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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
31 December 2014 (31.12.2014)

WIPO | PCT

(10) International Publication Number
WO 2014/210277 A1

(51) International Patent Classification:
G06F 21/71 (2013.01)

(21) International Application Number:
PCT/US2014/044284

(22) International Filing Date:
26 June 2014 (26.06.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/841,014 28 June 2013 (28.06.2013) US

(71) Applicant: **THE TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK** [US/US]; Mail Code 4308, 535 West 116th Street, New York, New York 10027 (US).

(72) Inventors: **SETHUMADHAVAN, Lakshminarasimhan**; MC 0401, 1214 Amsterdam Avenue, New York, New York 10027 (US). **SINHA, Kanad**; MC 0401, 1214 Amsterdam Avenue, New York, New York 10027 (US). **KEROMYTIS, Angelos**; MC 0401, 1214 Amsterdam Avenue, New York, New York 10027 (US). **PAPPAS, Vasilis**; MC 0401, 1214 Amsterdam Avenue, New York, New York 10027 (US). **KEMERLIS, Vasileios**; MC 0401, 1214 Amsterdam Avenue, New York, New York 10027 (US).

(74) Agents: **RABINOVITCH, Ido** et al.; Hunter Clark PLLC, Suite 213-T, 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))*

[Continued on next page]

(54) Title: DIVERSIFIED INSTRUCTION SET PROCESSING TO ENHANCE SECURITY

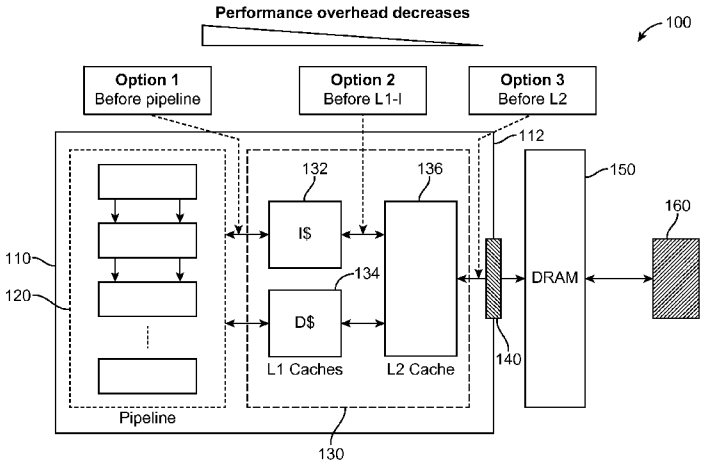

FIG. 1

(57) Abstract: Disclosed are devices, systems, apparatus, methods, products, and other implementations, including a method that includes receiving a block of information from non-processor memory at an interface between the non-processor memory and processor memory comprising two or more processor memory levels, determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code, and decrypting the block of information at the interface between the non-processor memory and the processor memory for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code. The block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

WO 2014/210277 A1

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

DIVERSIFIED INSTRUCTION SET PROCESSING TO ENHANCE SECURITY

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of, and priority to, U.S. Provisional Patent Application No. 61/841,014, entitled "PROCESSING OF DIVERSIFIED INSTRUCTION SETS TO ENHANCE SECURITY," and filed June 28, 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] Despite the variety of choices regarding hardware and software, to date a large number of computer systems remain identical. Characteristic examples of this trend are Windows on x86 and Android on ARM. This homogeneity, sometimes referred to as "computing oligoculture", provides a fertile ground for malware in the highly networked world of today.

[0004] One way to counter this problem is to diversify systems so that attackers cannot quickly and easily compromise a large number of machines. For instance, if each system has a different hardware/software interface, the attacker has to invest more time in developing exploits that run on every system manifestation. It is not that each individual attack gets harder, but the spread of malware slows down. Further, if the diversified instruction set architectures (ISA) are kept secret from the attacker, the bar for exploitation is raised even higher.

SUMMARY

[0005] In some variations, a method is provided that includes receiving a block of information from non-processor memory at an interface between the non-processor memory and processor memory comprising two or more processor memory levels, determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code, and decrypting the block of information at the interface between the non-processor memory and the processor memory for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code. The block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

[0006] 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.

[0007] The processor memory may include cache memory of a central processing unit (CPU) organized in two or more cache memory levels, and the non-processor memory may include random access memory.

[0008] Determining whether the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code may include determining whether the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code based on a request sent from the processor for the block of information.

[0009] Determining whether the block of information received from the non-processor memory corresponds to the encrypted instruction code may include determining that the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code when the request for the block of information is associated with a portion of the processor-memory configured to store instructions.

[0010] The method may further include associating an identifier with the received block of information, and setting the identifier to a value indicating one of instructions or data based on the determination of whether the block of information corresponds to the encrypted instruction code received from the non-processor memory.

[0011] The method may further include determining, in response to a fetch request made by an instruction portion of another level of the two or more levels of the processor memory for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of information stored at the one of the two or more levels of the two or more levels of the processor memory corresponds to the instructions, and performing one of, for example: transferring the at least the portion of the block of information stored at the one of the two or more levels of the processor memory to the instruction portion of the other of the two or more levels of the processor memory when the identifier associated with the block of information is determined to correspond to the instructions, or removing the block of information from the processor memory when the identifier associated with the block of information is determined to correspond to the data, and causing the block of information to be retrieved again from the non-processor memory in order to perform a decryption operation on the block of information retrieved again.

[0012] The method may further include determining, in response to a fetch request made by a data portion of another level of the two or more levels of the processor memory for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of information stored at the one of the two or more levels of the processor memory corresponds to the data, and performing one of, for example: transferring the at least the portion of the block of information stored at the one of the two or more levels of the processor memory to the data portion of the other of the two or more levels of the processor memory when the identifier associated with the block of information indicates that the block of information corresponds to the data, or removing the block of information from the processor memory when the identifier associated with the block of information indicates that the block of information correspond to instructions, and causing the block of information to be retrieved again from the non-processor memory in order to be transferred into the processor memory without being decrypted.

[0013] Decrypting the block of information corresponding to the encrypted instruction code at the interface between the processor-memory and the non-processor memory may include decrypting the block of information corresponding to the encrypted instruction code with one or more first keys assigned to a first controller device associated with a particular instruction

set, the one or more first keys being different from at least one key assigned to a second controller device associated with the particular instruction set so as to emulate instruction set diversification between the first controller device and the second controller device.

[0014] Decrypting the block of information corresponding to the encrypted instruction code with the one or more first keys assigned to the first controller device may include decrypting the encrypted instruction code with one or more first symmetric keys according to a counter mode encryption / decryption process. The one or more first symmetric keys may include one of, for example, a single symmetric key used for any encrypted instruction code retrieved by the first controller device, and/or multiple symmetric keys with each of the multiple symmetric key used for respective encrypted instructions from a respective one of different memory pages stored in the non-processor memory coupled to the first controller device.

[0015] In some variations, a system is provided that includes a processor including processor memory with two or more processor memory levels, and a decryption unit at an interface between non-processor memory and the processor memory. The decryption unit is configured to receive a block of information from the non-processor memory, determine whether the block of information received from the non-processor memory corresponds to encrypted instruction code, and decrypt the block of information for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code. The block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

[0016] 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.

[0017] The decryption unit configured to determine whether the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code may be configured to determine whether the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code based on a request sent from the processor for the block of information.

[0018] The decryption unit configured to determine whether the block of information received from the non-processor memory corresponds to the encrypted instruction code may

be configured to determine that the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code when the request for the block of information is associated with a portion of the processor-memory configured to store instructions.

[0019] The decryption unit may further be configured to associate an identifier with the received block of information, and set the identifier to a value indicating one of instructions or data based on the determination of whether the block of information corresponds to the encrypted instruction code received from the non-processor memory.

[0020] The processor may be configured to determine, in response to a fetch request made by an instruction portion of another level of the two or more levels of the processor memory for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of information stored at the one of the two or more levels of the two or more levels of the processor memory corresponds to the instructions, and perform one of, for example: transferring the at least the portion of the block of information stored at the one of the two or more levels of the processor memory to the instruction portion of the other of the two or more levels of the processor memory when the identifier associated with the block of information is determined to correspond to the instructions, or removing the block of information from the processor memory when the identifier associated with the block of information is determined to correspond to the data, and causing the block of information to be retrieved again from the non-processor memory in order to perform a decryption operation on the block of information retrieved again.

[0021] The processor may be configured to determine, in response to a fetch request made by a data portion of another level of the two or more levels of the processor memory for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of information stored at the one of the two or more levels of the processor memory corresponds to the data, and perform one of, for example, transferring the at least the portion of the block of information stored at the one of the two or more levels of the processor memory to the data portion of the other of the two or more levels of the processor memory when the identifier associated with the block of information indicates that the block of information corresponds to the data, or removing the block of information from the processor memory when the identifier associated

with the block of information indicates that the block of information correspond to instructions, and causing the block of information to be retrieved again from the non-processor memory in order to be transferred into the processor memory without being decrypted.

[0022] The decryption unit configured to decrypt the block of information corresponding to the encrypted instruction code at the interface between the processor-memory and the non-processor memory may be configured to decrypt the block of information corresponding to the encrypted instruction code with one or more first keys assigned to a first controller device associated with a particular instruction set, the one or more first keys being different from at least one key assigned to a second controller device associated with the particular instruction set so as to emulate instruction set diversification between the first controller device and the second controller device. The decryption unit configured to decrypt the block of information corresponding to the encrypted instruction code with the one or more first keys assigned to the first controller device may be configured to decrypt the encrypted instruction code with one or more first symmetric keys according to a counter mode encryption / decryption process. The one or more first symmetric keys may include one of, for example, a single symmetric key used for any encrypted instruction code retrieved by the first controller device, and/or multiple symmetric keys with each of the multiple symmetric key used for respective encrypted instructions from a respective one of different memory pages stored in the non-processor memory coupled to the first controller device.

[0023] 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 that include receiving a block of information from non-processor memory at an interface between the non-processor memory and processor memory comprising two or more processor memory levels, determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code, and decrypting the block of information at the interface between the non-processor memory and the processor memory for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code. The block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

[0024] 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.

[0025] In some variations, an apparatus is provided. The apparatus includes means for receiving a block of information from non-processor memory at an interface between the non-processor memory and processor memory comprising two or more processor memory levels, means for determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code, and means for decrypting the block of information at the interface between the non-processor memory and the processor memory for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code. The block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

[0026] 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.

[0027] 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 ±20% or ±10%, ±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 ±20% or ±10%, ±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.

[0028] 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.

[0029] 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.

[0030] 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

[0031] These and other aspects will now be described in detail with reference to the following drawings.

[0032] FIG. 1 is a schematic diagram of a system to implement instruction set randomization (ISR).

[0033] FIG. 2 is a diagram of a decryption counter mode cipher implementation.

[0034] FIG. 3 is a flowchart of an example procedure for instruction set randomization.

[0035] FIG. 4 is a schematic diagram of a generic computing system.

[0036] FIG. 5 is a diagram showing a portion of an Instruction Fetch Unit (IFU) at the front-end of the execution pipeline.

[0037] FIG. 6 includes a table listing common architecture parameters used for the simulations of the implementations described herein.

[0038] FIGS. 7 and 8 are graphs showing performance results for an instruction set randomization system with different placement choices for a decryption unit.

[0039] FIG. 9 is a flow diagram illustrating example software distribution operations in a System Instruction Set Randomization implementation.

[0040] FIG. 10 is a flow diagram illustrating example software distribution operations in a Page Instruction Set Randomization implementation.

[0041] Like reference symbols in the various drawings indicate like elements.

DESCRIPTION

[0042] Described herein are systems, devices, apparatus, computer program products, and other implementations to enhance security and protection of computing systems from malware and other malicious attacks on the computing system, including a method that includes receiving a block of information from non-processor memory (e.g., DRAM) at an interface between non-processor memory and processor memory comprising two or more processor memory levels, e.g., cache with a two-level memory configuration that includes an L1 level partitioned into an instruction portion, L1-I, and a data portion, L1-D, and further includes an L2 level that is the level closest (logically) to the interface. The method also includes determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code, and decrypting (e.g., a by a decryption unit employing, for example, a symmetric decryption key) the block of information at the interface between the processor non-memory and the processor memory for storage in one of the two or more levels (e.g., in the L2 level) of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code. The block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data. As will be discussed in greater details below, in some embodiments, the method may further include associating an identifier with the received block of information, and setting the identifier to a value indicating one of instructions or data based on the determination of whether the block of information corresponds to the encrypted instruction code received from the non-processor memory. When blocks of information retrieved from the non-processor memory are tagged / identified, the method may also includes determining, in response to a fetch request made by an instruction portion of another level of the two or more levels of the processor memory (e.g., the L1-I cache) for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of

information stored at the one of the two or more levels of the processor memory corresponds to instructions. When the identifier associated with the block of information is determined to correspond to the instructions, the at least the portion of the block of information stored at the one of the two or more levels of the processor memory is transferred to the instruction portion of the other of the two or more levels of the processor memory. When the identifier associated with the block of information is determined to correspond to data, the block of information is removed from the processor memory (i.e., the blocked is flushed from the processor memory), and the block of information is then re-read from the non-processor memory in order to perform a decryption operation on the block of information that is re-read.

[0043] The systems, methods, and other implementations described herein may be used to support instruction set diversification procedures and techniques to achieve minimal performance overhead. In some implementations, a native hardware support for diversifying systems by enabling instruction-set randomization (ISR) is provided, which aims to provide a unique random ISA for every deployed system. Software implementations are often too slow (70% to 400% slowdowns) and are generally are insecure because they use weaker encryption schemes and can be turned off. In some embodiments, to achieve instruction set randomization, different machines are assigned different encryption/decryption keys that are applied to instruction code generated for a particular instruction set architecture (ISA). By encrypting instruction code for a particular ISA with different secret keys, the benefits of having a different ISA for each individual machine / computer using the particular ISA can be emulated. As will be discussed below in greater details, a cryptographic scheme, such as AES (Advanced Encryption Standard adopted by the National Standards and Technology 2001) may be used to encrypt program binaries and decrypt the binary prior to execution by the processor. These procedures / methods for generating diversity through use of multiple different keys are simpler than actually customizing the decoder for every individual chip to implement random mappings or changing the microarchitecture machines.

[0044] Attacks typically follow the path of least resistance. Therefore, by breaking systems' homogeneity and providing diverse computing platforms, attackers are forced to develop custom exploits, thus drastically reducing their return-on-investment. In the present disclosure, only binary code-injection attack vectors, achieved via drive-by-download and social engineering attacks, are considered, but the methods, systems and other

implementations described herein can be applied to other types of attack. In code-injection attacks, an attacker's goal is to lure a user into running a malicious executable, or inject code on a local application (after exploiting a software vulnerability).

[0045] In addition to the instruction sets diversification approaches disclosed herein, various contemporary defense / security mechanisms may also be employed, e.g., address space layout randomization (ASLR) and software-only diversification solutions for mitigating Return-Oriented Programming (ROP) attacks. In some embodiments, support for hardware-aided protection may be used in processors in the form of NX bits, virtualization (AMD-V, VT-x), trusted execution (Intel TXT, ARM TrustZone), etc. The solutions described herein are agnostic and orthogonal to such technologies that behave homogeneously in all deployments. However, the diversification implementations described herein can make other supplemental protection solutions used more effective by raising the bar for an attacker that studies their deployment and tries to bypass them. Instruction set diversification can thus complement other installed defense mechanisms.

[0046] In implementing the various approaches described herein, it was assumed that NX protection (e.g., Intel's eXecute Disable bit, AMD's Enhanced Virus Protection, ARM's eXecute Never bit) may not always be effective against binary code-injection attacks. It was also assumed that a trustworthy software principal (referred to as a Trust Principal), which is part of a trusted computing base (TCB), is used in order to supply and manage keys in hardware. A Trusted Principal is generally a module in the hypervisor or OS, and can supply keys to authorized processes, and manage the keys during execution (if necessary). It is also assumed that remote key transfers are done via secure channels (e.g., IPSec).

[0047] When implementing an instruction-set randomization system / method, there are three objectives that should be met: (1) for correctness, encrypted instructions should be decrypted before execution; (2) to mitigate penalties, instruction decryption should be as far away from the energy and performance critical paths in the processor; and (3) for security, the instructions should be decrypted where an adversary does not have the ability to re-write the instruction stream, and the key(s) used for encryption should be available only to authorized users. As will become apparent below, these objectives may be achieved by placing a decryption unit of an ISR system at the interface between processor memory (e.g., cache with two or more cache memory levels) and non-processor-memory (e.g., DRAM).

[0048] Thus, with reference to FIG. 1, a schematic diagram of a micro-architectural system 100 to implement instruction set randomization is shown. The system 100 comprises a processor 110 (e.g., a general-purpose processor, an application-specific controller, etc.) that includes a processing pipeline 120 for staged processing of instructions. Electrically coupled to the processing pipeline 120 is a multi-level cache 130, which in the example of FIG. 1 includes two levels, namely, L1, which is partitioned into an instruction portion L1-I cache 132 to store instruction code and a data portion L1-D cache 134 to store non-instruction data. The multi-level cache 130 depicted in FIG.1 further includes a second level L2 cache 136 which is configured to store both instruction and non-instruction data. As noted, although the example processor 110 is depicted as including a 2-level processor memory (cache), in some implementations, the processor memory (such as the processor memory 130) may include additional levels.

[0049] As further illustrated in FIG. 1, the system 100 includes a decryption unit 140 (also referred to as a decryptor or a decryption module) that is placed at an interface (boundary) 112 between non-processor memory 150 (e.g., DRAM memory, or any other type of volatile or non-volatile computing storage device that may be a local or a remote storage device) and the outer-most processor memory level (i.e., the most remote, and possibly slowest cache memory that is accessed by the processor, with that level being accessed if there is a cache miss at a lower level of the cache memory). In some embodiments, the decryption unit 140 may comprise part of the processor 110, while in some embodiments, the decryption unit 140 may be a separate unit (e.g., not integral, modular unit) from the processor 110. As also shown in FIG. 1, there are several other placements options where the decryption unit may be positioned. For example, in placement Option 1 of FIG. 1, the decryption unit may be placed at the pipeline 120 (just before the fetch and decode stage). However, such an implementation may be energy inefficient and slow because a strong cryptographic algorithm, such as AES, would require between 20-120 cycles for decryption. While pipelining the decryption unit can mitigate some of the throughput concerns, the power and energy overheads may be significant as each instruction has to repeatedly go through the decryption unit. In placement Option 2, also depicted in FIG. 1, the decryption unit would be placed between L1-I and L2. It is to be noted that generally decryption can be performed anywhere on the processor chip as long as there is a clear distinction between instructions and data in the stream. One of the first structures in the microarchitecture where this becomes apparent is when instructions are filled from the L2 cache into the L1-I cache. If the

decryption unit is placed in the instruction miss fill path, the cost of decryption can be amortized due to locality available in the L1 cache. Because most cache accesses hit in the L1-I cache, ISR generally adds only to the L1-I miss penalty. If the processor supports simultaneous multithreading this latency can be mitigated even further. Placing the decryption unit at the interface 112 (also marked as Option 3) between the processor memory and non-processor memory additionally mitigates the latency problem.

[0050] In some embodiments, the decryption unit 140 is realized as a hardware-based implementation that may be part of the processor 110. In some embodiments, the decryption unit may be a unit separate from the processor 110, and may be a hardware implementation, a software implementation, or a hybrid hardware/software implementation. In some embodiments, the decryption key may be assigned to decryption unit at the time of manufacture of the chip (e.g., in hardware-based implementations of the decryption unit 140), or may be assigned (and/or periodically changed) at some time subsequent to the manufacture or implementation of the decryption key 140.

[0051] As noted, in the methods, systems, and other implementations described herein, to achieve diversification (through instruction set randomization), program instruction should be encrypted with secret keys in some fashion and decrypted before execution. In some embodiments, an encryption unit 160 may be coupled the non-processor memory (the DRAM 150 in the example of FIG. 1) that is configured to encrypt instruction code transferred to the non-processor memory for storage, and/or may be configured to periodically identify already stored content to apply encryption operation thereto. As also discussed below in reference to FIGS. 9 and 10, in some embodiments, encryption of instructions (and/or non-instruction data) may be performed at a remote device, e.g.,. a trusted server which encrypts instruction code and/or data with a secret key associated with a requesting computing device that includes the processor 110, and transfers to the requesting computing device encrypted instruction code (and/or encrypted data). In such embodiments, the receiving requesting computing device would include a communication module (e.g., wireless transceiver, a network interface, etc.) to communicate with the remote device. As with the decryption unit 140, in some embodiments, the particular encryption key may be assigned to the encryption unit 160 at the time of manufacture (thus making it more difficult to subsequently change the key post-manufacturing) or may be assigned (and be configured to be changed) at some later point of time (thus enabling periodical change to the secret key(s) a particular machine is

using in order to increase the difficulty of compromising the machine's / computer's security). Generally, it is desirable to use cryptographically strong encryption to prevent the attacker from guessing the ISA. To implement such a randomization scheme (i.e., use of multiple encryption/decryption keys) randomization is emulated by using different encryption keys on different devices / machines, and decrypting the encoded code at the decryption unit 140.

[0052] There are several possible ways to implement an encryption / decryption mechanism for each machine / device that offer a spectrum of diversity versus complexity tradeoffs. In some embodiments, randomization can be obtained by applying encryption for different message lengths, which correspond to different granularities of encrypted code. There are several possible approaches to implement encryption:

1) One key for the entire system – A simple implementation approach is to encrypt all software (and firmware) on the system with the same key (which would also be available for decrypting instruction code). Without the key, an attacker's binary code (the malicious payload) will be decoded as a garbled sequence of instructions, and will not behave as the attacker intended. One problem of such an approach is the excessive amount of trust put on the system-wide secret key. If the key is leaked and becomes known to the attacker, the one-key implementation will become ineffective from a security point of view.

2) One key per privilege level – Another approach to implement encryption is to use one key per privilege level, with the appropriate key enabled by the hardware upon privilege level changes. Here too, a leaked key will result in software of the active privilege level being compromised, and thus a new encryption key would need to be installed. Under this approach, it is important to ensure that installation of a new key in the hardware cannot be initiated at the various privilege levels (to prevent use of a stolen key for some privilege level to achieve unauthorized resetting of any of the keys).

3) One key per process – This approach provides a strong degree of protection against key leaks. Particularly, a leaked key for one application does not put another application in danger. However, such an approach complicates software implementation in several ways. First, some trusted entity has to store and manage keys for potentially a very large number of applications. Additionally, the software

component would need to operate at the operating system (OS) level, because the process abstraction is not readily available below the OS, and as such it will be inherently vulnerable to OS kernel vulnerabilities and attacks. Another problem with this approach is that it breaks code-sharing among applications. Thus, shared code regions, like libraries, will have to be encrypted under different keys. One simple workaround would be to statically compile all binaries.

4) One key per executable package – Under this approach, each package containing executable code (including shared libraries) is encrypted with different keys. Consequently, a running process might have multiple static and dynamically loaded code modules (corresponding to the executable and the shared libraries) in its address space, each encrypted with a different key. This approach has a similar order of key complexity as the "one-key-per-process" approach, except that now sharing common memory is possible at the granularity of executable modules, allowing the use of shared libraries as usual.

5) One key per page – In this approach, each memory page, even within the same executable or library, is allowed to have different encryption keys. This is similar to and compatible with the "one-key-per-executable-package" approach. However, this approach provides more diversity than the per-executable keys, since by knowing one key or figuring out a valid opcode, an attacker is limited to code injection within the respective memory page.

6) Smaller granularities – Code encryption in the methods, systems, and implementations described herein may also be implemented with even smaller encryption granularities (encryption at levels such as functions, basic blocks, cache block width, instruction-level, etc.) This could provide enormous diversity, but with a resultant overhead complexity that may make any such encryption granularity too high to be practical.

[0053] Thus, a particular machine (e.g., a computing device such as the system 100 of FIG. 1) may be configured to employ multiple encryption keys, based on the desired diversity and complexity required, and the decryption unit (such as the unit 140) would then be configured to apply decryption based on those encryption keys and encryption approaches implemented for the encryption unit. As noted, the particular key(s) used by one particular machine may, in some embodiments, be different from the key(s) used by another machine in order to

emulate instruction set randomization, and thus enable diversity through which machines / devices are rendered less vulnerable to attacks by rogue parties. As will be discussed in greater details below, two approaches that were used during testing and experimentations performed for the methods, systems, and other implementations described herein included the one-key-per-system approach and the one-key-per-page approach.

[0054] A further implementation choice for the encryption / decryption mechanism employed by a particular machine or device is the key storage implementation, specifically, where to store, for example, decryption keys. Generally, in implementations based on symmetric AES keys (e.g., AES keys with 128 bits), the keys are stored in hardware. When using the one-key-per system approach (also referred as "System ISR") which requires a single key, the key can be stored at the decryption unit itself because the key will generally not be modified (at least during normal operation). On the other hand, when using the one-key-per-page approach (also referred to as "Page ISR"), a logical place to hold the per-page AES keys (and a corresponding initialization vector, or IV, if required) is the Instruction Translation Lookaside Buffer (ITLB). On a page fault, the key corresponding to an encrypted page is installed into the ITLB SRAM (with or without OS assistance) as part of the fault handling mechanism. Subsequently, every time an instruction miss request is sent out the page keys are also read from the ITLB and sent out as part of the miss request so that the blocks can be unencrypted at the appropriate level.

[0055] Another implementation choice for encryption / decryption mechanisms is the encryption mode to be used. Although any type of cipher may be used (e.g., symmetric, asymmetric), in the implementations that were tested and evaluated, symmetric ciphers were chosen instead of asymmetric ones for encryption, in part because asymmetric ciphers are generally more expensive than symmetric ones. As noted, an example symmetric cipher is the AES cipher, but any other cipher type may be used instead.

[0056] Symmetric block ciphers can be used in a number of modes. An important requirement for the systems, methods, and other implementations described herein is that decryption of any random block in a message should not depend on any other block. This property of random access should be used since the decryption of blocks should be independent of other blocks. If this were not so, all required blocks would have to be made available to the decryption accelerator at once by fetching them from cache or memory, which will not only increase decryption latency but also pollute the cache. This requirement

means that, in some embodiments, CBC (cipher-block chaining), CFB (cipher feedback) and OFB (output feedback) cipher modes would not be used. Additionally, to achieve for stronger security, it is preferable not to use ECB (electronic codebook) mode since encryption using this mode is prone to dictionary and block-reuse attacks.

[0057] Thus, in some embodiments, a cipher mode that may be used for the systems, methods, and other implementations described herein is counter-mode ciphering. In this mode, a keystream block is generated by encrypting successive values of a counter that provides distinct outputs (at least during a particular period of time) so that each message block uses a unique counter value (else dictionary attacks can be mounted across ciphertext messages that use the same key and counter). FIG. 2 is a diagram of a decryption counter mode cipher implementation 200, in which in the decryption operations are performed with an encryption/decryption key and an initialization vector (IV) which may be used as a counter seed for encryption/decryption of a message block (the same IV would need to be used for encryption as well as decryption, but does not necessarily need to be kept secret). As depicted in FIG. 2, the decryption implementation 200 (which may be similar to the decryption unit 140 of FIG. 1) produces output based on the decryption key (which for symmetric ciphers may be the same as, or a function of, the encryption key), and the IV and/or the counter value it receives, and combines that resultant output with the ciphertext (using a combiner 202) to generate the plaintext. In some embodiments, the combining operation may be implemented as a XOR operation. Thus, in the counter mode, decryption is not performed on the ciphertext but on the counter. This enables simultaneous block fetch and decryption, assuming: (a) the appropriate counter value is known at request initiation, and (b) XOR operation on the block takes few, if any, cycles . There is, therefore, a performance overhead only when the decryption process takes longer than the fetch from the next stage, and if so, the overhead is determined according to the relationship (*decryption latency*) − (*fetch latency*) cycles. This implies that in the methods, systems, and other implementations described herein, in which decryption is performed at the processor memory and non-processor memory interface, there may be no overhead at all if decryption latency is less than memory fetch latency (which is generally the case).

[0058] In circumstances where counter mode (requiring an initialization vector, IV) is used in conjunction with symmetric encryption / decryption processing, there are some additional implementation choices that can be made. Particularly, when Page ISR is used (where each

code page of an application is encrypted with a unique key), in some implementation, for each code block in a page the counter may be the block offset in the page. One advantage of such a scheme is that, because the counter can be inferred from the virtual/physical address, the IVs need not be stored. In embodiments in which System ISR is implemented (where there is one key for the entire system), secure counter mode requires a unique counter for every code block. One way of achieving this is by assigning every package a unique ID in the range [1, 2^{73}]. The IV for each code page in a package may be the concatenation of the package ID and the absolute address of the page in the package. The actual counter may be derived as a concatenation of the IV and the page offset of the block. If one does not wish to maintain package ID, a simple implementation choice for System ISR is to use a constant IV for the entire system. This lowers the security in that in such an implementation, for a given key, the plaintext-ciphertext mapping remains constant. One could, thus, reuse code chunks with other code chunks from the same system without needing to know the key. Even so, all programs that wish to run on a system have to come from a trusted source which encrypts the binaries on the system, so this may not be an acceptable security choice in many situations. The vulnerability of this scheme is same as ECB mode, but with the performance advantages of CTR mode encryption, and implementation simplicity of system ISR mode. In the implementation choices discussed herein, simultaneous fetch and decryption can be enabled either by storing the IV for each page in the ITLB along with the key, inferring it (in hardware) from some page characteristic, or a combination of both.

[0059] With continued reference to FIG. 1, when the decryption unit 140 is positioned at the interface 112 (as depicted in FIG. 1), an implementation difficulty that arises is that higher-level processor memory have to determine whether a block that is to be retrieved from the non-processor memory should be processed as a block of instruction code (so that it can be decrypted) or as non-instruction data (in which case the block is not, in some embodiments, decrypted). In some implementations, this problem is resolved by determining the nature of the requested information block at the micro-architectural level. The originating source (instruction or non-instruction data) of a cache miss requiring a call to non-processor memory is preserved all the way to the memory controller to identify instruction or data miss requests. When a miss reply comes in from non-processor memory (e.g., DRAM), decryption operations are selectively applied to instruction fills. Upon a determination that an instruction block retrieved from the non-processor memory was retrieved as a result of a cache miss corresponding to instructions code (e.g., the cache miss originated from a cache

portion corresponding to instruction code, like the L1-I cache), the block of information retrieved from the non-processor memory is decrypted, e.g., by the decryption unit 140 in the example system 100 of FIG. 1, and the decrypted block of information is filled into the high level cache memory. This way, the decrypted instructions remain decrypted in the lower levels (L2 and L1), where most of the accesses hit, significantly reducing the penalty of ISR.

[0060] More particularly, when a cache miss occurs after a lower level request (e.g., from the L1 level in the example of FIG. 1) is received by the higher level cache (e.g., L2), and the higher level cache has to issue a request to the non-processor memory to obtain the required block of information, the point of origin (i.e., the source) from which the request received by the higher level cache is known. For example, when the L2 cache receives a request from the L1-I cache portion, it is known that in the event that the requested block is not in the L2 cache and a request needs to be sent to the non-processor memory, the block of information to arrive from the non-processor memory will correspond to instruction code, and thus will require (in the implementations described herein) decryption by the decryption unit 140. Conversely, when a request is received at the L2 cache from the L1-D cache portion, and the requested block is not available at the L2 cache, the block of information that will be retrieved from non-processor memory as a result of a consequential cache miss at the L2 cache will correspond to data (i.e., non-instruction data). Accordingly, in some embodiments, determining whether the block of information received from the non-processor memory at the interface between non-processor and processor memory corresponds to encrypted instruction code or to non-instruction data includes determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code based on a request sent from the processor for the block of information. The block of information received from the non-processor memory may be determined to correspond to encrypted instruction code comprises when the request for the block of information is associated with a portion of the processor-memory configured to store instructions (e.g., the L1-I cache portion).

[0061] In the embodiments described herein, the decryption implemented at the boundary / interface between a multi-level processor memory configuration and non-processor memory creates a security vulnerability. Suppose a cache block was fetched from DRAM into the L2 processor memory level in response to a data (i.e., non-instruction) load request. Consequently, the retrieved block, requested as non-instruction data, is not decrypted by the

decryption unit, and thus is stored into the high-level processor memory (e.g., the L2 level in the example of FIG. 1) as non-decrypted data. Suppose now that this retrieved non-instruction (and non-decrypted) data is also requested by the L1-I portion of the cache. In this scenario, the block, which came in as data without going through a decryption process, may be fed as-is to the L1-I cache portion and eventually may make its way into the pipeline, thus completely bypassing ISR. An attack vector to exploit this vulnerability might involve first loading the shell-code, crafted in native ISA, as data, and referencing the same locations for execution soon thereafter while that retrieved pernicious block of information is still in L2. Conversely, consider the case when a block that resides in L2 is fetched as an L1-I fill, i.e., as an instruction, but is eventually requested by the L1-D. The decrypted instruction (in a form which can be easily disassembled) is now being treated as data. In such a scenario, a well-crafted attack can be based on decrypted instructions from known locations. Given the location and the decrypted instruction, a reliable dictionary of instruction mapping could be constructed that can be used to construct valid pieces of ISRized code.

[0062] Therefore, in some implementations, to prevent such attack attempts, instructions and data cache blocks are tracked in all caches levels. Tracking of information blocks placed within processor memory may be performed by adding a bit to each cache block indicating whether it is instruction or data. This bit is set by tracking the first source of the miss, i.e., the instruction or the data cache. Thereafter, only blocks marked as instruction can subsequently be fed to the L1-I cache, and vice versa for data. Cross-sharing between the split caches is thus disallowed, either directly or through higher level cache memory (e.g., the L2 cache in the example of FIG. 1). Thus, in some embodiments, the processes implemented by the systems described herein include associating an identifier with the received block of information, and setting the identifier to a value indicating one of instructions or data based on the determination of whether the block of information corresponds to encrypted instruction code received from the non-processor memory.

[0063] With the blocks of information stored in the processor memory being identified (e.g., as instruction or non-instruction data), when the L2 cache level receives a request from the L1-I for a block marked data, that block has to be flushed from the processor memory and fetched again, this time going through the decryption process. Similarly, when a block marked instruction is requested as data, it is flushed and fetched again, only now the decryption module will be bypassed during the higher-level cache (e.g., L2) fill. Thus, in

some embodiments, processes implemented by the systems and other implementations described herein include determining (e.g., by the processor 110, or, alternatively, by the decryption unit 140, or by some control circuitry to control the decryption unit 140 and/or the processor 110), in response to a fetch request made by an instruction portion of a lower level of the multi-level processor memory for at least a portion of the block of information stored at a higher level of the multi-level processor memory, whether the identifier associated with the block of information stored at the higher level of the processor memory corresponds to instructions. Based on this determination, one of two procedures is performed. Specifically, the at least the portion of the block of information stored at the higher level of the processor memory is transferred to the instruction portion of the lower level of the processor memory when the identifier associated with the block of information is determined to correspond to the instructions. However, when the identifier associated with the block of information is determined to correspond to data (but the fetch request, as noted, was made by a portion of the cache dedicated to instructions), the block of information is removed (flushed) from the processor memory, and is retrieved again from the non-processor memory in order to perform a decryption operation on it.

[0064] In a similar manner, processes implemented by the systems and other implementations described herein may include determining (e.g., by the processor 110, or, alternatively, by the decryption unit 140, or by some control circuitry to control the decryption unit 140 and/or the processor 110), in response to a fetch request made by a data portion of a lower level of the multi-levels processor memory for at least a portion of the block of information stored at a higher level of the multi-level processor memory, whether the identifier associated with the block of information stored at the higher level of the processor memory corresponds to data. Based on this determination, one of two procedures may be performed. Specifically, the at least the portion of the block of information stored at the higher level of the processor memory may be transferred to the data portion of the lower level of the multi-level processor memory when the identifier associated with the block of information indicates that the block of information corresponds to the data. On the other hand, when the identifier associated with the block of information is determined to correspond to instruction code (but the fetch request, as noted, was made by a portion of the cache dedicated to non-instruction), the block of information is removed from the processor memory, and the block of information is retrieved again from the non-processor memory in order to be transferred into the processor memory without being decrypted.

[0065] With reference now to FIG. 3, a flowchart of an example procedure 300 to implement instruction set randomization is shown. As described herein, in some embodiments, instruction set diversification is achieved through instruction set randomization (ISR) implemented through use of different encryption keys with different devices that may otherwise have the same instruction set architecture. By employing different encryption keys (any type of key may be used, including symmetrical, or asymmetrical keys), instruction set diversification can be emulated. The procedure 300 includes receiving 310 a block of information from non-processor memory (e.g., DRAM) at an interface between the non-processor memory and processor memory comprising two or more processor memory levels. e.g., a cache with 2 levels, where the lowest level (which is farthest from the interface, and closest to the processor's pipeline) may be divided into an instruction portion (e.g., L1-I) to store instructions and a data portion (L1-D) to store data (e.g., non-instruction data).

[0066] Having received the block of information, a determination is made 320 whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code. As noted, this determination may be based on the origin of the cache request that resulted in the block of information being retrieved (fetched) from the non-processor memory. For example, if the request for the block of information originated from a portion of the cache memory dedicated to instructions (e.g., from the L1-I cache portion), then the retrieved block of information responsive to the request is deemed to correspond to instruction code.

[0067] The received block of information is decrypted 330 at the interface between the processor-memory and the non-processor memory (e.g., at the interface 112 depicted in the example system 100 of FIG. 1) for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code. Thus, if it is determined (e.g., at 320 of FIG. 3) that the received block of information corresponds to encrypted instruction code, the block of information is decrypted. The block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data. In some embodiments, decrypting the block of information corresponding to the encrypted instruction code includes decrypting the encrypted instruction code with one or more first keys assigned to the present controller device (e.g., computing device) on which the procedure 300 is performed. The one or more

first keys used to decrypt instruction code on the present processor may be different from at least one key assigned to a second, different, controller device (e.g., some other computing device) that is associated with the same particular instruction set architecture as the present controller device. As noted, the use of different keys for different devices enables achieving instruction set diversification without having to use different instruction set architecture for different individual devices. In some embodiments, the decryption process implemented by the decryption unit of the present device (and, by extension, an encryption unit coupled to, and/or in communication with, the same device) is based on use of one or more symmetric keys (e.g., AES keys) according to a counter mode encryption / decryption process. The encryption / decryption processing may be performed according to System ISR (i.e., one key for the entire system), page ISR (i.e., different keys for different memory pages), or according to any other encryption granularity scheme.

[0068] As also noted, to avoid attempts to circumvent the encryption / decryption-based security mechanism described herein by causing a block of information containing malicious instruction to be identified as non-instruction data at the time of processing by the decryption unit in order to subsequently use that block of information as instruction data, in some embodiments, block of information stored in the processor memory may be tagged or identified as 'instruction' or 'data.' When a block identified as 'data' is subsequently requested and fetched by a lower level cache to be used as instruction (e.g., by the L1-I portion 132 of the system 100 shown in FIG. 1), then the requested block of information is first removed (flushed) from the processor memory, and is retrieved again from the non-processor memory (which may be a storage device located locally or remotely) in order to perform a decryption operation on the block of information retrieved again. If the block of information is identified as 'instruction' and is in fact requested by a portion of the cache corresponding to instruction storage, the block of information is transferred to the requesting portion of the cache.

[0069] Performing at least some of the operations described herein may be facilitated by a processor-based computing system. Particularly, at least some of the various devices / systems / units described herein may be implemented, at least in part, using one or more processor-based devices. With reference to FIG. 4, a schematic diagram of a generic computing system 400 is shown. The computing system 400 includes a processor-based device 410 such as a personal computer, a specialized computing device, and so forth, that

typically includes a central processor unit 412 (which may be similar to, and/or be configured to perform operations similar to those performed by, the processor 110 depicted in FIG. 1). In addition to the CPU 412, the system includes main memory, cache memory and bus interface circuits (not shown in FIG. 4). As described herein, in some embodiments, coupled to the CPU 412 may be a decryption unit configured to decrypt (e.g., according to one or more pre-assigned symmetric keys) blocks of information received from non-processor memory based on a determination of whether the received blocks of information correspond to encrypted instruction code or to non-instruction data. The processor-based device 410 may include a mass storage element 414, such as a hard drive or flash drive associated with the computer 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) or LCD (liquid crystal display) monitor, that may be placed where a user can access them.

[0070] The processor-based device 410 is configured to perform at least some of the operations / procedures described herein. The storage device 414 may thus include a computer program product that when executed on the processor-based device 410 causes the processor-based device to perform operations / procedures described herein. The processor-based device 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. 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 processor-based 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 computing system 400. The processor-based device 410 may include an operating system, e.g., Windows XP® Microsoft Corporation operating system, Ubuntu operating system, etc.

[0071] 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.

[0072] 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.

[0073] 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.

[0074] In some embodiments, any suitable computer readable media can be used for storing instructions for performing the processes / operations / procedures described herein. For example, in some embodiments computer readable media can be transitory or non-transitory. For example, 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. As another example, transitory computer readable media can include signals on networks, in wires, conductors, optical fibers, circuits, any suitable media that is fleeting and devoid of any semblance of permanence during transmission, and/or any suitable intangible media.

[0075] In order to enable development of software and to test and evaluate at least some of the features of instruction set randomization (including encryption / decryption schemes), the

following system to support ISR was realized. An OpenSPARCT1 microprocessor core was used to implement hardware support for ISR. The OpenSPARC core was supported by the Xilinx XUPV5- LX110T FPGA development board, which enabled characterization of the system on live hardware without emulation. In the tested / evaluated implementation, the decryption module was placed below the L1-I memory interface (because, for this implementation, the FPGA does not have a multi-level cache). The Sun OpenSPARCT1 microprocessor is an open source variant of the Sun UltraSPARC T1 RISC architecture with its Verilog source available for free. It is a 64-bit processor and supports chip multi-threading (4 per core). The core SPARC pipeline includes six (6) stages and is part of the instruction fetch unit (IFU), which is responsible for maintaining instruction flow through the pipeline by controlling the modules feeding it. FIG. 5 shows a portion 500 of IFU at the front-end of the execution pipeline. Four pairs of addresses are maintained for each thread, which correspond to the current PC (in the thread instruction register, TIR) and the next PC (in the next instruction register, NIR) for four threads co-located at each core. Alternatively, the next address can also be determined based on an evaluated branch or interrupt. Once the next address has been determined, it is fetched from the L1-I cache which is coupled with a standard ITLB for faster address translation. Addresses that miss in the L1-I are stored in the missed instruction list (MIL) and eventually fed into the load-store unit, which is responsible for servicing misses. The retrieved instruction blocks are collected in the instruction fetch queue (IFQ) which fills the cache.

[0076] Ideally, for a single-key ISR used in ECB mode, the key should to be stored in some non-volatile memory, from where it can be loaded onto a register on boot time. This register can then be used to supply the key to the decryption unit for rest of system up-time. In the OpenSPARC implementation, the key was hard-coded at design time and was not made programmable. The key was used to decrypt the cache lines when they were filled into the L1.

[0077] In some of the tested embodiments, the decryption core was placed between the L1-I and L2 caches where it was configured to decrypt a full cache line of data as it came in from L2 using the key previously stored in the MIL at the time of L2 request generation (this also enabled implementation of a more complex cipher such as AES for enhanced security). The placement of the encryption core allows significant latency hiding, so a high-complexity procedure generally only minimally affects overall system performance.

[0078] The encryption key data (stored in the OS) is brought into the ITLB on a page fault by the miss handler of the OS or hypervisor. Stores are used to alternate address spaces to perform key writes. The alternate space 0x54, already supported in hardware, is used to write ITLB physical address data. At all the places where the miss handler already uses the STXA (store eXtended word to alternate address space) instruction to write ITLB physical address data, similar instructions are added to store key data in the same entry.

[0079] In some of the embodiments tested, instructions were decrypted as they were brought into L1-I from L2. This allowed for significant latency hiding, as the majority of instruction requests were serviced from the plaintext copy in L1-I and did not require decryption. The ITLB was augmented with a key storage memory structure that sits next to the physical address data memory. It is, therefore, easy to access the keys because when a memory translation in the ITLB is accessed, the corresponding key is used to decrypt the instruction at the particular address. Keys are brought into this structure on an ITLB miss by the miss handling mechanism operated by the OS/hypervisor. The missed instruction list (MIL) contains pending instruction requests that missed in L1-I and have been sent further out in the memory hierarchy for fulfillment. Each MIL structure was modified to add key storage, with enough space to store one full key per missed instruction. Key data was brought in from the ITLB at the same time as physical address data and the key in the corresponding MIL entry was used to decrypt the block before sending it on to the L1-I.

[0080] Aside from the crypto accelerator that had approximately 2500 lines of Verilog code, System ISR required less than 5 lines of Verilog code of the OpenSPARC processor. Implementing Page ISR in hardware was slightly more complex. Even so, only approximately 500 lines of Verilog code (a portion of which was wire-forwarding among different modules) needed to be added/modified. It is to be noted that even minimal changes in hardware code can enable very strong protections for software systems that have billions of lines of vulnerable code, thus validating an approach of using hardware-enhanced security to protect against attacks.

[0081] The impact that the placement choice of the decryption unit (e.g., at the interface between processor and non-processor memory, between L1 and L2 cache level, or just before the processing pipe line) has on performance of the methods, systems, and other implementations described herein was tested and evaluated. As will be further discussed

below, the testing and evaluation showed that even using strong and expensive encryption standard such as AES for security can be done with negligible overhead.

[0082] To further test and evaluate the impact the decryption unit placement has on performance, a gem5 simulator (a commonly used modular platform for computer system architecture research) was modified to simulate extra hit latencies corresponding to various implementation points. FIG. 6 includes a Table 600 providing the common architecture parameters used for the simulations. The simulator itself was run in system-call emulation mode with a SimpleTiming-CPU. To simulate latency corresponding to an AES decryption, 40 cycles were added (which is the latency of the AES core implemented) to the stage at which it is be to placed. For instance, for a decryption unit placed between the L1-I and the L2 caches, if the original hit latency in L2 for requests from L1-I is 100 cycles, the new hit latency became 140 cycles. The miss latency is increased by an equal amount as well. For all experiments conducted, the benchmarks from the SPEC CPU2006 benchmarks suite were used. These benchmarks were run to completion with the test input set. These benchmarks were chosen because they were the only ones in the SPEC CPU2006 suite that ran without functional-emulation errors on the ALPHA/gem5 emulator in syscall mode.

[0083] FIG. 8 shows the performance overhead for decryption unit placement at the L1-I and L2 boundary / interface. For the sake of completeness, the scheme where the decryption module is placed at the head of the execution pipeline, i.e., when every instruction fetched from the L1-I is decrypted, was also examined (results shown in FIG. 7). Performance results for the case where decryption occurs between the L2 and memory (i.e., processor memory and non-processor memory interface) were not obtained because decryption latency is much lower than memory fetch latency, resulting in no performance loss. As can be seen from the performance results shown in graphs 700 and 800 of FIGS. 7 and 8, placing the decryption unit before the pipeline is generally prohibitively expensive (from a computational effort perspective), whereas the other choice incurs nominal overhead. From a high level perspective, execution overhead will be correlated to how often instruction decryption occurs. Placing the decryption module at the cache interfaces results in minimal overhead because of locality at the caches. This seems to be especially true for benchmarks like *gcc*, *gobmk* and *sjeng* which suffer from a lot of L1-I misses. In such cases, it is advantageous to implement the third choice for placement of the decryption unit, namely, at boundary / interface between the processor-memory (cache) and the non-processor memory, where the performance

overhead is low regardless of the cache miss profile (but at the possible cost of more complex processing).

[0084] It is also to be noted that physical-area requirements for implementations of the methods, systems, and other implementations described herein generally do not require significant area overhead. Particularly, when the implementations described herein were synthesized for a 32nm technology node, it was determined that adding an extra 16-entry 128-bit RAM (and the associated circuitry) to the ITLB for storing encryption / decryption keys resulted in a 81% increase in the size of the ITLB, e.g., from 0.295 mm^2 to 0.536 mm^2. This should be manageable, however, because the actual core is orders of magnitude bigger than the ITLB. Additionally, a fully-pipelined AES decryption unit was implemented that took 40 cycles to produce a 128-bit plaintext. The synthesized implementation took up 0.407 mm^2 of die space for the same technology.

[0085] A security analysis of the security properties and weaknesses of the ISR processes described herein was conducted for various types of attack, as more particularly discussed below. One type of attack for which the methods, systems, and other implementations described herein were analyzed is "Code Injection" attack. Code-injection attacks are among the most widely known type of software abuse. In particular, these attacks rely on depositing arbitrary executable code inside the address space of a victim process (typically in the form of data), and then transferring control into that code, e.g., by overwriting a function pointer, a dispatch table, or a function return address. Instruction-set randomization generally defeats binary code-injection attacks (when using a cryptographic process not susceptible to known plaintext attacks). The underlying principle behind this ability to defeat code injection attacks is that because an execution engine expects opcodes tailored to a particular ISA, it will not understand those of any other arbitrary ISA and thus fail to execute payloads comprising of such instructions. Even if a valid ISA for a target was to be determined by exploiting some (local) flaw or information leak, the diversity of ISAs in the ecosystem ensures that other systems will not fall prey.

[0086] Another type of attack for which the methods, systems, and other implementations described herein were analyzed is the "Code-Reuse" attack. Code-reuse techniques are the attackers' response to the increased adoption of system hardening mechanisms, like ASLR and NX, from commodity systems. The main idea behind code-reuse is to construct the malicious payload by reusing instructions already present in the address space of a vulnerable

process. For example, if the attacker manages to diverge the control flow to the beginning of a library function such as libc's *system*, then the attacker can spawn a shell without the need to inject any code. In most cases, however, invoking a single function is not enough for a successful compromise – multiple return-to-libc calls need to be "chained" together by reusing short instruction sequences. Return-Oriented Programming (ROP) takes this idea to the extreme by using only a carefully selected set of short instruction sequences, known as gadgets, for performing arbitrary computations. Note that ROP has been shown to be Turing complete, thus eliminating the need for calling libc functions. This technique gives the attacker the same level of flexibility offered by arbitrary code injection without injecting any new code at all. The malicious payload comprises just a sequence of gadget addresses intermixed with any necessary data arguments. However, in most publicly available exploits so far, attackers do not rely on a fully ROP-based payload. Typically, ROP code is used only as a first step for bypassing NX: it allocates a memory area with write and execute permissions, by calling a library function like VirtualAlloc or mprotect, copies into it some plain shellcode, and finally transfers control to the copied shellcode that now has execute permissions. ISR cannot protect against pure code-reuse attacks, since the attacker can construct his payload without knowing the underlying ISA. However, there are no pure practical code-reuse attacks, and the methods, systems, and other implementations described herein thwart multi-stage code-reuse exploits that rely on code injection. Additionally, the methods, systems, and other implementations described herein are orthogonal to many techniques for mitigating pure code-reuse attack.

[0087] Another type of attack is one based on BIOS/Boot vulnerabilities. Before an OS enables privilege-protection (often with assistance from hardware), a typical system boots into a BIOS and eventually to a bootloader. These two modules are promising targets for attackers, since they offer virtually unrestricted access to system resources with few (if any) monitors on their activity. Employing ISR at these stages using predetermined secret keys ensures that they cannot be patched or replaced with unverified code.

[0088] A further attack type is one that employs self-modifying code (SMS). ISR complicates applications that employ self-modifying code (SMC), such as just-in-time (JIT) compilers. To counter these types of attacks, such applications can be either completely disallowed, or can be selectively permitted to run with a significant reduction on the security guarantees offered. Allowing SMC requires different approaches depending on the mode of

ISR under consideration. For System ISR, if the SMC agent is considered part of the TCB, it could be allowed access to the system key. This significantly weakens system security by increasing the attack surface substantially and, therefore, is not very appealing. In case of Page ISR, the SMC (or JIT) engine could either execute the generated code pages with null-keys, or encrypt it with a set of self-generated keys. Either way, some trust has to be offloaded to the SMC agent.

[0089] Although the Instruction Set Randomization (ISR) methods, systems, and other implementations described herein provide diversity and thus raise the bar for wide-scale exploitation, ISR is susceptible to some forms of attack. ISR's weaknesses include:

a) Trusted OS – Page ISR requires a trusted system (e.g., verified OS module) component to perform the key management for each page (per process). Unfortunately, commodity OSes are vulnerable to compromises, and are nowadays increasingly attractive targets.

b) Data-Only Attacks – Since ISR does not guarantee data integrity, it is susceptible to data-only attacks (attacks that do not touch control data or modify code execution). Moreover, for applications that leak information through processed data, program state may be gleaned by observing its data.

c) Denial-of-Service (DoS) Attacks – DoS attacks may be carried out by repeatedly patching a program with illegal code, and crashing it with "illegal instruction" exceptions. However, under most circumstances this is considered acceptable behavior.

In spite of the above indicated weaknesses / drawbacks, additional security and integrity measures can be adopted on top of ISR to overcome these (or future) issues and make it more robust.

[0090] The challenges for implementing ISR-based security extend beyond merely diversification at the hardware level, and require careful implementation of software distribution schemes and systems. Two different models for software distribution were considered. One was a model for System ISR with a constant IV, and the other was Page ISR with unique key per page. These represent two different points on the implementation spectrum for software distribution.

[0091] Starting first with a System ISR implementation, application distribution for this case can be explained in the context of the app-store model (note that if regular counter-mode was to be used, system ISR can also be implemented in the non app-store model). When a user downloads an app from the app-store, the app-store encrypts the binary with the ISR key for that device. FIG. 9 is a flow diagram 900 illustrating example software distribution operations, which include:

(1) The chip-key as well as the unencrypted application binaries are collected and archived in trusted key and app databases respectively.

(2) When a user chooses to install an application, the user sends a request (at operation 910) to the gateway server.

(3) The gateway then sends fetch requests for the corresponding binary (at operation 920) as well as the user's key (at operation 930), which is retrieved (at operations 940 and 950). Retrieval of the user's key is based on a unique ID that may have been sent as part of the user's request.

(4) The binary is encrypted with the fetched key (at operation 960) and sent to the user's machine (at operation 970).

[0092] Key Management for System ISR is performed as follows. Each processor chip is programmed with a unique random AES key. One simple way for producing chips with different keys is to include some nonvolatile memory on a die and then program it with a serial number and secret symmetric keys after manufacturing, perhaps using a chip tester. In the example methods, systems, and other implementations described herein, the chip manufacturer (or its proxy) is trusted to safely hold the keys and provide encryption service to authorized serial numbers. In the event the key is inadvertently leaked or lost, the manufacturer may remotely re-program the non-volatile memory using, for example, secret microcode instructions (which are generally not ISRized). To Bootstrap with a System ISR implementation, the System ISR implementation can be turned on right from boot-up. In such a case, even the BIOS, firmware and bootloader may be encrypted. No other major code or design changes need be done for these modules, but rather a simple binary rewriting pass with the AES key is sufficient to ISRize the system. It is to be noted that diversification does not hinder or change existing secure boot schemes where the bootloader and/or OS are verified for integrity (e.g., via comparison with some known or expected measurement of

hashes on the binary image). Importantly, however, because of ISR, the measurement to be verified will not be the same across systems.

[0093] Two risks associated with a System ISR implementation include the fact that use of a single key increases risk of misuse / loss, and the fact that applications are required to go through the app store, which might give rise to a host of undesirable problems (absence of single trusted distribution authority, less flexibility, privacy, etc.). Page ISR implementation resolves these problems with a flexible key utilization system allowing use of arbitrary keys.

[0094] When a developer wants to distribute an application under a Page ISR implementation, the developer would be using the page-level granularity of encryption, i.e., encrypt different code pages with different keys (however, in some embodiments, one key could correspond to more than one page). FIG. 10 is a flow diagram 1000 showing a sequence of operations to distribute software in a Page ISR implementation. The sequence of operations includes the following operations:

(1) Users send (at operation 1010) download request to the app server. The public key is embedded in the request.

(2) For every download request, an app server receives for his application, the app server encrypts (at operation 1020) the binary with a symmetric encryption process such as AES using developer-generated keys. These key-to-address mappings are then encrypted using the public key from user request, packaged with the binary (as an ELF section, for instance) and forwarded (also illustrated at operation 1020) to the user.

Because the binary code is encrypted with developer-generated keys, which are in turn encrypted with the chip's public key, code confidentiality is guaranteed because the chip's private key is never revealed (generally not even to its owner). Moreover, even if the symmetric keys for an application were to be leaked, they would not compromise the rest of the system. Alternatively, if code confidentiality is not necessary, applications can be downloaded in plaintext and encrypted locally.

[0095] To Bootstrap with a Page ISR implementation, the BIOS initializes the bootloader, which then initializes the OS, which enables paging. To enable ISR during this process, one key for the entire BIOS program may be used with the physical load addresses acting as the counter values. The BIOS and bootloader can be prepared the same way as single key ISR

applications are prepared for execution. When the OS gains control, it initializes the page tables of its own code pages with keys that were used to encrypt it. From this point onwards (i.e., after virtual memory has been set up), the system starts operating in Page ISR.

[0096] When a program is executed, the loader extracts the encrypted key-to-page mapping from the binary and passes this information to the OS. The OS then either decrypts this mapping using the system's private key itself or asks the hardware to do so in case trusting the OS with the private key is not desirable. Once it has the decrypted mappings, it is responsible for installing the keys corresponding to each code page, both in the hardware and software address-translation mechanisms. One way to facilitate hardware translation is to extend existing page table data fields with an additional field for the 128-bit key. The OS can install the appropriate keys to each page-table entry on a page fault, which can then be transferred to the ITLB by the hardware on every ITLB miss. The decryption process uses lower order page offset bits as the counter. This can be automatically inferred at runtime and the need to store IVs does not arise.

[0097] In conclusion, as described herein, once a security flaw is discovered, attacks exploiting it can be quickly and widely deployed due to the homogeneity of computing ecosystem. Diversity is a desirable feature to limit the damage such security lapses cause. In the methods, systems, and other implementations described herein, a solution for randomizing instruction sets that is secure and has minimal performance overheads was described. In the approaches discussed herein, instruction streams are encrypted with AES and are decrypted before execution (e.g., at a decryption unit placed at an interface / boundary between processor memory and non-processor memory). The encryption key(s) (e.g., based on AES) is kept secret to prevent development of unauthorized binaries.

[0098] The implementations described herein include micro-architectural optimizations to implement the above-described approaches with minimal overhead. One feature of the ISR solution described herein is that simple modifications in hardware enable broad security coverage. For example, the implementations described herein can protect against various binary code-injection attacks.

[0099] Also explored herein was the issue of how entire ecosystems could be developed around the different types of proposed ISR. An advantage of these diversification approaches based on ISR is that they employ practical cryptographic primitives and can be built on top of

established models. Moreover, they are orthogonal to, and can be used alongside, most other code-signing and security measures to further harden systems.

[00100] ISR provides a foundation for further security improvements. One such direction is offloading key-management to hardware so that the requirement of a trusted OS can be done away with, thus evicting thousands of lines of software code from the sphere of trust. Yet another direction is to extend encryption to data items (e.g., non-instruction data), which can open the way for truly obfuscated and practical "dark execution," which is particularly significant when executing programs on remote clouds.

[00101] 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:

receiving a block of information from non-processor memory at an interface between the non-processor memory and processor memory comprising two or more processor memory levels;

determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code; and

decrypting the block of information at the interface between the non-processor memory and the processor memory for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code, wherein the block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

2. The method of claim 1, wherein the processor memory comprises cache memory of a central processing unit (CPU) organized in two or more cache memory levels, and wherein the non-processor memory comprises random access memory.

3. The method of claim 1, wherein determining whether the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code comprises:

determining whether the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code based on a request sent from the processor for the block of information.

4. The method of claim 3, wherein determining whether the block of information received from the non-processor memory corresponds to the encrypted instruction code comprises:

determining that the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code when the request for the block of information is associated with a portion of the processor-memory configured to store instructions.

5. The method of claim 3, further comprising:

associating an identifier with the received block of information; and

setting the identifier to a value indicating one of instructions or data based on the determination of whether the block of information corresponds to the encrypted instruction code received from the non-processor memory.

6. The method of claim 5, further comprising:

determining, in response to a fetch request made by an instruction portion of another level of the two or more levels of the processor memory for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of information stored at the one of the two or more levels of the two or more levels of the processor memory corresponds to the instructions; and

performing one of:

transferring the at least the portion of the block of information stored at the one of the two or more levels of the processor memory to the instruction portion of the other of the two or more levels of the processor memory when the identifier associated with the block of information is determined to correspond to the instructions; or

removing the block of information from the processor memory when the identifier associated with the block of information is determined to correspond to the data, and causing the block of information to be retrieved again from the non-processor memory in order to perform a decryption operation on the block of information retrieved again.

7. The method of claim 5, further comprising:

determining, in response to a fetch request made by a data portion of another level of the two or more levels of the processor memory for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of information stored at the one of the two or more levels of the processor memory corresponds to the data; and

performing one of:

transferring the at least the portion of the block of information stored at the one of the two or more levels of the processor memory to the data portion of the other of the two or more levels of the processor memory when the identifier associated with the block of information indicates that the block of information corresponds to the data; or

removing the block of information from the processor memory when the identifier associated with the block of information indicates that the block of information correspond to instructions, and causing the block of information to be retrieved again from the non-processor memory in order to be transferred into the processor memory without being decrypted.

8. The method of claim 1, wherein decrypting the block of information corresponding to the encrypted instruction code at the interface between the processor-memory and the non-processor memory comprises:

decrypting the block of information corresponding to the encrypted instruction code with one or more first keys assigned to a first controller device associated with a particular instruction set, the one or more first keys being different from at least one key assigned to a second controller device associated with the particular instruction set so as to emulate instruction set diversification between the first controller device and the second controller device.

9. The method of claim 8, wherein decrypting the block of information corresponding to the encrypted instruction code with the one or more first keys assigned to the first controller device comprises:

decrypting the encrypted instruction code with one or more first symmetric keys according to a counter mode encryption / decryption process, wherein the one or more first symmetric keys comprises one of: a single symmetric key used for any encrypted instruction code retrieved by the first controller device, or multiple symmetric keys with each of the multiple symmetric key used for respective encrypted instructions from a respective one of different memory pages stored in the non-processor memory coupled to the first controller device.

10. A system comprising:

a processor comprising processor memory with two or more processor memory levels; and

a decryption unit at an interface between non-processor memory and the processor memory, the decryption unit configured to:

receive a block of information from the non-processor memory;

determine whether the block of information received from the non-processor memory corresponds to encrypted instruction code; and

decrypt the block of information for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code, wherein the block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

11. The system of claim 10, wherein the processor memory comprises cache memory of a central processing unit (CPU) organized in two or more cache memory levels, and wherein the non-processor memory comprises random access memory.

12. The system of claim 10, wherein the decryption unit configured to determine whether the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code is configured to:

determine whether the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code based on a request sent from the processor for the block of information.

13. The system of claim 12, wherein the decryption unit configured to determine whether the block of information received from the non-processor memory corresponds to the encrypted instruction code is configured to:

determine that the block of information received from the non-processor memory at the interface corresponds to the encrypted instruction code when the request for the block of information is associated with a portion of the processor-memory configured to store instructions.

14. The system of claim 12, wherein the decryption unit is further configured to:

associate an identifier with the received block of information; and

set the identifier to a value indicating one of instructions or data based on the determination of whether the block of information corresponds to the encrypted instruction code received from the non-processor memory.

15. The system of claim 14, wherein the processor is configured to:

determine, in response to a fetch request made by an instruction portion of another level of the two or more levels of the processor memory for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of information stored at the one of the two or more levels of the two or more levels of the processor memory corresponds to the instructions; and

perform one of:

transferring the at least the portion of the block of information stored at the one of the two or more levels of the processor memory to the instruction portion of the other of the two or more levels of the processor memory when the identifier associated with the block of information is determined to correspond to the instructions; or

removing the block of information from the processor memory when the identifier associated with the block of information is determined to correspond to the data, and causing the block of information to be retrieved again from the non-processor memory in order to perform a decryption operation on the block of information retrieved again.

16. The system of claim 14, wherein the processor is configured to:

determine, in response to a fetch request made by a data portion of another level of the two or more levels of the processor memory for at least a portion of the block of information stored at the one of the two or more levels of the processor memory, whether the identifier associated with the block of information stored at the one of the two or more levels of the processor memory corresponds to the data; and

perform one of:

transferring the at least the portion of the block of information stored at the one of the two or more levels of the processor memory to the data portion of the other of the two or more levels of the processor memory when the identifier associated with the block of information indicates that the block of information corresponds to the data; or

removing the block of information from the processor memory when the identifier associated with the block of information indicates that the block of information correspond to instructions, and causing the block of information to be retrieved again from the non-processor memory in order to be transferred into the processor memory without being decrypted.

17. The system of claim 10, wherein the decryption unit configured to decrypt the block of information corresponding to the encrypted instruction code at the interface between the processor-memory and the non-processor memory is configured to:

decrypt the block of information corresponding to the encrypted instruction code with one or more first keys assigned to a first controller device associated with a particular instruction set, the one or more first keys being different from at least one key assigned to a second controller device associated with the particular instruction set so as to emulate instruction set diversification between the first controller device and the second controller device.

18. The system of claim 17, wherein the decryption unit configured to decrypt the block of information corresponding to the encrypted instruction code with the one or more first keys assigned to the first controller device is configured to:

decrypt the encrypted instruction code with one or more first symmetric keys according to a counter mode encryption / decryption process, wherein the one or more first symmetric keys comprises one of: a single symmetric key used for any encrypted instruction code retrieved by the first controller device, or multiple symmetric keys with each of the multiple symmetric key used for respective encrypted instructions from a respective one of different memory pages stored in the non-processor memory coupled to the first controller device.

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

receiving a block of information from non-processor memory at an interface between the non-processor memory and processor memory comprising two or more processor memory levels;

determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code; and

decrypting the block of information at the interface between the non-processor memory and the processor memory for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code, wherein the block of information is stored at

the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

20. The computer readable media of claim 19, wherein decrypting the block of information corresponding to the encrypted instruction code at the interface between the processor-memory and the non-processor memory comprises:

decrypting the block of information corresponding to the encrypted instruction code with one or more first keys assigned to a first controller device associated with a particular instruction set, the one or more first keys being different from at least one key assigned to a second controller device associated with the particular instruction set so as to emulate instruction set diversification between the first controller device and the second controller device.

21. An apparatus comprising:

means for receiving a block of information from non-processor memory at an interface between the non-processor memory and processor memory comprising two or more processor memory levels;

means for determining whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code; and

means for decrypting the block of information at the interface between the non-processor memory and the processor memory for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code, wherein the block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.

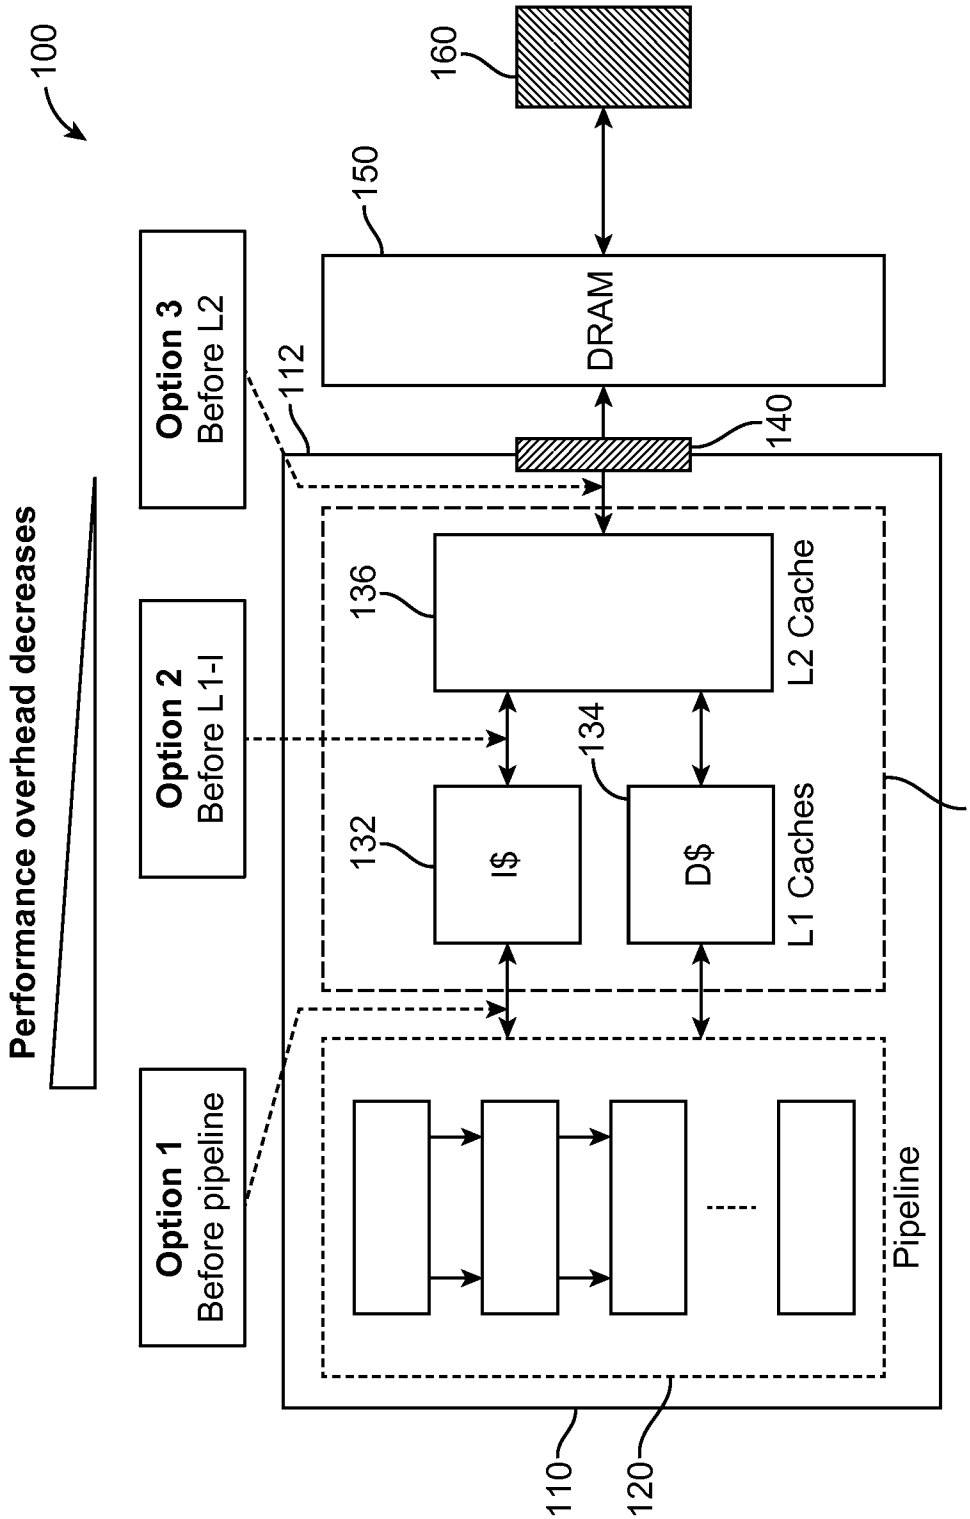
100
160
150
DRAM
112
140
Option 3
Before L2
Option 2
Before L1-I
Option 1
Before pipeline
Performance overhead decreases
136
L2 Cache
134
132
I$
D$
L1 Caches
130
Pipeline
110
120

FIG. 1

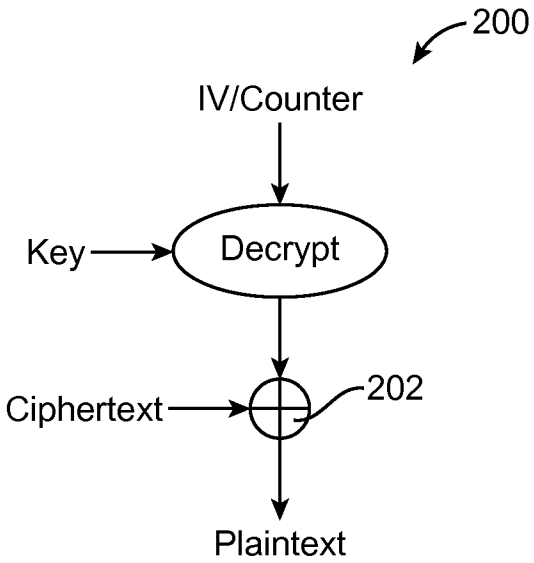

FIG. 2

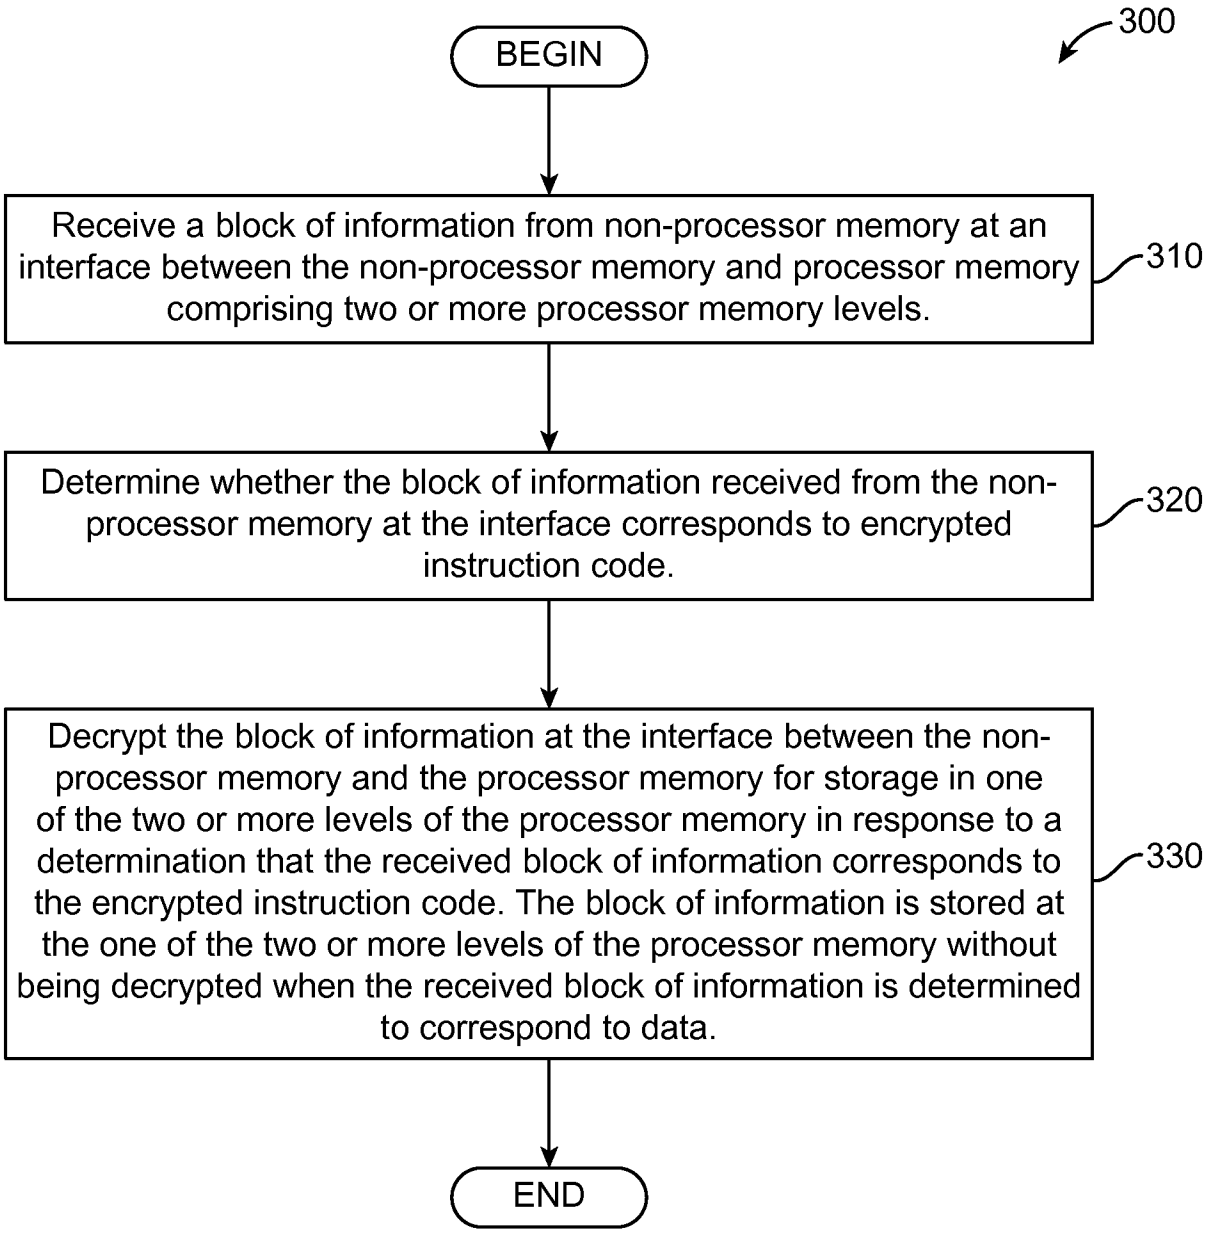
300
BEGIN
Receive a block of information from non-processor memory at an interface between the non-processor memory and processor memory comprising two or more processor memory levels.
310
Determine whether the block of information received from the non-processor memory at the interface corresponds to encrypted instruction code.
320
Decrypt the block of information at the interface between the non-processor memory and the processor memory for storage in one of the two or more levels of the processor memory in response to a determination that the received block of information corresponds to the encrypted instruction code. The block of information is stored at the one of the two or more levels of the processor memory without being decrypted when the received block of information is determined to correspond to data.
330
END

FIG. 3

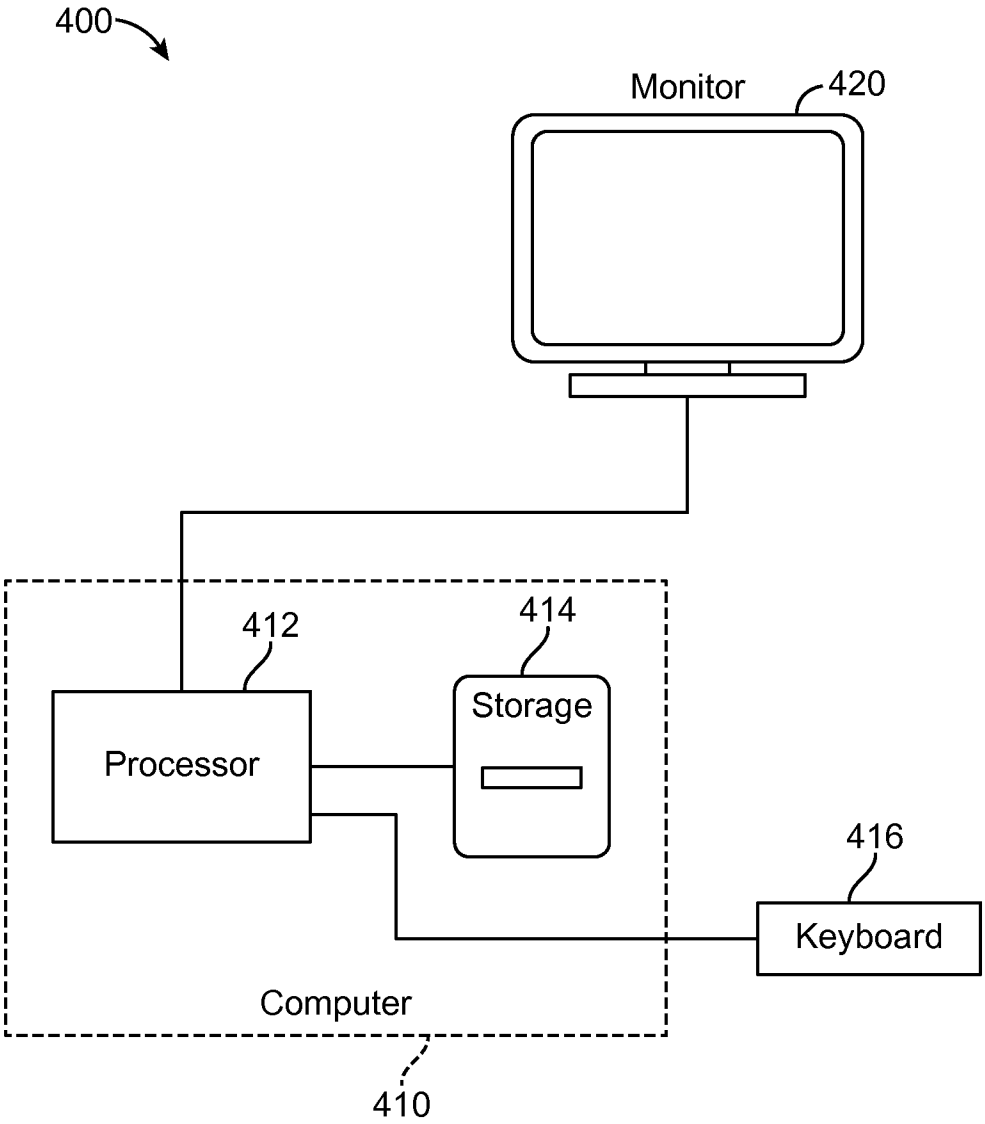
400
Monitor
420
412
414
Storage
Processor
416
Keyboard
Computer
410

FIG.4

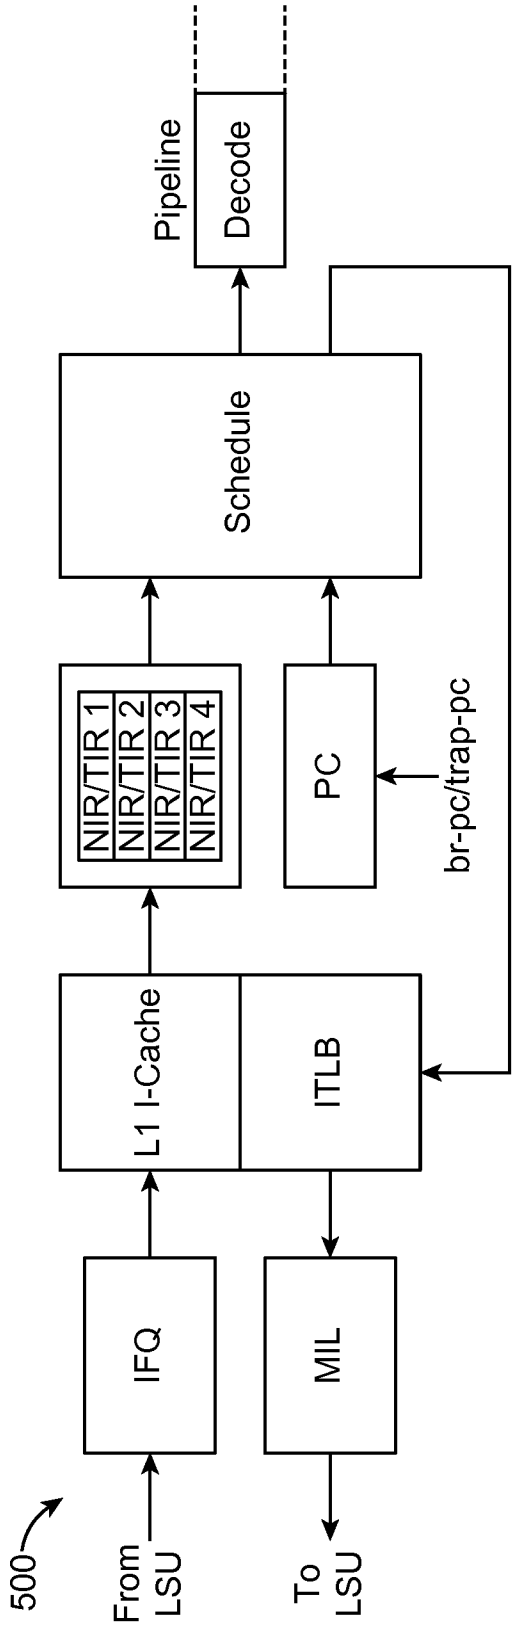

FIG.5

600

Parameters	Specifications
Number of cores	1
L1 I-cache	32kB, 2-way, 64B line
L1 D-cache	64kB, 2-way, 64B line
L1 access latency	2 cycles
L2 cache	2MB, 8-way, 64B line
L2 access latency	20 cycles
AES latency	40 cycles
Memory access latency	60 cycles

FIG.6

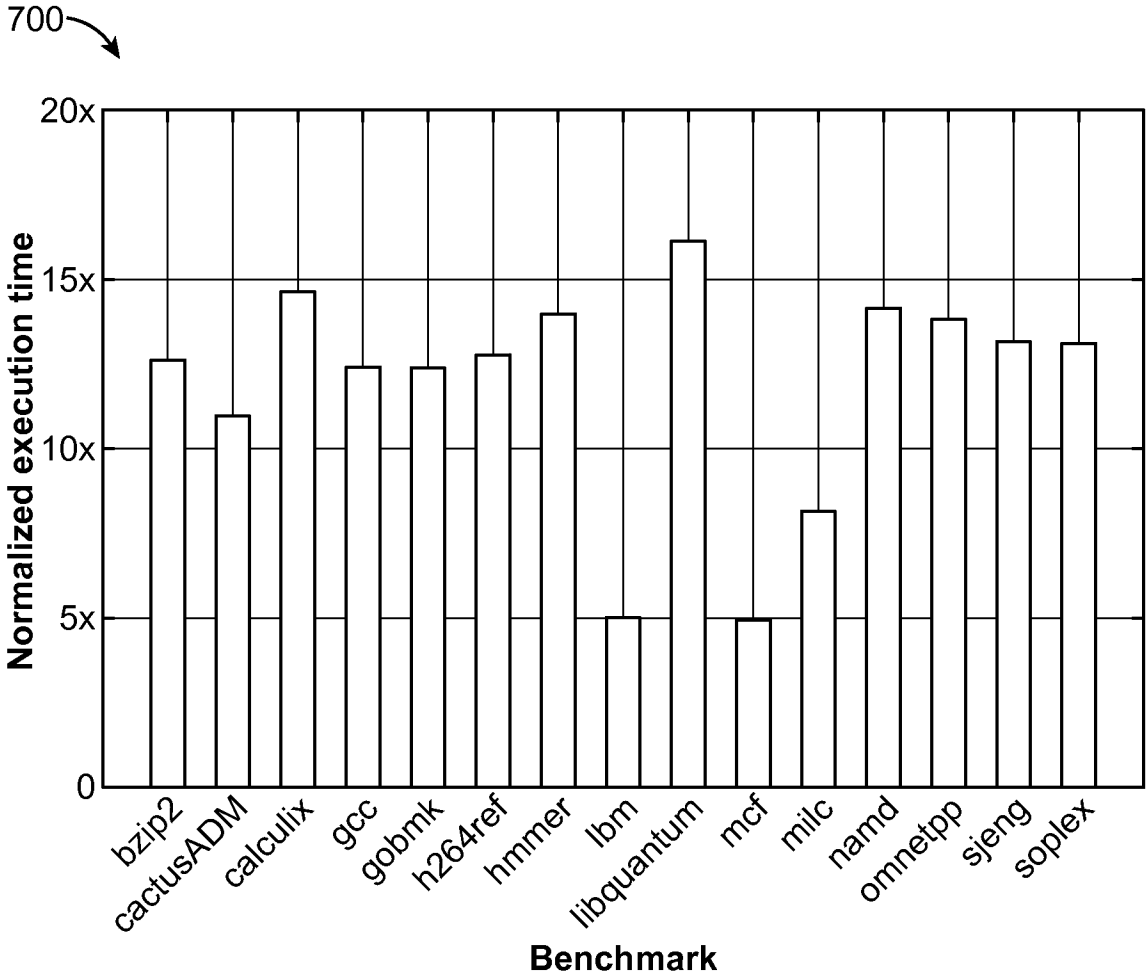
700
20x
15x
10x
5x
0
Normalized execution time
bzip2
cactusADM
calculix
gcc
gobmk
h264ref
hmmer
lbm
libquantum
mcf
milc
namd
omnetpp
sjeng
soplex
Benchmark

FIG.7

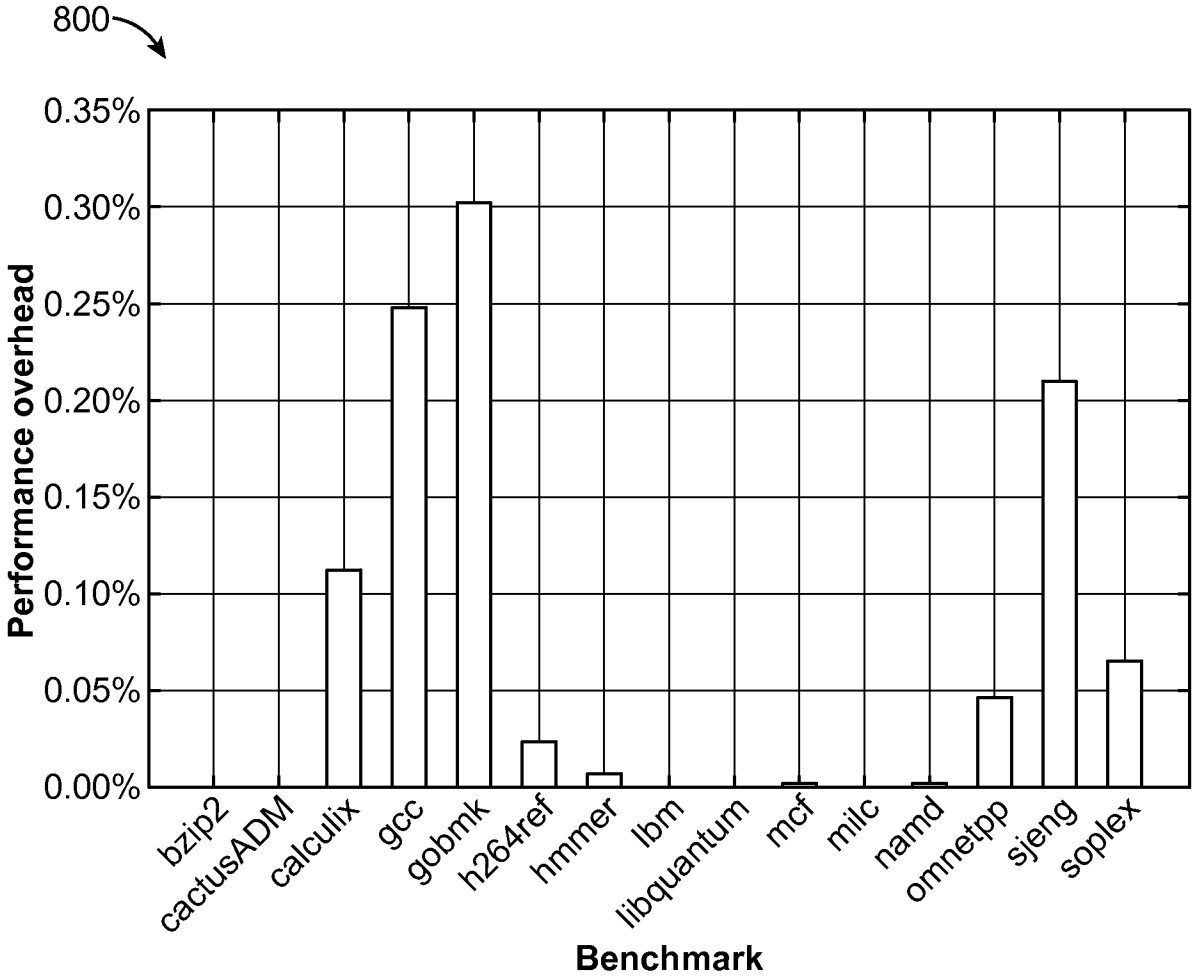
800
0.35%
0.30%
0.25%
0.20%
0.15%
0.10%
0.05%
0.00%
Performance overhead
bzip2
cactusADM
calculix
gcc
gobmk
h264ref
hmmer
lbm
libquantum
mcf
milc
namd
omnetpp
sjeng
soplex
Benchmark

FIG.8

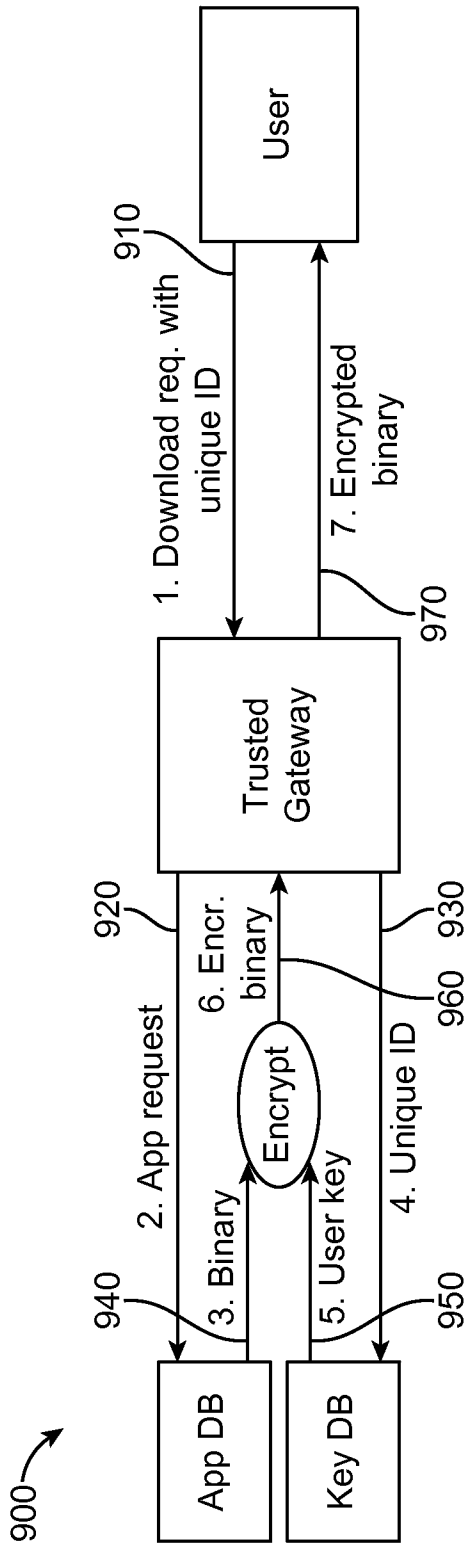
900
User
910
1. Download req. with unique ID
7. Encrypted binary
970
Trusted Gateway
920
930
2. App request
6. Encr. binary
960
Encrypt
3. Binary
5. User key
4. Unique ID
940
950
App DB
Key DB

FIG. 9

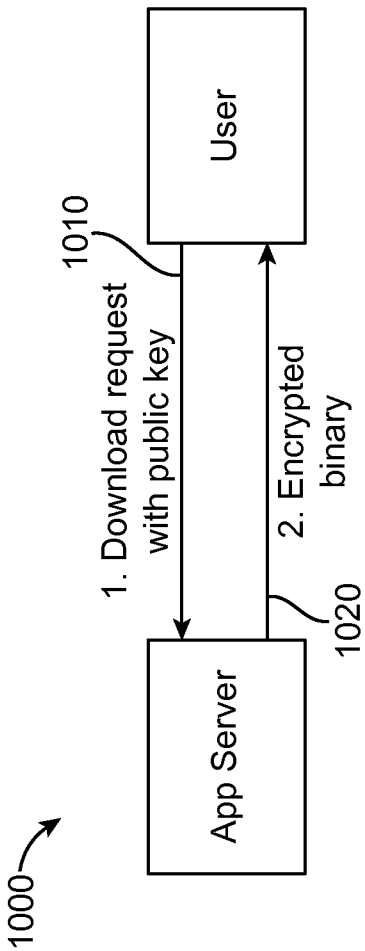
1000
User
1010
1. Download request with public key
2. Encrypted binary
1020
App Server

FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US14/44284

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G06F 21/71 (2014.01)
CPC - G06F 9/30178
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) Classification(s): G06F 21/72, 21/72; H04L 9/14 (2014.01)
CPC Classification(s): G06F 9/30178, 9/30189; H04L 9/0637

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
MicroPatent (US-G, US-A, EP-A, EP-B, WO, JP-bib, DE-C,B, DE-A, DE-T, DE-U, GB-A, FR-A); ProQuest; IEEE/IEEExplore; Google/Google Scholar, Espacenet; Keywords: encryption, decryption, binary, instruction, machine code, cache, L1, L2, symmetric key, counter mode, trust

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 2012/0036351 A1, (PEET, JR., et al.), 09 February 2012; figures 1, 4; paragraphs [0020], [0037], [0040].	1-5, 10-14, 19, 21 --- 6-9, 15-18, 20
Y	US 2012/0096282 A1, (HENRY, G. et al.), 19 April 2012; figure 1, paragraphs [0144], [0157].	6-7, 15-16
Y	US 2010/0299480 A1, (HOBSON, L. et al.), 25 November 2010; paragraph [0022].	7, 16
Y	US 6,920,562 B1, (KERR, D. et al.), 19 July 2005; figures 6, 9; column 10, lines 47-67.	8-9, 17-18, 20
Y	US 2010/0122095 A1, (JONES, M. et al.), 13 May 2010; abstract; figure 2, paragraphs [0032], [0061].	9, 18

☐ 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	Date of mailing of the international search report
29 September 2014 (29.09.2014)	24 OCT 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: Shane Thomas PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/210 (second sheet) (July 2009)