible.

[0060] In one embodiment, the processor hardware may apply a counting heuristic of counting “JMP reg” instances when it sees a recent higher rate of stack push or pop operations. This implies keeping track of multiple processor instruction events and applying a similar signaling to the BT software 130 or anti-malware software 140 when such control transfers that may have been associated with a variable derived from the stack. To help reduce the rate of false positives, knobs may be provided to BT software 130 or anti-malware software 140 to control the instruction count window and threshold values for both stack related operations and register based control transfers. In some embodiments adjust the values may be adjusted dynamically based on code context and a degree of trust associated with the currently running code.

[0061] As an ROP style program executes, address pointers ultimately must be popped off the stack as the program progresses, causing the stack pointer value to increase over time. This can be more generically detected by the processor keeping track of changes in the stack pointer with respect to the execution of any control transfer instruction that makes use of an indirect register value (that could have come from the stack such as “JMP reg”) or references the stack directly (i.e. RET). By tabulating for each such control transfer if the stack pointer has increased, decreased, and remain unchanged, one can detect rough equivalence. For example a “RET” or equivalent “POP reg; JMP reg” may be detected roughly as a control transfer with an increasing stack pointer by at least the program counter width from the previous control transfer. Likewise a “CALL” or equivalent “PUSH reg; JMP reg” may be detected roughly as a control transfer with a decreasing stack pointer by at least the program counter width. Over a sufficiently long window, ROP-style code will see many more control transfers in which the stack point increases, while legitimate code may have other manipulations on the stack due to other variable usages, but ultimately stack space usage is bounded (assuming no memory leaks). As a result, more noise may exist for distinguishing legitimate code from ROP-style code, but through the use of hardware control knobs previously described for the window size, biasing, and threshold, embodiments may detect this usage and flag them as a branch anomaly.

[0062] One variant ROP style technique is to take advantage of stack pivot sequence code that can be naturally found in code as part of management of the stack in a separate frame for the target routine. Such stack pivot sequence code can naturally end with a restore of the stack pointer (ESP) from some other saved value prior to execution of the “RET” that retrieves the pushed instruction pointer from the original stack. Such code may end with a sequence like the following:

[0063] mov esp, eax

[0064] ret

[0065] An attacker could take advantage of this to create a gadget where they have created a ROP pointer sequence in a separate data region from the original program stack. This may be useful if the existing stack is limited in space due to some size or boundary constraints. The attack would require assembling other code via a ROP gadget to initialize register ‘eax’ in the example above to the attacker controller data region.

[0066] To address use of such a stack pivot sequence in one embodiment, the ROP heuristic detection hardware 115 may look for such pairing of instructions (instruction 1: stack pointer is the destination operand, instruction 2: return), and count these with additional weighting (e.g., increment the biased accumulator by 2 instead of 1) versus an ordinary non-pivoting return, on the basis that a higher frequency of stack pivoting operations is more unusual in ordinary code versus ROP-style attack code. A separate configuration control may also be provided to control the weighting for such a count.

[0067] Returning to FIG. 1, the workstation or client 100 employs hardware features 115 such as a PMU and counters 200 or ROP heuristic hardware for collecting data in hardware for use in detecting ROP exploits. The hardware features 115 may generate an interrupt, such as a performance management interrupt (PMI) to signal an ROP heuristic detection driver 120, typically implemented as part of the OS controlling the workstation 100. Any other technique for signaling a potential ROP exploit may be used, including non-interrupting signaling. The ROP heuristic driver 120 may also be used for configuring the hardware features 115. In one embodiment, the ROP heuristic detection driver 120 provides information extracted from the hardware features 115 to a BT software 130, which can perform execution profiling. If the BT software 130 determines that an ROP exploit has been detected, it may inform the AV software

140, which acts as a security agent, so that the AV software 140 may take whatever action is desired in response.

[0068] The BT software 130 and AV software 140 are illustrated in FIG. 1 as installed and executing on the workstation 100. Embodiments may allow one or more of the BT software 130 or AV software 140 to execute at least in part on other programmable devices. In one embodiment, the BT software 130 may send information about the possible exploit to a Security Information and Event Management (SIEM) system 150 that may be deployed in a remote server, for example.

[0069] In the discussion below, AV software 140, for example the McAfee Host Intrusion Prevention System, is assumed present on the target system 100 to manage the overall security policies; however, implementations may provide the AV software 140 on a separate device. The AV software 140's job is to attach the BT software 130 to the target application and configure and activate the security policy with APIs exposed by BT software 130. An appropriate level of security policy can be determined per application based upon the malware risk of the application.

[0070] The BT software 130 provides instruction level execution runtime analysis after the heuristic event and may apply a variety of instruction level execution analysis methods to make a final deposition (ROP malware versus false positive). The BT software 130 may be attached to the target application at invocation time or on demand after ROP heuristic event. The actual policy of when the BT software 130 is loaded and attached into the application may be implementation or policy specific.

[0071] The BT software 130 may expose a set of standardized APIs to the AV software 140 to configure the security policy such as an ROP detection policy and call back conditions, etc. Examples of those include low heuristic hardware thresholds for maximum protection coverage, higher thresholds for low overhead, whitelisted applications and address ranges, application specific configurations, etc. With the APIs exposed by the BT software 130, the AV software 140 can customize the ROP heuristic policy to meet specific usage models.

[0072] An ROPh device driver 120 configures the ROP heuristic hardware counters 115 and interfaces with the BT software 130 to signal a branch anomaly event. In implementations where dedicated ROP heuristic hardware counters do not exist on the processor 110, a set of model specific performance monitoring counters may be used to monitor and detect execution

anomalies of branch executions. The BT software 130 may use IOCTL APIs exposed by the ROPh device driver 120 to configure the heuristic hardware configuration 115 to match with the security policy requested from the AV software 140.

[0073] When a ROP heuristic counter is equal or greater than the threshold value, the ROP heuristic hardware may generate a ROP heuristic event (an interrupt) and the ROPh Device Driver 120 may suspend the affected application and signal the event to the BT software 130. When the BT software 130 has yet to be attached to the affected application, the BT software may dynamically be loaded and attached on demand prior to receiving the ROP heuristic event. Upon receiving a ROP heuristic signal, the BT software 130 may also receive the interrupted application register state and the recorded branch history recorded by the LBR 300 of the processor from the ROPh device driver 120.

[0074] Thus, an implementation according to one embodiment comprises the following components, illustrated in FIG. 4: the ROP heuristic hardware 115 (which in one embodiment includes general purpose PMU hardware 200 and LBR hardware 300); ROP heuristic device driver 120; BT software 130, and AV software 140. Details of embodiments of each are provided in FIGs. 5–8 and described below.

[0075] FIG. 5 is a flowchart illustrating the actions taken by the ROP detection hardware 115 according to one embodiment. In block 510, the counters (such as the PMU counters 230 of FIG. 2) continuously count branch anomalies. In some embodiments, additional computational resources may be deployed to compute for branch anomaly heuristic modeling. The counters may be configured from the ROP heuristic device driver 120, described further below with reference to FIG. 6. In block 520, the hardware determines that the counters have a predetermined relationship with a predetermined threshold value. For example, the hardware may determine whether the counter is greater than equal to the threshold value, and if so in block 530 generate an ROP heuristic interrupt to the ROP heuristic device driver. Other relationships, such as greater than, instead of greater than or equal may be used. Alternate embodiments may count down from the predetermined threshold value, detecting the event of reaching zero. Although illustrated in FIG. 5 as an interrupt fielded by the ROP device driver at block 630, any type of event signaling technology may be used.

[0076] FIG. 6 is a flowchart illustrating the behavior of an ROP heuristic device driver 120 according to one embodiment. In block 610, the device driver 120 may be used to configure

the ROP detection hardware 115 as described in block 510 above. In block 620, the ROP heuristic device driver 120 receives an interrupt or other signal from the ROP detection hardware 115. In one embodiment, if the interrupt occurs during a known safe condition, for example, the interrupted application or an address range corresponding to the interrupt is listed in a whitelist, the ROP heuristic device driver 120 may configure the ROP detection hardware 115 appropriately. Any desired configuration in such a situation may be performed, such as disabling the ROP detection feature, changing the threshold value, or resetting counters in the ROP heuristic hardware 115. If the interrupt occurs while not in a known safe condition as determined in block 620, then in block 630 the ROP device driver 120 may capture or collect available information useful for the BT analysis. In one embodiment, this useful information may include information such as a context of the interrupted application such as register states, last branch record history, etc. For example, the ROP device driver 120 may read a block of memory at the address of the interrupt, read the content of the LBR stack 300, and read the content of memory pointed to by the LBR entries (from and to addresses). Any available information may be collected as desired. In some embodiments, the collected information may be stored in an ROP detection cache in memory. In other embodiments, the collected information may be written to a file in a file system. Any technique for capturing and preserving the collected information for use by the BT software 130 may be used. Any or all of the collected may be directly provided to the BT software as part of the control flow from the ROP heuristic device driver 120 or may be indirectly provided by the ROP heuristic device driver storing the information and providing access to that stored information to the BT software 130.

[0077] In block 640 the ROP heuristic device driver 120 may then signal an ROP heuristic event to the BT software 130, indicating a branch anomaly that may mean an ROP exploit has been detected. Any desired technique for signaling the BT software 130 may be used. In some embodiments, the BT software 130 may be loaded on demand and attached to the interrupted application prior to receiving the ROP heuristic event notification from the ROP heuristic device driver 120.

[0078] In some embodiments, the ROPh device driver 120 may be implemented as part of the OS kernel. However, in some embodiments, at least some of the driver 120 may be implemented in user mode. The BT software 130 is typically a user mode program, but parts of the BT software 130 may be implemented in the OS if desired.

[0079] FIG 7 is a flowchart illustrating one embodiment of the BT software 130. In block 710, upon receiving a signal from the ROPh device driver 120, the BT software 130 determines whether branch history data is available. If the branch history data is available, then in block 720 the BT software 130 performs branch history based binary analysis.

[0080] The branch history based analysis may include, for example, 1) whether the recorded RET target address has a preceding CALL instruction, 2) a check for known ROP gadget behavioral signatures (e.g., a return to critical APIs). These analyses are illustrative and by way of example only, and other branch history based analysis may be performed as desired. The recorded CALL address history may also be used for conducting CALL/RET matching during the forward execution analysis. Other types of backward code inspection may also be performed, including static code flow analysis, looking at previous LBR entries prior to the anomaly, or processor trace records.

[0081] In block 730, if the branch history based analysis can determine that the possible exploit is a malware attack, then the BT software in block 740 may suspend the affected application and notify the AV software than an ROP exploit has been detected.

[0082] When the history-based analysis is not available or is not sufficient to determine the positive-ness of malware attack, the BT software 130 may perform instruction level forward execution analysis in block 750 to make an ROP event disposition as an ROP exploit or a false positive. Instruction level execution runtime analysis may be accomplished by regenerating the binary with security check instrumentations, producing a translated version of the original binary, prior to execution. The translation is dynamically produced and may be stored into a translated cache storage on an as-needed basis by the BT software 130. Once translated, the execution of the program occurs from the translated version of the original code until there is a translation cache miss.

[0083] Instruction level forward execution analysis continuously follows the execution after the ROP heuristic event and may perform the same binary level analysis as the history based analysis listed above, applying the analysis dynamically during execution. In addition, forward execution analysis may conduct 1) dynamic branch target validation, 2) dynamic CALL/RET matching check using CALL history when available, 3) stack pivot check, 4) detection of AV evasion (hook hopping), etc.

[0084] In one embodiment, to reduce the overhead of instruction level execution analysis, the BT software 130 may generate a translated version of the original binary with in-lined security checks, thus eliminating frequent exits (context switches) to the BT software 130 for analysis. This may involve the replacement of certain instructions, such as indirect branch instructions, with stub instructions that when executed cause instrumentation code of the BT software 130 to execute. As part of translating the original binary, the target addresses of direct branch instructions may be directly incorporated into those instructions as portions of routines are placed in a translation cache. For indirect branch instructions that incorporate an identifier of their intended targets, the translation may attempt to determine the their target addresses by using those identifiers to refer to (for example) an entry point table for a library, allowing the target addresses to be retrieved. Alternately, the translation performed as part of instruction level execution analysis may retrieve indications of target addresses known to be valid from other tables, such as whitelists. One or more sizable libraries are commonly included with an operating system that may normally be stored at predictable addresses locations in storages across a variety of computing devices. Due to the reliance of what may be a great many other routines on the function routines of those libraries, those libraries typically change only very infrequently for fear of causing unforeseen and undesirable effects on the other routines that use them. Thus constructing a viable whitelist of valid target addresses for valid entry points of the function routines to which a branch instruction may validly jump becomes feasible. Where libraries are not consistently stored at predictable addresses, similar techniques may derive tables of valid addresses, such as valid entry points, once the addresses of the libraries become known. Although referred to as tables, storage of valid address information may be implemented as any of a variety of types of data structures in which indications of valid target addresses may be stored in any of a variety of formats.

[0085] If the target address derived during execution of the translated portion of the code placed in the translation cache is modified, a lookup of the target address in the tables of valid addresses may not generate a match. This may be an indication of an ROP attack. In some embodiments, both target addresses of CALL and RET instructions may be analyzed as described above, because a hallmark of ROP attacks is to modify the target addresses of return instructions as part of causing execution to jump among pieces of legitimate code to perform a malicious function.

[0086] Where an indirect branch instruction is a CALL instruction to a function routine that is expected to later be paired with a RET instruction following completion of that function routine (e.g., a function routine of a library), the BT software 130 may additionally derive the return target address of that RET instruction and store that return target address in a return look-up table as a known valid target address for that RET instruction. During normal operation, the return target address should indicate the address of the instruction immediately following the CALL instruction in whatever routine in which the CALL instruction exists. Given that the valid target address of a RET instruction that corresponds to a CALL instruction is invariably the address of whatever instruction that immediately follows that CALL instruction, it is possible to determine the valid target address of the corresponding RET instruction with certainty.

[0087] If the target of an indirect branch instruction derived during execution up to a stub instruction matches a known valid target address (however derived), then that target address to which the indirect branch instruction would direct the flow of execution is deemed to have been validated, and that indirect branch instruction may be allowed to execute. If the target address does not match, that may be considered an invalid target address, and thus may positively indicate the presence of an ROP exploit to be signaled to the AV software 140.

[0088] In some embodiments, the BT software 130 may monitor the frequency with which RET instructions occur to determine whether a ROP attack may be underway. Since ROP attacks are based on repeatedly manipulating target addresses associated with RET instructions to direct a flow of execution about portions of legitimate code to perform a malicious function, ROP often results in the frequent execution of RET instructions. Example techniques may involve repeatedly comparing the frequency with which return instructions are executed to determine whether that frequency exceeds a threshold of frequency of execution of RET instructions such that the observed use of RET instructions is deemed consistent with a ROP attack in progress. Such a threshold may be specified in any desired way, such as a number of executions of RET instructions per a specified time period or per a specified number of clock or processing cycles of the 110.

[0089] As has been discussed, it is envisioned that routines are translated a portion at a time, with each translated portion being placed in a translation cache where the processor component later executes it. The BT software 130 may further apply dynamic profiling

techniques, identifying frequently executed execution paths and removing security checks for already validated paths. As a result, the instruction level execution analysis can be performed at almost original binary execution speed.

[0090] If the instruction level execution analysis determines in block 760 that an ROP attack has been detected, then in block 740 the BT software 130 may suspend the affected application and notify the AV software 140. Any desired technique for notifying the AV software 140 may be used in block 740, such as calls to an API defined by the AV software 140. Any desired collection of information may be passed from the BT software 130 to the AV software 140, including any or all of the information obtained from the hardware 115 via the device driver 120.

[0091] As illustrated in block 770, in one embodiment, after predetermined threshold condition corresponding to the forward execution analysis is provided, the BT software 130 may switch back to the native program execution in block 780 if no malware attack is detected. The predetermined threshold may be specified in any desired way, for example, a number of events, a number of blocks, a predetermined time, etc. If the threshold condition has not been reached, then the BT software 130 may continue the forward execution analysis, continuing to execute the instrumented application binary in order to try to determine whether an ROP exploit is underway.

[0092] Advanced analytics in addition may take into account additional contextual data and implement extra checks based on other factors, such as:

[0093] 1. Distribution of from/to addresses.

[0094] 2. Uniqueness of from, to and from/to addresses.

[0095] 3. Matching of from/to addresses and other PMU counters to a distribution that characterizes the specific process (software fingerprinting).

[0096] By taking into account the address of the instruction causing the PMI (which is stored on the stack) raised by reaching threshold of counter, the BT software 130 may determine which process was responsible for the PMI, and may limit the analysis to specific monitored processes. For example, the BT software 130 may filter only addresses belonging to the address space of the monitored process. In some embodiments, the data about process location in memory is available from the OS thru Process Walking or enumerating processes.

Embodiments may exclude certain processes to suppress incorrect detections or to improve system performance. The BT software 130 may analyze the time sequence of specific counters for a selected process as well as the distribution of the addresses of instructions causing those events. In addition, the distribution of branch misprediction instructions may be used to form a software fingerprint.

[0097] FIG. 8 is a flowchart illustrating the AV software 140 according to one embodiment. The AV software 140 in block 810 may perform further inspection based upon the reported information and decide on what actions need to be taken. For example, the AV software 140 may decide to terminate the affected application. In some cases, the AV software 140 may decide to take further security actions because of the analysis, such as disabling vulnerable software components such as plug-ins and blacklisting URLs associated with the affected application.

[0098] The anti-malware software 140 may also take subsequent actions to disposition the code in which the exploit is detected, such as terminating the process and attempting to perform a repair to remove the software from the platform, notifying a backend server of this detection, monitoring the application in a sandbox for isolation to learn more about its behavior, or if the application is legitimate, allow the program to continue after restarting execution.

[0099] Referring now to FIG. 9, a block diagram illustrates a programmable device 900 that may be used for implementing the techniques described herein in accordance with one embodiment. The programmable device 900 illustrated in FIG. 9 is a multiprocessor programmable device that includes a first processing element 970 and a second processing element 980. While two processing elements 970 and 980 are shown, an embodiment of programmable device 900 may also include only one such processing element.

[0100] Programmable device 900 is illustrated as a point-to-point interconnect system, in which the first processing element 970 and second processing element 980 are coupled via a point-to-point interconnect 950. Any or all of the interconnects illustrated in FIG. 9 may be implemented as a multi-drop bus rather than point-to-point interconnects.

[0101] As illustrated in FIG. 9, each of processing elements 970 and 980 may be multicore processors, including first and second processor cores (i.e., processor cores 974a and 974b and processor cores 984a and 984b). Such cores 974a, 974b, 984a, 984b may be configured to execute instruction code. However, other embodiments may use processing elements that are

single core processors as desired. In embodiments with multiple processing elements **970**, **980**, each processing element may be implemented with different numbers of cores as desired.

[0102] Each processing element **970**, **980** may include at least one shared cache **946**. The shared cache **946a**, **946b** may store data (e.g., instructions) that are utilized by one or more components of the processing element, such as the cores **974a**, **974b** and **984a**, **984b**, respectively. For example, the shared cache may locally cache data stored in a memory **932**, **934** for faster access by components of the processing elements **970**, **980**. In one or more embodiments, the shared cache **946a**, **946b** may include one or more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, a last level cache (LLC), or combinations thereof.

[0103] While FIG. 9 illustrates a programmable device with two processing elements **970**, **980** for clarity of the drawing, the scope of the present invention is not so limited and any number of processing elements may be present. Alternatively, one or more of processing elements **970**, **980** may be an element other than a processor, such as an graphics processing unit (GPU), a digital signal processing (DSP) unit, a field programmable gate array, or any other programmable processing element. Processing element **980** may be heterogeneous or asymmetric to processing element **970**. There may be a variety of differences between processing elements **970**, **980** in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like. These differences may effectively manifest themselves as asymmetry and heterogeneity amongst processing elements **970**, **980**. In some embodiments, the various processing elements **970**, **980** may reside in the same die package.

[0104] First processing element **970** may further include memory controller logic (MC) **972** and point-to-point (P-P) interconnects **976** and **978**. Similarly, second processing element **980** may include a MC **982** and P-P interconnects **986** and **988**. As illustrated in FIG. 9, MCs **972** and **982** couple processing elements **970**, **980** to respective memories, namely a memory **932** and a memory **934**, which may be portions of main memory locally attached to the respective processors. While MC logic **972** and **982** is illustrated as integrated into processing elements **970**, **980**, in some embodiments the memory controller logic may be discrete logic outside processing elements **970**, **980** rather than integrated therein.

[0105] Processing element **970** and processing element **980** may be coupled to an I/O subsystem **990** via respective P-P interconnects **976** and **986** through links **952** and **954**. As illustrated in FIG. 9, I/O subsystem **990** includes P-P interconnects **994** and **998**. Furthermore, I/O subsystem **990** includes an interface **992** to couple I/O subsystem **990** with a high performance graphics engine **938**. In one embodiment, a bus (not shown) may be used to couple graphics engine **938** to I/O subsystem **990**. Alternately, a point-to-point interconnect **939** may couple these components.

[0106] In turn, I/O subsystem **990** may be coupled to a first link **916** via an interface **996**. In one embodiment, first link **916** may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another I/O interconnect bus, although the scope of the present invention is not so limited.

[0107] As illustrated in FIG. 9, various I/O devices **914**, **924** may be coupled to first link **916**, along with a bridge **918** that may couple first link **916** to a second link **920**. In one embodiment, second link **920** may be a low pin count (LPC) bus. Various devices may be coupled to second link **920** including, for example, a keyboard/mouse **912**, communication device(s) **926** (which may in turn be in communication with the computer network **903**), and a data storage unit **928** such as a disk drive or other mass storage device which may include code **930**, in one embodiment. The code **930** may include instructions for performing embodiments of one or more of the techniques described above. Further, an audio I/O **924** may be coupled to second link **920**.

[0108] Note that other embodiments are contemplated. For example, instead of the point-to-point architecture of FIG. 9, a system may implement a multi-drop bus or another such communication topology. Although links **916** and **920** are illustrated as busses in FIG. 9, any desired type of link may be used. In addition, the elements of FIG. 9 may alternatively be partitioned using more or fewer integrated chips than illustrated in FIG. 9.

[0109] Referring now to FIG. 10, a block diagram illustrates a programmable device **1000** according to another embodiment. Certain aspects of FIG. 10 have been omitted from FIG. 10 in order to avoid obscuring other aspects of FIG. 10.

[0110] FIG. 10 illustrates that processing elements **1070**, **1080** may include integrated memory and I/O control logic ("CL") **1072** and **1082**, respectively. In some embodiments, the **1072**, **1082** may include memory control logic (MC) such as that described above in connection

with FIG. 9. In addition, CL 1072, 1082 may also include I/O control logic. FIG. 10 illustrates that not only may the memories 1032, 1034 be coupled to the CL 1072, 1082, but also that I/O devices 1044 may also be coupled to the control logic 1072, 1082. Legacy I/O devices 1015 may be coupled to the I/O subsystem 1090 by interface 1096. Each processing element 1070, 1080 may include multiple processor cores, illustrated in FIG. 10 as processor cores 1074A, 1074B, 1084A and 1084B. As illustrated in FIG. 10, I/O subsystem 1090 includes point-to-point (P-P) interconnects 1094 and 1098 that connect to P-P interconnects 1076 and 1086 of the processing elements 1070 and 1080 with links 1052 and 1054. Processing elements 1070 and 1080 may also be interconnected by link 1050 and interconnects 1078 and 1088, respectively.

[0111] The programmable devices depicted in FIGs. 9 and 10 are schematic illustrations of embodiments of programmable devices that may be utilized to implement various embodiments discussed herein. Various components of the programmable devices depicted in FIGs. 9 and 10 may be combined in a system-on-a-chip (SoC) architecture.

[0112] The combined hardware and software techniques described above may be implemented as part of any desired type of anti-malware system, such as an intrusion protection system, to detect relatively difficult-to-detect ROP without the need for a specific signature of the exploit and with less performance impact than a purely software-based technique as has been discussed in the literature previously. Furthermore, proper design of the analytical engine may avoid the negative impact of false positives in the analysis.

[0113] The following examples pertain to further embodiments.

[0114] Example 1 is a machine readable medium, on which are stored instructions, comprising instructions that when executed cause a programmable device to: configure a processor of the programmable device to detect code flow anomalies; receive an indication of a code flow anomaly from the processor; perform binary translation of a portion of an application corresponding to the code flow anomaly; and detect a return-oriented programming exploit responsive to the binary translation.

[0115] In Example 2 the subject matter of Example 1 optionally includes wherein the instructions further comprise instructions that when executed cause the programmable device to: notify an anti-malware software of the return-oriented programming exploit.

[0116] In Example 3 the subject matter of any of Examples 1 or 2 optionally includes wherein the instructions that when executed cause the programmable device to configure the processor comprise instructions that when executed cause the programmable device to: configure a performance monitoring unit of the processor to detect mispredicted branches.

[0117] In Example 4 the subject matter of any of Examples 1 or 2 optionally includes wherein the instructions that when executed cause the programmable device to configure the processor comprise instructions that when executed cause the programmable device to: configure a return-oriented programming heuristic unit of the processor to detect code flow anomalies that indicate a possible return-oriented programming exploit.

[0118] In Example 5 the subject matter of any of Examples 1 or 2 optionally includes wherein the instructions that when executed cause the programmable device to receive an indication of a code flow anomaly from the processor comprise instructions that when executed cause the programmable device to: receive an interrupt indicating a code flow anomaly from the processor.

[0119] In Example 6 the subject matter of any of Examples 1 or 2 optionally includes wherein the instructions that when executed cause the programmable device to receive an indication of a code flow anomaly from the processor comprise instructions that when executed cause the programmable device to: receive the indication of the code flow anomaly from the processor in a device driver; collect information by the device driver corresponding to the code flow anomaly; and provide the indication of the code flow anomaly and the collected information to a binary translation module.

[0120] In Example 7 the subject matter of any of Examples 1 or 2 optionally includes wherein the instructions that when executed cause the programmable device to perform binary translation comprise instructions that when executed cause the programmable device to: analyze branch history information corresponding to the code flow anomaly;

[0121] perform forward instruction analysis; and ignore code flow anomalies that do not indicate a return-oriented programming exploit.

[0122] In Example 8 the subject matter of any of Examples 1 or 2 optionally includes wherein the instructions further comprise instructions that when executed cause the programmable device to: take an anti-malware action responsive to a determination that a return-oriented programming exploit has occurred.

[0123] Example 9 is a programmable device programmed to detect malware exploits, comprising: a processor, comprising: a code flow anomaly detector; a device driver module; and a binary translation module, wherein the device driver module is configured to: receive an interrupt from the code flow anomaly detector; and configure the code flow anomaly detector; and notify the binary translation module of the code flow anomaly, and wherein the binary translation module is configured to: receive the notification of the code flow anomaly from the device driver module; perform binary translation of an application corresponding to the code flow anomaly; and detect a return-oriented programming exploit responsive to the binary translation.

[0124] In Example 10 the subject matter of Example 9 optionally includes wherein the code flow anomaly is a code flow anomaly indicating a possible return-oriented programming exploit.

[0125] In Example 11 the subject matter of Example 9 optionally includes further comprising: an anti-malware module, configured to: receive a notification of a return-oriented programming exploit from the binary translation software; and take an action responsive to the notification.

[0126] In Example 12 the subject matter of any of Examples 9–11 optionally includes wherein the code flow anomaly detector comprises: a performance monitoring unit, configured to detect mispredicted branches; and a last branch record unit, configured to capture information regarding branches.

[0127] In Example 13 the subject matter of any of Examples 9–11 optionally includes wherein the code flow anomaly detector comprises: heuristic-based hardware configured to detect return-oriented programming code flow anomalies.

[0128] In Example 14 the subject matter of any of Examples 9–11 optionally includes wherein the device driver module is further configured to: collect information corresponding to the code flow anomaly; and provide the collected information to the binary translation module.

[0129] In Example 15 the subject matter of any of Examples 9–11 optionally includes wherein binary translation module is further configured to suspend an application corresponding to the code flow anomaly responsive to detection of a return-oriented programming exploit.

[0130] In Example 16 the subject matter of any of Examples 9–11 optionally includes binary translation module is configured to detect a return-oriented programming exploit responsive to: analysis of a branch history corresponding to the code flow anomaly; and forward instruction analysis of code of an application associated with the code flow anomaly.

[0131] Example 17 is a method of detecting malware exploits, comprising: detecting code flow anomalies by a code flow anomaly detector unit of a processor; generating an indication of a code flow anomaly; performing binary translation of a portion of an application corresponding to the code flow anomaly; and detecting a return-oriented programming exploit responsive to the binary translation.

[0132] In Example 18 the subject matter of Example 17 optionally includes wherein detecting code flow anomalies comprises: detecting mispredicted branches.

[0133] In Example 19 the subject matter of Example 18 optionally includes wherein detecting code flow anomalies further comprises: detecting an imbalance between call and return instructions.

[0134] In Example 20 the subject matter of Example 17 optionally includes further comprising: suspending the application; and notifying an anti-malware software of the detected return-oriented programming exploit.

[0135] In Example 21 the subject matter of any of Examples 17–20 optionally includes wherein performing binary translation comprises: analyzing a branch history of the application.

[0136] In Example 22 the subject matter of any of Examples 17–20 optionally includes wherein performing binary translation comprises: performing instruction-level forward analysis of the application.

[0137] In Example 23 the subject matter of any of Examples 17–20 optionally includes further comprising: configuring the code flow anomaly detector unit by a device driver.

[0138] In Example 24 the subject matter of any of Examples 17–20 optionally includes wherein the indication of the code flow anomaly comprises an interrupt generated by the code flow anomaly detector unit.

[0139] In Example 25 the subject matter of any of Examples 17–20 optionally includes further comprising: taking an anti-malware action responsive to detecting a return-oriented programming exploit.

[0140] Example 26 is a programmable device for detecting malware, comprising: means for configuring a processor of the programmable device to detect code flow anomalies; means for receiving an indication of a code flow anomaly from the processor; means for performing binary translation of a portion of an application corresponding to the code flow anomaly; and means for detecting a return-oriented programming exploit responsive to the binary translation.

[0141] In Example 27 the subject matter of Example 26 optionally includes wherein further comprising: means for notifying an anti-malware software of the return-oriented programming exploit.

[0142] In Example 28 the subject matter of any of Examples 26 or 27 optionally includes wherein the means for configuring the processor comprise: means for configuring a performance monitoring unit of the processor to detect mispredicted branches.

[0143] In Example 29 the subject matter of any of Examples 26 or 27 optionally includes wherein the means for configuring the processor comprise: means for configuring a return-oriented programming heuristic unit of the processor to detect code flow anomalies that indicate a possible return-oriented programming exploit.

[0144] In Example 30 the subject matter of any of Examples 26 or 27 optionally includes wherein the means for receiving an indication of a code flow anomaly from the processor comprise: means for receiving an interrupt indicating a code flow anomaly from the processor.

[0145] In Example 31 the subject matter of any of Examples 26 or 27 optionally includes wherein the means for receiving an indication of a code flow anomaly from the processor comprise: means for receiving the indication of the code flow anomaly from the processor in a device driver; means for collecting information by the device driver corresponding to the code flow anomaly; and means for providing the indication of the code flow anomaly and the collected information to a binary translation module.

[0146] In Example 32 the subject matter of any of Examples 26 or 27 optionally includes wherein the means for performing binary translation comprise: means for analyzing branch history information corresponding to the code flow anomaly;

[0147] means for performing forward instruction analysis; and means for ignoring code flow anomalies that do not indicate a return-oriented programming exploit.

[0148] In Example 33 the subject matter of any of Examples 26 or 27 optionally includes further comprising: means for taking an anti-malware action responsive to a determination that a return-oriented programming exploit has occurred.

[0149] Example 34 is a programmable device programmed to detect malware exploits, comprising: a processor, comprising: a code flow anomaly detector; a device driver module; and a binary translation module, wherein the device driver module is configured to: receive an interrupt from the code flow anomaly detector; and configure the code flow anomaly detector; and collect information corresponding to a code flow anomaly; provide the collected information to the binary translation module; notify the binary translation module of the code flow anomaly, and wherein the binary translation module is configured to: receive the notification of the code flow anomaly from the device driver module; perform binary translation of an application corresponding to the code flow anomaly; and detect a return-oriented programming exploit responsive to the binary translation.

[0150] In Example 35 the subject matter of Example 34 optionally includes further comprising: an anti-malware module, configured to: receive a notification of a return-oriented programming exploit from the binary translation software; and take an action responsive to the notification.

[0151] In Example 36 the subject matter of any of Examples 34–35 optionally includes wherein the code flow anomaly detector comprises: a performance monitoring unit, configured to detect mispredicted branches; and a last branch record unit, configured to capture information regarding branches.

[0152] In Example 37 the subject matter of any of Examples 34–35 optionally includes wherein the code flow anomaly detector comprises: heuristic-based hardware configured to detect return-oriented programming code flow anomalies.

[0153] In Example 38 the subject matter of any of Examples 34–35 optionally includes binary translation module is configured to detect a return-oriented programming exploit responsive to: analysis of a branch history corresponding to the code flow anomaly; and forward instruction analysis of code of an application associated with the code flow anomaly.

[0154] It is to be understood that the above description is intended to be illustrative, and not restrictive. For example, the above-described embodiments may be used in combination with each other. Many other embodiments will be apparent to those of skill in the art upon

reviewing the above description. The scope of the invention therefore should be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

CLAIMS

What is claimed is:

1. A machine readable medium, on which are stored instructions, comprising instructions that when executed cause a programmable device to:
 - configure a processor of the programmable device to detect code flow anomalies;
 - receive an indication of a code flow anomaly from the processor;
 - perform binary translation of a portion of an application corresponding to the code flow anomaly; and
 - detect a return-oriented programming exploit responsive to the binary translation.
2. The machine readable medium of claim 1, wherein the instructions further comprise instructions that when executed cause the programmable device to:
 - notify an anti-malware software of the return-oriented programming exploit.
3. The machine readable medium of any of claims 1 or 2, wherein the instructions that when executed cause the programmable device to configure the processor comprise instructions that when executed cause the programmable device to:
 - configure a performance monitoring unit of the processor to detect mispredicted branches.
4. The machine readable medium of any of claims 1 or 2, wherein the instructions that when executed cause the programmable device to configure the processor comprise instructions that when executed cause the programmable device to:
 - configure a return-oriented programming heuristic unit of the processor to detect code flow anomalies that indicate a possible return-oriented programming exploit.
5. The machine readable medium of any of claims 1 or 2, wherein the instructions that when executed cause the programmable device to receive an indication of a code flow

anomaly from the processor comprise instructions that when executed cause the programmable device to:

receive an interrupt indicating a code flow anomaly from the processor.

6. The machine readable medium of any of claims 1 or 2, wherein the instructions that when executed cause the programmable device to receive an indication of a code flow anomaly from the processor comprise instructions that when executed cause the programmable device to:

receive the indication of the code flow anomaly from the processor in a device driver;

collect information by the device driver corresponding to the code flow anomaly; and

provide the indication of the code flow anomaly and the collected information to a binary translation module.

7. The machine readable medium of any of claims 1 or 2, wherein the instructions that when executed cause the programmable device to perform binary translation comprise instructions that when executed cause the programmable device to:

analyze branch history information corresponding to the code flow anomaly;

perform forward instruction analysis; and

ignore code flow anomalies that do not indicate a return-oriented programming exploit.

8. The machine readable medium of any of claims 1 or 2, wherein the instructions further comprise instructions that when executed cause the programmable device to:

take an anti-malware action responsive to a determination that a return-oriented programming exploit has occurred.

9. A programmable device programmed to detect malware exploits, comprising:
a processor, comprising:

a code flow anomaly detector;

a device driver module; and

a binary translation module,

wherein the device driver module is configured to:

- receive an interrupt from the code flow anomaly detector; and
- configure the code flow anomaly detector; and
- notify the binary translation module of the code flow anomaly, and

wherein the binary translation module is configured to:

- receive the notification of the code flow anomaly from the device driver module;
- perform binary translation of an application corresponding to the code flow anomaly; and
- detect a return-oriented programming exploit responsive to the binary translation.

10. The programmable device of claim 9, wherein the code flow anomaly is a code flow anomaly indicating a possible return-oriented programming exploit.

11. The programmable device of claim 9, further comprising:
an anti-malware module, configured to:

- receive a notification of a return-oriented programming exploit from the binary translation software; and
- take an action responsive to the notification.

12. The programmable device of any of claims 9–11, wherein the code flow anomaly detector comprises:

- a performance monitoring unit, configured to detect mispredicted branches; and
- a last branch record unit, configured to capture information regarding branches.

13. The programmable device of any of claims 9–11, wherein the code flow anomaly detector comprises:

- heuristic-based hardware configured to detect return-oriented programming code flow anomalies.

14. The programmable device of any of claims 9–11, wherein the device driver module is further configured to:

collect information corresponding to the code flow anomaly; and
provide the collected information to the binary translation module.

15. The programmable device of any of claims 9–11, wherein binary translation module is further configured to suspend an application corresponding to the code flow anomaly responsive to detection of a return-oriented programming exploit.

16. The programmable device of any of claims 9–11, binary translation module is configured to detect a return-oriented programming exploit responsive to:

analysis of a branch history corresponding to the code flow anomaly; and
forward instruction analysis of code of an application associated with the code flow anomaly.

17. A method of detecting malware exploits, comprising:
detecting code flow anomalies by a code flow anomaly detector unit of a processor;
generating an indication of a code flow anomaly;
performing binary translation of a portion of an application corresponding to the code flow anomaly; and
detecting a return-oriented programming exploit responsive to the binary translation.

18. The method of claim 17, wherein detecting code flow anomalies comprises:
detecting mispredicted branches.

19. The method of claim 18, wherein detecting code flow anomalies further comprises:
detecting an imbalance between call and return instructions.

20. The method of claim 17, further comprising:
suspending the application; and
notifying an anti-malware software of the detected return-oriented programming exploit.

21. The method of any of claims 17–20, wherein performing binary translation comprises:
analyzing a branch history of the application.
22. The method of any of claims 17–20, wherein performing binary translation comprises:
performing instruction-level forward analysis of the application.
23. The method of any of claims 17–20, further comprising:
configuring the code flow anomaly detector unit by a device driver.
24. The method of any of claims 17–20, wherein the indication of the code flow anomaly comprises an interrupt generated by the code flow anomaly detector unit.
25. The method of any of claims 17–20, further comprising:
taking an anti-malware action responsive to detecting a return-oriented programming exploit.

1/10




FIG. 1





FIG. 2
(Prior Art)

FIG. 3
(Prior Art)







FIG. 4







FIG. 6







FIG. 8






FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/053889**A. CLASSIFICATION OF SUBJECT MATTER****G06F 21/56(2013.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 21/56; G06F 9/54; G06F 21/79; G06F 21/54; G06F 12/08; G06F 21/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & keywords: code flow anomaly, binary translation, return-oriented programming exploit

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2014-0372701 A1 (QUALCOMM INCORPORATED) 18 December 2014 See paragraphs [0068]-[0076]; and figure 7.	1-25
Y	US 2015-0067763 A1 (GREGORY WILLIAM DALCHER et al.) 05 March 2015 See paragraphs [0012]-[0024]; claims 23-29; and figure 1.	1-25
A	US 2012-0167120 A1 (DAAVID HENTUNEN) 28 June 2012 See paragraphs [0032]-[0047]; and figures 2-3.	1-25
A	US 2014-0229717 A1 (ASHISH VENKAT et al.) 14 August 2014 See paragraphs [0028]-[0030]; and figures 4-5.	1-25
A	WO 2014-152469 A1 (THE TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK) 25 September 2014 See paragraphs [0047]-[0053]; and figure 1.	1-25



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

10 January 2017 (10.01.2017)

Date of mailing of the international search report

11 January 2017 (11.01.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea



Facsimile No. +82-42-481-8578

Authorized officer

CHIN, Sang Bum

Telephone No. +82-42-481-8398



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/053889

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0372701 A1	18/12/2014	CN 103946855 A EP 2776971 A2 EP 3062259 A1 IN 3071CHN2014 A JP 05944520 B2 JP 2014-532944 A US 2013-0117843 A1 US 8839429 B2 US 9262627 B2 WO 2013-070773 A2 WO 2013-070773 A3	23/07/2014 17/09/2014 31/08/2016 31/07/2015 05/07/2016 08/12/2014 09/05/2013 16/09/2014 16/02/2016 16/05/2013 12/12/2013
US 2015-0067763 A1	05/03/2015	CN 105408911 A EP 3039608 A1 WO 2015-030748 A1	16/03/2016 06/07/2016 05/03/2015
US 2012-0167120 A1	28/06/2012	GB 2499932 A US 2015-0161396 A1 US 8997218 B2 WO 2012-084507 A1	04/09/2013 11/06/2015 31/03/2015 28/06/2012
US 2014-0229717 A1	14/08/2014	CN 104903906 A EP 2956885 A1 EP 2956885 A4 US 9135435 B2 WO 2014-126762 A1	09/09/2015 23/12/2015 05/10/2016 15/09/2015 21/08/2014
WO 2014-152469 A1	25/09/2014	CN 105229612 A CN 105247532 A KR 10-2015-0138229 A KR 10-2016-0008509 A US 2016-0275289 A1 WO 2014-149080 A1	06/01/2016 13/01/2016 09/12/2015 22/01/2016 22/09/2016 25/09/2014