

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

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
24 February 2011 (24.02.2011)

(10) International Publication Number
WO 2011/022040 A1

(51) International Patent Classification:
G06F 17/30 (2006.01) *G06F 9/455* (2006.01)
G06F 21/00 (2006.01)

(74) Agent: **WAGNER, Robyn**; Van Pelt, Yi & James LLP,
10050 N. Foothill Blvd., Suite 200, Cupertino, CA 95014
(US).

(21) International Application Number:
PCT/US2010/002214

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

(22) International Filing Date:
11 August 2010 (11.08.2010)

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/234,604 17 August 2009 (17.08.2009) US
12/580,891 16 October 2009 (16.10.2009) US
61/257,043 2 November 2009 (02.11.2009) US
61/286,369 14 December 2009 (14.12.2009) US
12/714,547 1 March 2010 (01.03.2010) US

(71) Applicant (for all designated States except US):
FATSKUNK, INC. [US/US]; 590 Mariposa Ave.,
Mountain View, CA 94041 (US).

Published:

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

(72) Inventors: **JAKOBSSON, Bjorn Markus**; 590 Mariposa Ave., Mountain View, CA 94041 (US). **JOHANSSON, Karl-Anders, R.**; 590 Mariposa Ave., Mountain View, CA 94041 (US).

(54) Title: AUDITING A DEVICE




FIG. 1

(57) **Abstract:** The auditing of a device that includes a physical memory is disclosed. One or more hardware parameters that correspond to a hardware configuration is received. Initialization information is also received. The physical memory is selectively read and at least one result is determined. The result is provided to a verifier.

AUDITING A DEVICE

CROSS REFERENCE TO OTHER APPLICATIONS

[0001] This application claims priority to U.S. Provisional Patent Application No. 61/234,604 entitled DETECTION OF MALWARE filed August 17, 2009, U.S. Provisional Patent Application No. 61/257,043 entitled AUDITING A DEVICE filed November 2, 2009, and U.S. Provisional Patent Application No. 61/286,369 entitled AUDITING A DEVICE filed December 14, 2009, each of which is incorporated herein, in its entirety, by reference for all purposes.

BACKGROUND OF THE INVENTION

[0002] Existing techniques for detecting the presence of unauthorized programs are typically resource-intensive. For example, they generally require constant updates (e.g., of blacklists) and periodic or continuous scans for problems. The situation is exacerbated if the device being protected by such techniques has limited resources, such as limited memory, or by being powered by a battery. As one example, a device with limited resources may not be able to store definitions for detecting all known unauthorized programs. As another example, scanning for unauthorized programs is typically a power-intensive act, and may quickly deplete the battery of a battery-powered device. In some environments, a central authority is used to facilitate the discovery of unauthorized programs. One drawback of this approach is that it typically requires that the device being protected compile detailed logs of device activities. Generating such logs is resource-intensive (e.g., requiring large amounts of disk storage; processing power to assemble the log data; and the bandwidth to deliver the log data to the central authority) and can also present privacy problems.

[0003] Existing techniques for detecting the presence of unauthorized programs are also generally vulnerable to attempts by such programs to cause incorrect reporting. For example, a rootkit can “listen in” to requests by applications to the operating system, and may modify these requests and their responses. If an application requests information about what processes are running, a malicious rootkit application can avoid detection by removing information about itself from the report that is returned by the operating system.

[0004] Existing techniques for screening against the installation or execution of unauthorized programs are also known to be vulnerable to new instances of malware that may not immediately be detectable due to a lack of information about their structure and functionality. Therefore, and irrespective of the resources available to the device, if the unauthorized program is sufficiently sophisticated and/or has not previously been encountered, it can evade detection and cause undetected harm. And, if the unauthorized program has intentionally been installed by the user to bypass detection (e.g., to facilitate software piracy), traditional techniques may fail to locate the unauthorized program, or any other unauthorized activities.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] Various embodiments of the invention are disclosed in the following detailed description and the accompanying drawings.

[0006] Figure 1 illustrates an embodiment of an environment in which device auditing is provided.

[0007] Figure 2 illustrates an embodiment of a device.

[0008] Figure 3 illustrates an embodiment of a process for performing a device audit.

[0009] Figure 4 illustrates an embodiment of a process for performing a device audit.

[0010] Figure 5A illustrates a representation of a memory prior to an execution of the process shown in Figure 3.

[0011] Figure 5B illustrates a representation of a memory while the process shown in Figure 3 is occurring.

[0012] Figure 6 illustrates an embodiment of a process for performing a device audit.

[0013] Figure 7 illustrates an example of pseudo code for use in conjunction with auditing a device.

[0014] Figure 8 illustrates an example of a process for performing a device audit.

[0015] Figure 9 illustrates an embodiment of an environment in which device auditing is provided.

[0016] Figure 10 illustrates an embodiment of a portion of a device.

[0017] Figure 11 illustrates an embodiment of a process for performing a device audit.

[0018] Figure 12 illustrates a portion of memory being read in accordance with a step.

[0019] Figure 13 illustrates an embodiment of an implementation of a process for selectively reading memory.

[0020] Figure 14 illustrates an embodiment of an implementation of a process for timing a portion of a device audit.

DETAILED DESCRIPTION

[0021] The invention can be implemented in numerous ways, including as a process; an apparatus; a system; a composition of matter; a computer program product embodied on a computer readable storage medium; and/or a processor, such as a processor configured to execute instructions stored on and/or provided by a memory coupled to the processor. In this specification, these implementations, or any other form that the invention may take, may be referred to as techniques. In general, the order of the steps of disclosed processes may be altered within the scope of the invention. Unless stated otherwise, a component such as a processor or a memory described as being configured to perform a task may be implemented as a general component that is temporarily configured to perform the task at a given time or a specific component that is manufactured to perform the task. As used herein, the term 'processor' refers to one or more devices, circuits, and/or processing cores configured to process data, such as computer program instructions.

[0022] A detailed description of one or more embodiments of the invention is provided below along with accompanying figures that illustrate the principles of the invention. The invention is described in connection with such embodiments, but the invention is not limited to any embodiment. The scope of the invention is limited only by the claims and the invention encompasses numerous alternatives, modifications and equivalents. Numerous specific details are set forth in the following description in order to provide a

thorough understanding of the invention. These details are provided for the purpose of example and the invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the invention is not unnecessarily obscured.

[0023] Figure 1 illustrates an embodiment of an environment in which device auditing is provided. In the example shown, device 102 is a cellular telephone. Device 102 is in communication (e.g. via network 104) with verifier 106. In Figure 1, device 102 communicates with verifier 106 via a 3G network. Verifier 106 is under the control of a carrier, such as a provider of telephony service to device 102. Verifier 106 includes a database of hardware configuration information, including an entry corresponding to device 102 and the amount of RAM included on device 102.

[0024] As will be explained in more detail below, device 102 can be audited such that any evasive programs (e.g., malware) present on the device can be detected and/or removed. In some embodiments this is accomplished through the undertaking of a sequence of modifications to physical memory included on device 102. Results associated with the performance of the memory modifications are verified by verifier 106. Once device 102 is determined to be free of influence of such evasive programs, additional scans can be performed, which are also described in more detail below. For example, in addition to detecting malware (e.g., software installed without a user's knowledge and/or consent), the techniques described herein can detect "jailbreaking" actions (e.g., privilege escalations) taken by the user, such as to circumvent digital rights management installed by a carrier or hardware manufacturer.

[0025] A variety of devices can be used in conjunction with the techniques described herein. For example, in some embodiments device 102 is a video game console. The video game console is configured to communicate with a verifier under the control of the manufacturer of the game console via the Internet (104). If the owner of device 102 makes an unauthorized change to device 102 (e.g., by using a modification chip), verifier 106 will be able to detect the modification accordingly.

[0026] Other examples of devices that can be used in conjunction with the techniques described herein include desktop computers, notebook computers, netbooks, personal digital

assistants, video playback devices (e.g. televisions, DVD players, portable video players), routers, access points, settop boxes, medical devices, and virtually any other device that includes a processor and a memory.

[0027] In various embodiments, verifier 106 is controlled by a user of device 102, instead of by a separate entity. For example, a desktop computer owned by the user of device 102 can be configured to provide verification services to device 102. In that scenario, device 102 can be configured to communicate with the verifier via a local network. Device 102 can also be configured to communicate with verifier 106 directly (e.g., via a dedicated cable) and network 104 is omitted as applicable.

[0028] In some embodiments a verifier is collocated with or otherwise directly coupled to device 102. For example, a subscriber identity module (“SIM”) card inserted into a cellular phone can be configured to provide the functionality of verifier 106 to the cellular phone. As another example, the functionality of verifier 106 can be integrated into a power cord used to charge a cellular phone. In such embodiments, an external verifier can either be omitted, or can be used in addition to the verification services provided by the collocated/coupled verifier. As one example, suppose device 102 is a personal video player with integrated WiFi capabilities. A power cord used to charge the device can be configured to provide verification services to the device each time it is charged. In addition, if the WiFi radio is active, the device can be configured to communicate periodically with a verifier provided by the manufacturer of the device. As another example, a verifier 106 can be included on a USB device that is periodically inserted by a user into a laptop 102. In addition, whenever a user of laptop 102 attempts to conduct banking transactions with an online bank, the bank can also provide verification services to the laptop 102 prior to granting access to the user’s account. As yet another example, a network operator or service provider can require a user to have his or her machine audited before he or she is allowed on the network or allowed to access a service. User can also initiate an audit, for example, after realizing that he or she has been exposed to a potentially risky situation. One way a user can initiate an audit is to select a menu option on the device. Another example way is for the user to request an audit from verifier 106 (e.g., by submitting an online request through a web form).

[0029] Figure 2 illustrates an embodiment of a device. In the example shown, device 102 includes a processor 202, a first memory 204, a second memory 206, and a

communications interface 208. As one example, device 102 includes a 528 Mhz ARM processor (202), 128MB of RAM (204), a micro SD card slot into which a user has inserted a 1GB micro SD card (206), and a 3G modem (208). Memory 204 is also referred to herein as “fast” memory. Memory 206 is also referred to herein as “slow” memory. However, memories 204 and 206 need not be different speeds. Other components may also be included in device 102, such as a GPS receiver (not shown). Elements, such as second memory 206, may also be omitted as applicable. One may refer to RAM that can contain active programs as fast, and consider RAM that can only store data as slow.

[0030] Using the auditing techniques described herein, the absence of active processes in fast memory can be verified. And, after that verification has been completed, all memory (e.g., both fast and slow) can be scanned to identify, classify, report and potentially modify the contents of the fast and slow memory, or portions thereof. The distinction between fast and slow memory can be made in a variety ways. For example, on a device with RAM, flash memory and a hard drive, it is possible to treat only the RAM as fast memory and flash memory and the hard drive as slow memory. It is also possible to treat both the RAM and the flash memories as fast memory and the hard drive as slow memory. It is also possible to consider all memory physically located on a given device as being fast, and all external memory accessible (or potentially accessible) by the device as slow. The turnaround time to communicate to external components will cause such external accesses to be slower, irrespective of the type and actual local access speed of the external memory. Depending of what types of memory are treated as fast vs. slow, the selection of parameters would be done accordingly.

[0031] As will be described in more detail below, the existence of unauthorized modifications to device 102 can be detected by configuring device 102 to run a series of modifications to memory 204 and examining the results. If for example, the time it takes to perform the modifications exceeds a tolerance of a predetermined length of time, or if a result determined in conjunction with the modifications does not match an expected result, the presence of an evasive program may be indicated. In various embodiments, the memory modifications are performed across all memory on a device (e.g. both memory 204 and memory 206), instead of being run only on fast memory such as memory 204.

[0032] Figure 3 illustrates an embodiment of a process for performing a device audit. In various embodiments, the process shown in Figure 3 is performed by device 102. The

process shown in Figure 3 can be initiated in a variety of ways. For example, the process can be initiated every time the user charges the device (e.g., by configuring the device to initiate the process when it detects a power supply). The process can also be initiated in response to the occurrence of a particularly large or unusual transaction, in response to a concern that the user is at risk (e.g., in response to the carrier receiving notification that a new vulnerability has been released by a nefarious individual), in response to the elapsing of a certain amount of time, etc. Additional examples of events that can trigger the initiation of the process shown in Figure 3 include an attempt by the user of device 102 to make a payment or otherwise engage in a financial transaction, an authentication attempt (e.g., the user of the device attempting to access to a bank account), and an access request being performed (e.g., a request for the download of a movie to the device).

[0033] The process begins at 302 when one or more hardware parameters that correspond to a hardware configuration is received. Example hardware parameters include the amount and the speed of fast memory 204. For example, in the case of the device shown in Figure 2, the hardware parameters would include “amount=128M” and “speed= 300Mhz.” Additional parameters that can be used include the number of cores, the type of bus, etc.

[0034] The hardware parameters can be received in a variety of ways. As one example, the SIM of a cellular phone can be configured to detect the amount and speed of installed memory. As another example, if a proprietary cable is used to connect device 102 to a power source (or to a computer or other device), the parameters may be known (and thus “received”) by virtue of the cable only working in conjunction with a device having a specific amount and speed of memory. As yet another example, a serial number of a device may indicate the amount and speed of fast memory 204 installed on a device. In various embodiments, the user (or a representative thereof) is requested to input memory parameters in a web form or a configuration file. Assumptions can also be made about the likely memory configuration of the device and a benchmarking program can be run to confirm whether the assumption is likely to be correct.

[0035] At 304, a sequence of modifications to a physical memory is performed. Examples of ways in which such modifications can be performed are described in more detail below. In some embodiments the sequence of modifications to be performed is determined by the verifier. The set of modifications to be made can be provided to the device in a variety of ways. For example, the sequence can be constructed on the device based on a seed value.

The sequence can be pre-loaded onto the device at time of manufacture, at time of delivery to the supplier or carrier, or at the time of purchase. It can also be loaded by user choice or by a service provider at any time after purchase (e.g., as an over-the-update or as a firmware update), or when needed to perform an audit. The parameterization can be performed by the manufacturer or supplier or carrier, given known specifications. It can also be performed through a lookup, e.g., of serial number, by a user or service provider. The parameters can be associated with the model or device name. If the device is reconfigured, e.g., by replacement or addition of components, then these new components can carry information about the new or additional parameterization. The components can also carry the entire set of instructions, instead of just the parameters. Alternatively, the serial numbers, names, or types of components can indicate the needed change in parameters. If it is believed that the client device is secure at the time of installation of the algorithm or a new component, then the client machine can also inquire what components are installed (as is typically done as a system is booted up), and set the parameters accordingly.

[0036] In various embodiments, device manufacturers offer to preload non-activated auditing software at no cost, and later request payment to activate auditing services (and/or the additional scanning services described in more detail below. The auditing software can subsequently be activated by carriers, on request by end users or service providers. The carrier collects payment for the activation and optionally forwards portions of the payment to the handset manufacturer, providers of auditing software, providers of additional scanning software (e.g., antivirus detection services), and any other parties involved in the transaction.

[0037] At 306, one or more results of the portion of the process performed at 304 are reported to a verifier. In some embodiments results are provided to proxy 906, which timestamps the results and provides them to the verifier. As will be described in conjunction with Figure 5, in some embodiments multiple iterations of modifications to the memory and communications with the verifier are made, and the processes shown in Figures 3 and 4 are adapted accordingly.

[0038] Figure 4 illustrates an embodiment of a process for performing a device audit. In various embodiments, the process shown in Figure 4 is performed by verifier 106. As explained above, in some embodiments the process shown in Figure 4 is performed by an entity separate from device 102 (such as on a verifier controlled by a carrier). In other

embodiments the process is performed by a verifier located on or otherwise physically coupled to device 102.

[0039] The process begins at 402 when results are received. For example, when device 102 reports results at 306, those results are received by a verifier at 402.

[0040] At 404, a determination is made as to whether the results received at 404 indicate that an expected sequence of physical modifications was made. Verifier 106 is configured with information such as the amount of time the execution of a sequence of memory modifications should take on device 106 (assuming no authorized modifications have been made). In some embodiments verifier 106 is also be configured to store additional information, such as seed values and the results of computations performed by device 102.

[0041] If the expected sequence of physical memory modifications is determined to have been made (e.g., device 106 performed the sequence of memory modifications), it is concluded (406) that no unauthorized modifications have been made to the device. And, any evasive processes that might previously have been active on device 102 have been neutralized. If the expected sequence of physical memory modifications is determined to have not been made (e.g., because the amount of time to perform the sequence is off, or computed results are incorrect), it is concluded (406) that an unauthorized modification has been made to the device. (e.g., that an evasive process is present on the device and is attempting to avoid detection). In various embodiments, error correcting codes are used to avoid errors due to network noise. Message-authentication codes and other authentication techniques can be used to avoid active tampering with contents. Encryption techniques can be used to obfuscate contents and make it impossible for eavesdroppers to determine the plaintext messages being transmitted.

[0042] Figure 5A illustrates a representation of a memory prior to an execution of the process shown in Figure 3. In the example shown, kernel 502, authorized program 504, unauthorized program (e.g., a malware agent) 508, and an auditor program 506 are loaded in RAM. Typically, in order to remain resident on a device, an evasive program needs to do one of two things. It must either remain active in RAM (or swap space), or it must modify a legitimate program, data, or configuration of the device to allow the malware agent to gain control after a scan has been performed. As will be explained in more detail below, using the techniques described herein, the presence of the malware agent can be detected, irrespective

of the techniques it employs to evade detection. In addition, using the techniques described herein, the presence of the malware agent can be detected even if auditor 506 is loaded after malware agent 504.

[0043] Figure 5B illustrates a representation of a memory while the process shown in Figure 3 is occurring. As will be explained in more detail below, auditor 506 is configured to clear memory RAM (and any swap space) except for the space used by auditor 506. In various embodiments, a minimalistic set of other services is also permitted to occupy RAM. For example, if device 102 supports 3G communications, the area of RAM occupied by a 3G driver/module is not cleared, so that auditor 506 can use the 3G modem to communicate with verifier 106. As another example, in some embodiments a microkernel is permitted to occupy a portion of RAM while auditor 506 clears the remainder of the RAM.

[0044] Figure 6 illustrates an embodiment of a process for auditing a device. The process begins at 602 when an auditor process running on device such as device 102 clears all portions of memory 204 (and any swap space) that is not claimed for use by the auditor. In some embodiments, this includes unloading the kernel, various drivers, and all other processes. In various embodiments, the unclaimed memory space is overwritten by a sequence instead of being cleared (e.g., with zeros). One example sequence is a pseudo-random sequence that is combined with the original memory contents, such as by using the XOR operation. This allows the unclaimed memory space to later be reconstituted by the repeated combination with a pseudo-random sequence that complements or equals the previously used pseudo-random sequence. The unclaimed memory space can also be overwritten with contents in a way that clears it, but which does not correspond to the typical erasing operation of the device. For example, it is possible to clear unclaimed memory by writing a sequence of 01010101 to it, or any other appropriate sequence.

[0045] In some embodiments, the auditor code comprises two components: a loader and a variable algorithm segment. The task of the loader is to load algorithm segments from non-RAM storage (e.g., something other than memory 204), and hand over control to the loaded algorithm segment. After an algorithm segment has completed, it hands back the control to the loader.

[0046] At 604, contents of memory 204 are reported to verifier 106. In some embodiments the entire contents are reported. In other embodiments, only descriptions of changes since the last audit are communicated.

[0047] At 606, the device receives a cryptographic seed from the verifier. The seed is expanded to a pseudorandom string and the string is written to RAM. An example technique for writing a string to RAM in accordance with portion 606 of process 600 is provided below.

[0048] At 608, the device receives a cryptographic key from the verifier.

[0049] At 610, the device uses the received key to compute a keyed hash of the entire contents of the device's RAM.

[0050] At 612, the device reports the resulting value to the verifier. Verifier 106 evaluates the results, e.g., according to the process shown in Figure 4.

[0051] In various embodiments, device 102 reports state information from computations at 606 and 610 at time intervals set by verifier 106. The use of such intervals provides assurance that the computations performed by device 102 are being performed within memory 204 (and not, e.g. a portion of memory 206).

[0052] Device 102 obtains updates, from verifier 106, of the seed and respective key on an as-needed basis. The use of updates provides assurance that device 102 is not outsourcing computation to an external fast resource. For example, in order to outsource the computation, an evasive program would have to forward the seed and key updates to the external device, which would introduce a measurable delay.

[0053] Verifier 106 verifies that both the final function value and partial results are correct and are reported by device 102 to the verifier within acceptable time bounds. An example technique for evaluating the time it takes an auditor to perform its tasks is provided below. As mentioned above, in some embodiments verifier 106 is external to device 102 and is operated by a party other than the owner of the device. In other embodiments, verifier 106 is under the control of the user of device 102.

[0054] After the process shown in Figure 6 has been completed, auditor 506 can restore the contents of the device, whether fully or partially, and return control to previously active processes or to a process performing further scans of memory contents. The contents

of the fast memory can be restored if they were swapped out to slow memory prior to the execution of the timed computation, or if the original contents were combined with a string, the latter allowing a similar combination to be performed, thereby recovering the previous state. It is also possible to restart the device by loading a “starting” state. It is further possible to first hand over control to a process that scans, reviews, reports and modifies the contents of memory, or any subset of these operations (described in more detail below). The reporting can be presented to verifier 106, or to a third party, such as one in charge of managing the processing of memory contents. In the latter case, verifier 106 may be in charge of assuring that there is no active malicious process, and the second verifier could be in charge of processing the memory of the device to determine whether it complies with a particular policy, which may be related to malware detection, digital rights management, or another policy identifying what device memory contents are desirable.

[0055] Example Adversarial Strategies

[0056] In order for an evasive program to avoid being detected, e.g., during portion 604 of the process shown in Figure 6, it must be active in RAM, either as a unique process (504) or as part of a corrupted version of auditor 506. The following are six example ways in which an evasive program such as malware agent 504 can attempt to remain active:

[0057] Strategy 1: Outsource storage.

[0058] The malware agent can stay active in RAM and attempt to remain undetected by causing auditor 106 to not clear the appropriate space (e.g., at 602) and rely on non-RAM storage or external storage to store the corresponding portion of the pseudo-random string generated at 606. The computation at 610 would then be modified to use the outsourced storage instead of the space where the malware agent resides.

[0059] Strategy 2: Compute missing data.

[0060] Instead of outsourcing storage of portions of the pseudo-random string, the malware agent can store a modified representation of the string (e.g., a compressed version, or a version that is missing portions), and reconstitute relevant portions of the string as they are needed during the computation of the keyed hash at 610. Since the malware agent has the seed from which the pseudo-random string is generated, it can use this – or later states of the pseudo-random generator – to regenerate required portions of data.

[0061] **Strategy 3: Outsource computation.**

[0062] The malware agent can forward relevant data to an external device (assuming the necessary communications infrastructure, such as a WiFi connection is still enabled). The external device receives data from device 102 and computes the values needed to report to verifier 106, feeding these values to the malware agent on device 102.

[0063] **Strategy 4: Modify detection code.**

[0064] The malware agent can attempt to replace the code of auditor 506 with modified code. This replacement code may be designed to suppress reports of compromised memory contents, or contain a hook for malware code to be loaded after the audit completes. The malware agent can attempt to incorporate such changes to auditor 506 without taking up more space by swapping out or compressing portions of the auditor code and loading or unpacking it again as it is needed.

[0065] **Filling Fast Memory**

[0066] This section describes an example technique that can be used in conjunction with portion 606 of the process shown in Figure 6.

[0067] Figure 7 illustrates an example of pseudo code for use in conjunction with auditing a device. In the example shown, the subroutine *get_permutation* returns a vector indicating a random permutation of *number_blocks* items, ranging from 0 to *number_blocks*–1, where *number_blocks* is the number of portions of size equal to a flash block that the RAM comprises, minus those needed by the auditor. The subroutine *next_string_chunk* returns a pseudo-randomly generated chunk of bits; the term *chunk* is used to refer to the amount of data that can be sent on the memory bus. As one example, for an Android G1 phone, a chunk is 32 bits.

[0068] Both *get_permutation* and *next_string_chunk* use the most recently provided seed as input. The pseudo-random string can be computed as $segment_i \leftarrow hash(segment_{i-1})$, i.e., in a way that *cannot* be computed using random access. One example is a function based on iterated application of the hash function, given the non-homomorphic properties of hash functions. A variety of hash functions may be used. One example is MD6 in 512-bit mode.

[0069] The constant *rounds* is the number of times a pseudo-random chunk is XORed into the contents of a cell, using the function *modify_memory*. The choice of *rounds* controls the amount of work an adversary has to perform to carry out the second adversarial strategy (computing missing data), while also incurring an increasing cost to the honest execution of the algorithm for large values. In the example shown, *rounds* = 2, which results in a noticeably greater cost to the adversary than *rounds* = 1, since the value of each cell will come to depend on two other cells. This can confound memory management strategies of an adversary. In the example shown, *chunks_per_block* is the number of chunks contained in a flash block, equaling 32768 (=128kB/ 32 bits) for an example G1 phone, while *number_blocks* = 1024 (= 128 MB / 128kB).

[0070] The function *modify_memory(pos, string)* XORs the contents of position *pos* with the value *string*, where *pos* = 0 describes the first chunk of RAM to be operated on, and *pos* = *number_blocks* × *chunks_per_block* - 1 is the last chunk.

[0071] The memory access structure described in conjunction with Figure 7 causes accesses to individual pages of randomly ordered blocks, if forced to use flash (e.g., memory 206) instead of RAM (204). This will cause the flash to be cleared with an overwhelming probability, and the pseudo-random access order prevents the adversary from scheduling the memory accesses to avoid this drawback. The cost of a flash-bound computation in comparison to the RAM-bound alternative available to the honest execution of the algorithm is noticeably more time consuming.

[0072] In some embodiments, one hash function application is used to generate several invocations of *next_string_chunk*. This reduces the computational burden associated with the auditing process, which emphasizes the contribution of the memory access in terms of the time to perform the task.

[0073] In various embodiments the input to the hash function is a constant number of previous outputs; this complicates storage for a malware agent wishing to reconstitute the state of a given portion of the pseudo-random generator, and is thus useful to further frustrate any attempt to use strategy 2 (compute missing data).

[0074] **Performing Timing**

[0075] This section describes an example technique that can be used for timing the execution of auditing tasks. For example, in some embodiments the technique is employed by verifier 106 as described in conjunction with the text corresponding to Figure 6.

[0076] Verifier 106 is configured to time the execution of portions 606 and 610 of the process shown in Figure 6, e.g., to identify attempts to outsource storage; compute missing data; and outsource computation.

[0077] In some embodiments verifier 106 is configured to obtain state information from device 102 at frequent intervals (e.g., that are set by verifier 106). One example of state information is the memory contents of the memory chunk that was last updated, which vouches for that device 102 has reached this stage of the computation. Verifier 106 sends update requests to device 102 at regular intervals. In some embodiments the update requests correspond to updates of the state of the pseudo-random generator used to compute the output of the subroutine *next_string_chunk*. If the output of the subroutine *next_string_chunk* is generated by selecting an unused portion from an already generated pseudo-random string, the string can be cleared at the same time, thus forcing the new seed to affect the state immediately.

[0078] An evasive program employing adversarial strategy 3 (i.e., outsourcing computation), must transmit the update of the pseudo-random string to the external device that performs the computation, after which the external device has to compute the resulting next value to be reported by device 102 to verifier 106 and transmit this to the evasive program. This incurs a round-trip delay. If the round-trip delay exceeds the time between timing checkpoints, the cheating will be detected. Here, an assumption is made that seeds and keys, along with other status information, is communicated securely between the client device and the verifier. Various cryptographic techniques can be used to achieve this.

[0079] In various embodiments, the device-specific time between checkpoints is chosen so that there is not enough time to outsource computation using communications equipment (e.g., WiFi) included on device 102, pessimistically assuming congestion-free environments.

[0080] The execution time of *modify_memory* is determined by the parameter selections described above and what hash function to use to compute *next_string_chunk*. For

example, the MD6 hash function can be configured to different output sizes, from 224 to 512 bits. As explained above, in some embodiments a 512-bit version is used. The time per invocation of *modify_memory* is noticeably less than the time between checkpoints determined above.

[0081] Examples of Detecting Various Evasive Programs

[0082] The following section provides examples of how evasive programs employing the various strategies described above can be detected using the techniques described herein.

[0083] Defending against adversarial strategy 1 – outsource storage.

[0084] Assume an empty SD card has been inserted into device 102. The corresponding write speed could reach up to 5MB/s. The size of a block processed by *modify_memory* as described above is chosen, in this example, to be 128kB. The time to write the data to the SD card would be 25ms. In comparison, suppose RAM on device 102 has a write speed of 100MB/s. The corresponding write time would be 1.25ms. The additional delay can be readily detected. And, if multiple accesses to the SD card are made between two checkpoints, the additional delay will be even more readily detected.

[0085] Defending against adversarial strategy 2 – compute missing data.

[0086] As mentioned above, the pseudo-random string can be computed in a way that cannot be computed using random access. To compute the value of a certain output, the corresponding input needs to be computed from stored data. Since *rounds* > 1, the data stored in RAM is not this needed state, but a combination of the states of the two or more rounds. The state needs to be explicitly stored (in RAM) by the malware agent, as part of its code space, and the needed state computed from this. This forces the malware agent to compute at least (and in fact, *much* more) than *rounds* × *number_blocks* × *chunks_per_block* hash operations during the execution of portion 610 of the process shown in Figure 6, in addition to the computation it needs to perform to “legitimate” computation. For the selected parameter choices, that is more than 100 million hash operations. Given an approximate time of 10 μ s for computing a hash function invocation, this is about 1000 s, which is approximately 3 orders of magnitude more than the expected and can be detected accordingly.

[0087] A sample proof corresponding to the defense against strategy 2 will now be provided. Assume that an evasive program resides in memory 204 and takes up at least portions of some c 32-bit chunks for itself and its variables. A pessimistic assumption can be made that *all* of this space can effectively be used to store variables, which is not possible, but which gives a lower bound on the amount of work that the malware has to perform to remain undetected. In reality, its effort is greater as not all c chunks can be used for storage, but some are needed to store its code.

[0088] For each one of the c hits to RAM chunks that do not contain the values needed to compute the function, the malware agent has to compute the expected contents. It is assumed in this example that the original contents – before the RAM-filling was performed – were zero. If this is not so, the effort of the malware agent would be greater, so making this assumption establishes a lower bound on the effort of the malware agent. To compute the expected updates to this cell that would have been performed by the RAM-filling algorithm, the malware agent needs to compute the values for all the *rounds* passes on the memory chunk in question. The values XORed in to memory come from a pseudo-random sequence. And, it is only possible to reconstitute the state of the chain in the cell where it is missing by computing the value *next_string_chunk* from a value that is stored by the malware agent in part of the c chunks of storage. It is assumed that the variables are stored solely in RAM, or the malware agent needs to *also* succeed with strategy 1 (outsource storage).

[0089] As explained above, the pseudo-random generator cannot be computed using a random-access approach. It is the case that $L = 16$ chunks is needed to store a state, given a chunk size of 32 bits and a state size (= MD6 output size) of 512 bits. The malware agent has to recompute the sequence of hash function invocations from a position of RAM associated with this state (which does not have to be where the malware agent *stored* this state.)

[0090] Given the random permutation over cells during the writing of the memory (the order which cannot be anticipated by the malware agent), the expected length of the run to the string position corresponding to the stored state is at least $rounds \times n / (c / L)$, where $n = number_blocks \times chunks_per_block$ corresponds to the number of chunks that RAM consist of, $rounds \times n$ is the length of the pseudo-random string, and where c / L are the number of pseudo-random states stored by the malware agent. Thus, for each hit to a “bad” cell, the malware agent has to perform an expected $rounds \times n \times L / c$ invocations of

next_string_chunk, which corresponds to $rounds \times n \times c$. There are c such hits, not counting hits to “bad” cells that occur as the malware agent attempts to compute one of the expected states. Therefore, the malware agent has to perform at least $rounds \times n$ hash operations to compute the contents of the c bad blocks from the stored contents. The approximate time to do this (according to an example implementation) is at least between 100,000-1,000,000 times slower than the legitimate client which indicates that any attempt to compute missing data will be detected.

[0091] If the computation of the chain causes the access to a cell that has been used to store a value for another pass of the computation, then this causes another hit to be incurred. It would happen with an approximate probability

$(c - c / rounds) / c \times c / number_blocks = (c - c / number_blocks) / number_blocks \approx c / number_blocks$ for each memory access, and therefore, with approximate probability

$1 - (1 - c / number_blocks)^{number_blocks \times rounds^2 / c}$ for a given first bad cell hit, as described above. A rough approximation of this quantity is $1 - e^{-rounds^2}$. For $rounds = 2$, this is more than 98% probability. This additional cost would increase with increasing values of c . An adversary would therefore do best to make c small.

[0092] In the following, assume that the adversary uses $c = L = 16$ cells only, all 16 to store one value. With this configuration, the adversary would fail to compute the value (unless using external memory) in those situations where the chain leads in a direction that does not make it possible to compute the value corresponding to the “programcell” from the value in the “storage cell”. For $rounds = 2$, this failure occurs with probability 75%. In the remaining 25% of the cases, the adversary would simply be slowed down. (To always succeed to compute the value, the adversary needs to store at least $round = 2$ values, each 512 bits long.

[0093] **Defending against adversarial strategy 3 – outsource computation.**

[0094] In some embodiments the time between checkpoints is chosen so that there is no time to outsource computation using the communications equipment on the device. The time between checkpoints can be chosen by verifier 106 in a way that makes this immediately detectable. A strategy involving outsourcing of computation will fail, as the roundtrip has to be completed between two checkpoints for the right value to be provided by the adversary.

This is independent of the speed of communication between the client device and the verifying party.

[0095] Defending against adversarial strategy 4 – modify detection code.

[0096] Suppose unauthorized program 508 corrupts the execution of some steps (e.g., as described in conjunction with Figure 6), then willingly loads legitimate code and removes itself. Such an adversary could potentially corrupt portions 602 and 604 of the process, but will not be able to corrupt portion 606. Specifically, it needs to corrupt portion 602 of the process (clearing swap space and RAM) in order to maintain active. It can then cause a misreporting of state at 604. However, this will be detected when the keyed hash of the memory contents are computed (610). This is both due to the assumed collision-freeness of the hash function used, and the fact that the key is not disclosed to the device until 608. Portion 608 cannot be corrupted without being active during 606, which in turn would cause detection, as described above. And, the evasive program will be unable to compute the correct values to be reported at 612 without executing portion 606 of the process shown in Figure 6.

[0097] Combinations of the four adversarial strategies will also fail, since each of them will be detected and combinations of them do not change the underlying device-specific limitations.

[0098] Additional Processing

[0099] Figure 8 illustrates an example of a process for performing a device audit. In various embodiments, the auditing processes described above form one phase (802) of a two (or more) phase process. Once the techniques described above have been employed with respect to a device such as device 102, an assumption can be made that no evasive software is active in the RAM of the device. Arbitrary additional processing can then be performed on the device (804). Examples of additional processing that can be performed are described below.

[00100] Example: Malware

[00101] After the processing of 802 is performed, at 804, device 102 performs traditional antivirus software to identify known bad software, such as may be stored in

memory 206. At 804 device 102 can also be configured to report the entire contents of memory 206 or portions of memory to verifier 106 or to another device.

[00102] Example: Jailbreaking

[00103] After the processing of 802 is performed, at 804 device 102 determines whether its operating system loader has a particular hash, and/or otherwise determines whether the operating system loader has been altered away from a desirable state.

[00104] Example: Phone Unlocking

[00105] After the processing of 802 is performed, at 804 device 102 determines whether its operating system loader has been altered and also determines whether any information associated with a service provider has been altered.

[00106] Example: Software Piracy

[00107] After the processing of 802 is performed, at 804 device 102 determines whether any software included in memory 206 has been modified from an expected configuration, ascertains any associated serial numbers for the software, and/or otherwise determines whether the included software is being used in an unauthorized/unlicensed manner. In some embodiments device 102 reports the contents of memory 206 or portions thereof to verifier 106.

[00108] Example: Media Piracy

[00109] Suppose that media files (e.g., music, video, or image files) are customized using watermarks during distribution, and that these watermarks are cryptographically authenticated, e.g., using a MAC or digital signature. At 804 it can be determined which files present on device 102 have legitimate watermarks, and whether these contain valid authenticators. The determination can be made either locally to device 102 or centrally (e.g., on verifier 106).

[00110] In various embodiments, applications (such as a music player installed on device 102) record usage and other data (forming a log of activity) and associate the information with the appropriate media (e.g., song file). The logs can be read by the verifier 106 at 804.

[00111] Example: Chains of Custody / Usage Logs

[00112] Suppose an application (or data file) has an associated log file that is used to record transactions. One example is a log file that records the occurrence of financial transactions, including stored-value information. The legitimacy of changes made to the log file can be verified as follows. First, the processing of 802 is performed. Then, at 804 a determination can be made (e.g., by comparing hashes of program images) as to whether or not the application (or data file) has been altered, and thus, whether the log file is genuine.

[00113] One approach to the processing performed at 804 in this example is as follows: First, memory 206 is scanned and a list of applications and data files associated with the applications is created. Next, a list of descriptors for the applications and data files is determined. An example of a descriptor is a hash of the file, along with the name and type of file, and an identifier stating what string sequence(s) it matched. Next, a second list of any descriptions of applications or data that is not already reported on in the first list is made. The description created here may include all or parts of the code for an application, or of a description of what type of input files it processes and output files it produces. The second list is transmitted to an external party, such as verifier 106, where it is verified. The second list can also be processed locally using any policy obtained from a policy verification server.

[00114] The outcome of the verification can be used to affect the permissions to the applications and data, and can be used to control how external servers interact with the device, including whether it is granted access to network resources (such as the Internet, 3G networks, corporate networks, etc.). As another example, the software allowed to run on the device can be restricted, and notify the user of lack of compliance, attempt to remove or repair or otherwise modify files, etc.

[00115] Example: Parental Control Filters and other Monitoring Features

[00116] After the processing of 802 is performed, in various embodiments, additional middleware is installed that can be configured to log (and/or block) various events associated with the device. Examples include:

[00117] (a) determining what photos were generated on the device and later transmitted out (e.g., to prevent “sexting”).

[00118] (b) determining (e.g., based on device activity and GPS changes) whether the device was used (e.g., for texting or watching video clips) while travelling at a speed greater than 20 miles per hour.

[00119] (c) determining (e.g., based on installation activity) whether alternative applications (such as a second instant messaging program in addition to a default program) has been installed, and then creating a log file for the alternative application.

[00120] (d) determining (e.g., based on browser history information) what URLs a user has visited including which URLs were manually navigated to and which URLs were referred to in other HTML documents that were accessed. One benefit of this logging is to identify whether a person is likely to have fallen victim to phishing; has visited a web site known to distribute unwanted content, including malware; and whether the device is likely to be involved in click-fraud. Such abuse is possible to achieve without infection of the device itself, e.g., by use of JavaScript, cascading style sheets, and/or other related scripting languages.

[00121] **Example: Additional Applications**

[00122] In addition to the above examples, yet more uses of the techniques described herein are possible. For example, device auditing can be used in vehicular black-boxes for metering usage, purposes of insurance, tariffs, taxes, tolls, etc. - both to identify malware and intentional tampering.

[00123] The device auditing technique can be included as a component in other applications, allowing these applications to temporarily suspend themselves to perform the scan, and later be given control again, in the known clean state.

[00124] As yet another example, the techniques can be used in medical devices, to determine that they are free from infection, correctly configured and maintained, and in order to audit usage in special cases when it becomes valuable to know who had access to data and equipment. The devices in question may log usage information at all time, in a way that preloaded application cannot interfere with; the audit process would include a memory-printing scan to assert that the preloaded applications are still in a good state, and that no conflicting applications or configurations are present.

[00125] Finally, the techniques can be used for detection of malware in situations where there is no need to remediate, or where that is not the primary goal. One such context is for online games, to detect the absence of modules to cheat in the games.

[00126] **Preserving Privacy**

[00127] In some embodiments descriptions of all state (e.g., the contents of memory 204) is communicated to the verifier 106. However, some data should preferably not be transferred off device 102, such as private keys and non-executable private data. In the following section, techniques preserving the privacy of such data are described.

[00128] Assume that a first random number is called x , and that it is selected from some space of possible values, $1 \dots \text{max}_x$. It is possible that x encodes malware apart from providing an input to the auditing process for which it was intended. A legitimate program computes a one-way function value y from the input data x and some system parameters, which is called (g_1, n_1) . One example way of doing this is by computing $y = g_1^x \text{ modulo } n_1$, where g_1 generates a large subgroup of G_{n1} .

[00129] Let the program then compute a second one-way function value z from the value y and some system parameters, which is called (g_2, n_2) . One example way of doing this is by computing $z = g_2^y \text{ modulo } n_2$, where g_2 generates a large subgroup of G_{n2} .

[00130] Next, it is assumed that the client machine proves (e.g., using a zero-knowledge proof) that there is a value x such that $z = g_2^{g_1^x \text{ modulo } n_1} \text{ modulo } n_2$, where (z, g_1, g_2, n_1, n_2) are known by the verifier, but (z, x) are not. The device (the “prover”) then erases the value x but stores (y, z) and the parameters (g_1, g_2, n_1, n_2) .

[00131] At later times, the device has to prove that the value y that it stores, but which is secret, corresponds to the value z . (Here, z can be stored on device 102, but can also be stored by verifier 106.) One example proof that can be used is a zero-knowledge proof.

[00132] If the second proof concludes and verifier 106 accepts it, then the verifier knows that the unknown value z that the client stores is of a format that cannot be used to hide a significant amount of data of value to a malware agent.

[00133] Here, z can be used to encrypt some other data, which is referred to as m , and whose ciphertext is referred to as c . Thus, $c = E_z(m)$ for an encryption algorithm E . Assuming

symmetric encryption, $m=D_z(c)$ for some decryption algorithm D . The device contents can be verified, but m remains unknown by the party receiving c . This party would not know z , but only that z is of some acceptable form that cannot hide large amounts of malware data. Since the auditing process described herein allows the verifier party to be assured that only legitimate programs exist in the RAM of the client device, it can be known that the programs – using the secret value z – can access m , given c . However, the verifier cannot.

[00134] Since it is known that the accessing program is legitimate, it is also known that m will only be accessed in an approved manner. For example, if m is data and not code, then it is the case that the accessing program will not try to execute the data.

[00135] **Using A Pseudo Random String Generator**

[00136] Figure 9 illustrates an embodiment of an environment in which device auditing is provided. In the example shown, device 902 includes, in addition to the components illustrated in Figure 2, a SIM configured to serve as a proxy (906) for external verifier 904. As will be described in more detail below, a monolith kernel stored in the instruction cache of device 102 (where it fits in its entirety) swaps out all other processes (excepting any processes it chooses to except) when it is activated, and performs an auditing process. The monolith kernel has an associated working space that is located in the data cache (and registers). The cache is typically implemented using RAM, and is considered as being part of it herein. As used herein, “free RAM” is the portion of RAM which ought to be free after all applications – including the regular kernel – have been swapped out. In some embodiments, “free RAM” is defined as the segment of RAM that is not taken up by a set of approved routines and data. For example, the regular kernel may be an approved routine, as may common and whitelisted applications. Moreover, approved data may correspond to data that is known by the external verifier, and may be of any format, as long as it is whitelisted (i.e., believed to be safe.). In such cases, the approved programs need not be swapped out to secondary storage (as described in more detail below) but may instead remain resident during the memory reading portion of the audit (e.g., 1108).

[00137] In some embodiments the monolith kernel corresponds to a program F_ε , parameterized for a known execution environment ε . As explained above, the execution environment corresponds to a hardware configuration of the device. Executing F_ε on input x produces a sequence of outputs $F_{\varepsilon i}(F_\varepsilon, x)$, each within a time $t_i(F_\varepsilon, x)$ from the start of the

execution and produces an ending state $s(F_\epsilon, x)$. In this example, $x \in X$, where X is the set of all legitimate inputs.

[00138] Proxy 906 is used to reduce latency variance from the device, and in various embodiments is implemented as a tethered cellular phone, a cell phone tower, etc., instead of or in addition to a SIM. In some embodiments external verifier 904 performs an initial computation (described in more detail below) and communicates (e.g., via a secure channel), part of the information to proxy 906, using device 902 as an intermediary. Proxy 906 times computations performed by the monolith kernel and reports the timing measures back to external verifier 904. In some embodiments, external devices, such as tethered cell phones or computers, base stations, or additional external verifiers are used instead of or in addition to proxy 906. It is also possible to use software proxies that are believed to be tamper resistant, or to use special-purpose hardware proxies.

[00139] Figure 10 illustrates an embodiment of a portion of a device. As mentioned above, “free” RAM is defined as being the portion of RAM that ought to be free after all applications and the standard kernel have been swapped out. The width of the bus is a word. The size of memory is also describable in words. For example, a 512 byte memory page as shown in Figure 10 has a size 128 words on a standard handset, where a word is 32 bits. As used herein, a “chunk” is the length of the cache line. In the example shown, the cache line corresponds to 8 words, each of which is 32 bits, and the chunk is 256 bits accordingly.

[00140] Figure 11 illustrates an embodiment of a process for performing a device audit. In various embodiments, the process shown in Figure 11 is performed by device 902. The process is configured such that its computations are expected to complete in a particular amount of time. Any change to the amount of free RAM evaluated and any attempts to access secondary storage 1004 will result in an observable lengthening of the amount of time the computations take to complete. Similarly, changing the contents of any of the whitelisted programs or associated data will cause a delay or the computation of the incorrect responses.

[00141] The process shown in Figure 11 can be initiated in a variety of ways, such as the ways described in conjunction with the process shown in Figure 3. As additional examples, the audit process can be included in a shutdown or boot route. It is also possible for an application to initiate the audit process. The application would be deactivated, the processing would be performed, and control would be handed back to the application when

complete. In some embodiments an application queries a central authority (or the device) for information about how recent the last scan was performed. The SIM card can store information about when a scan was performed. If the SIM card has the functionality that allows it to constantly measure time, it can give an actual time as the answer. Otherwise, it can give a time estimate based on the number of seen transactions, many of which will be known to be periodic. Such information can be used to assess the duration since the last scan.

[00142] The process begins at 1102 when contents of memory 1002, except monolith kernel 1006 (and any processes it deems acceptable to retain) are swapped to secondary storage 1004. In some embodiments portion 1102 of the process includes swapping out the normal kernel, or parts thereof. Crucial features, such as device drivers for serial communication, are re-implemented in monolith kernel 1106. In various embodiments, the contents are swapped out verbatim, or compact descriptions of the contents are swapped out or stored on the proxy, external verifier, or other trusted device, or stored in RAM in a state that cannot be used for active code. (E.g., it is possible to store instructions in a portion of cache not intended for instructions, but only data.) In some embodiments, no “free” space exists and portion 1102 of the process shown in Figure 11 is omitted.

[00143] At 1104, one or more hardware parameters that correspond to a hardware configuration are received. This portion of the process is similar to portion 302 of the process shown in Figure 3. Also at 1104, initialization information, such as a seed that can be used to generate a pseudo-random string, is received. Other examples of initialization information include a step value and a key value, as described in more detail below.

[00144] At 1106, the free RAM is overwritten. In some embodiments the output of a pseudorandom string generated using the seed is used to overwrite the free RAM. One technique for overwriting the free RAM is to generate an n^2 -bit pseudorandom string with the particular property that the computation of any one of the output bits will take at least half as long as the computation of the entire block of bits or at least 512 applications of MD6 in 512 bit mode. The technique uses three phases, and repeats (with different values of aux) until the output strings fill the entire free RAM:

[00145] 1. Generating: Using a hash function h (e.g., MD6) with an output size of “ n ” bits, generate the value $x_i = h(seed, i, aux)$ for $0 \leq i \leq n-1$, and some value aux. This generates n^2 pseudorandom bits.

[00146] 2. Shuffling: Compute $y_j = \prod_{i=0}^j 2^i \text{BIT}_j(x_i)$, $0 \leq j \leq n-1$, where BIT_j is a function that returns the j^{th} most significant bit of the input. This shuffles the bits in a manner that requires computation of all n hash function applications to reconstitute any one of the values.

[00147] 3. Blending: Compute $z_j = h(y_j)$, for $0 \leq j \leq n-1$. This assures that each bit of the output is a function of all the n input bits, each one of which required one hash function evaluation to compute.

[00148] In various embodiments, additional shuffling and blending is performed to further increase the cost of computing any part of the final string. In addition, other techniques for overwriting free RAM can be used instead of the example technique described in conjunction with portion 1106 of the process shown in Figure 11.

[00149] At 1108, memory 1002 (or portions thereof) is read in a manner determined by the “step” value. Results are accumulated and the computation is keyed using a key. In various embodiments, the processing of portion 1108 is performed by a memory access scheduler and an accumulator, each of which will now be described in more detail.

[00150] **Memory Access Scheduler**

[00151] Let “sRAM” be the size of RAM 1002, measured in its entirety, in chunks. External verifier 904 will select a random value step in the range $page < step < sRAM - page$, such that “step” is an odd value. Here, “page” denotes the size of one memory page in secondary storage, also measured in chunks. In the case where there are several page sizes (e.g., if there are several components that constitute secondary storage), in various embodiments the largest of the page sizes is used.

[00152] Performing the processing of 1108 includes a loop in which memory is accessed and the results combined to form a keyed memory checksum. For each iteration of the loop, the access location is increased by the value step, modulo sRAM. Because “step” and sRAM are relatively prime, all RAM memory locations will be accessed exactly once. Further, the access order will be unknown to an adversary until the value “step” is disclosed. An illustration of memory 1002 being read in accordance with a “step” is provided in Figure 12.

[00153] In the example shown in Figure 9, device 902 includes a single, single-core CPU. In systems such as laptop computers that include multi-core processors and/or multiple processors, the processing of 1108 can be constructed in a way that is either inherently serial (and which will therefore obstruct the use of multiple processors) or adapted to take advantage of multiple processors. As one example of the latter, several computations can be started with offsets such that each thread corresponds to a different portion of memory, and where there are no collisions.

[00154] **Accumulator**

[00155] Memory contents can be accumulated in a register using a simple non-linear function that combines the previous register contents (referred to herein as a “state”) with the newly read memory contents (data), one by one. Examples of accumulating functions include hash functions (e.g., MD6); a non-linear shift-back register; and simpler functions.

[00156] One example of a simpler function is $\text{state} \leftarrow \text{ROR}(\text{state} + \text{data})$. The latter function corresponds to a function $\text{ROR}(\dots(\text{ROR}(\text{state}_0 + \text{data}_1) + \text{data}_2) \dots + \text{data}_n)$, where “+” refers to regular addition, and “ROR” rotates the contents of the register one bit to the right. In this case, the function itself may not be nonlinear, but when combined with the a priori unknown step size and the tight timing requirements, it is nonetheless sufficient to satisfy the processing requirements needed.

[00157] As mentioned above, in various embodiments, the accumulation process is keyed. One way to accomplish this is to offset the value “state” with a new value “key” (obtained from the external verifier or the proxy) at regular intervals. The offset can be performed by adding the current value state to the new value key.

[00158] Further, while the process described in conjunction with 1108 is based on reading memory, in some embodiments a write operation is included to cause further flash-based slowdowns. As one example, sequences of “1”s are written, as causing an erase of the entire block, should the data be stored in flash. To simplify the scheduling of where to write (and with it the monolith kernel), the location can be obtained from the proxy at the same time a new key value is obtained.

[00159] Various other sequences of memory accesses can also be performed at 1108. For example, it is possible to use two step values instead of one, where these step values may

both be even numbers, but where they cause mostly all of the space to be covered. It is also possible to use a collection of numbers or parameters that determine a function that selects the sequence of locations. It is possible to think of this as a maximum-length sequence, where the outputs are locations, and the maximum-length sequence includes all values in a given range, corresponding to memory positions. It is possible to offset such values to avoid accessing certain areas (e.g., the monolith kernel), should that be desirable. In the case of a maximum-length sequence, the key provided by the external verifier or proxy can be the initial state, or the weights associated with the various cells of the LFSR.

[00160] At 1110, the keyed computation is provided to an external verifier 904. If the external verifier approves of the results, device 902 is deemed to be in a safe state.

[00161] At 1112, device 902 executes any functions that are to be executed in a safe state. Examples include setting up an SSL connection, casting a vote, entering a password, scanning secondary storage for malicious/unwanted programs, etc. In various embodiments, if the code of the safe-state function is in secondary storage (i.e. it is not part of the monolith kernel), a digest of the function is compared to a value stored in the monolith kernel (or on the proxy). The function is activated only if the values match. In various embodiments, if the proxy can perform the message digest computation, the monolith kernel need not contain code to do the same.

[00162] At 1114, the state of RAM 1002 is restored by loading contents (swapped out at 1102) of secondary storage 1004.

[00163] A large portion of the potential load of the process shown in Figure 11 involves swapping out applications and data from RAM to secondary storage, and swapping it back. It is possible to avoid doing this, e.g., in order to save time. This could be done by killing the applications. If an external verifier or other resource knows what applications are running, and potentially also their state or parts thereof, it is possible for this party to assist in restarting selected applications after the auditing process has run. It is possible for secondary storage or a SIM card or other on-board unit to maintain some of this information. It is possible to identify applications and data not by their full strings, but by shorter identifiers to save space and time. It is possible to have approximate algorithms that largely re-generate the same state after the detection algorithm has run. For example, this may restart a browser, but fail to recover the browser contents.

[00164] Further, it is not necessary to swap out active applications if they only take up some portion of RAM. For example, suppose they only take up the lower half of RAM. For each cell (number i) in free RAM, copy the contents of that cell to a position higher up (position $2i$). This is preferably done starting from the end (higher numbered positions.) This effectively slices the applications, and makes them reside only in even positions. Now pseudorandom values need only be written into the odd-numbered positions, and only need to perform the nonlinear accumulation of the odd-numbered cells. Note that it is not possible for any functional malware to remain active. It is possible for malware to exist, but only if its instructions are “jump to next open space” and that is where the next instruction is. Since all the space that is not overwritten by pseudorandom material will be jumps only (there is not space for more in consecutive space), it is knowable that the malware cannot achieve anything. It is possible to make the distance between slices larger, especially if the space is not predictable by an adversary. The distance may be predicted by a sequence generator, for example, where different distances are different length. The spreading out of data and programs within the RAM can be combined with offsetting these with random strings. The microkernel (the program in charge of the malware detection)—would not be spread out in this manner, as it needs to remain in a state that allows it to execute.

[00165] In conjunction with Figure 11, a description was made as to how RAM could be read using a step size that is not a priori known. It is also possible to write (free) RAM using a step size other than one, where the step size may be unknown a priori. One effect of this is that it causes writes to secondary storage for an attacker wishing to store information in secondary storage. These delays are greater than read delays if secondary storage uses flash. It is possible to use a simple step increment modulo an integer corresponding to the range to be written -- plus an offset if the monolith kernel does not reside in high addresses. It is also possible to use no particular modulo for the arithmetic -- which corresponds to using a modulo corresponding to the addressable RAM space -- and to identify whether the cell to be written is in a range that should not be written.

[00166] Figure 13 illustrates an embodiment of an implementation of a process for selectively reading memory. In some embodiments Figure 13 is used to implement a portion of portion 1108 of the process shown in Figure 11.

[00167] In some embodiments, full use of the data cache is made. Specifically, data is processed in chunks the width of the cache line, whose size typically ranges from 8 to 512

bytes. In the following, the parameter “ α ” denotes the width of the cache line, in words. The parameter “period” corresponds to the length of the timing periods. This is set to correspond to a time that is shorter than an optimistic roundtrip latency from device 902 to a potential external helper. The reported values, and the time at which they are reported, are recorded by proxy 906.

[00168] Examples of Detecting Various Evasive Programs

[00169] The following section provides examples of how evasive programs employing the various strategies described above can be detected using the techniques described herein. To succeed in evading detection, an adversary has to provide the correct responses $F_{\epsilon_i}(x)$ within the expected time $t_i(F_{\epsilon}, x) + \delta$. There are various ways this can be attempted by an adversary, each of which will now be described.

[00170] Defending against adversarial strategy 1 – external support.

[00171] The values “key” and “step” are communicated to the proxy over a secure channel. The value of “state” is reported for each iteration of the outer loop. The computation of this value depends on the value “key.” Therefore, it is not possible for an external device to help with the generation of the responses if the roundtrip latency (from sending out the value key to receiving the result of the computation) exceeds the time threshold $t_i(F_{\epsilon}, x) + \delta$, where $t_i(F_{\epsilon}, x)$ is the time for F_{ϵ} to generate the response.

[00172] Defending against adversarial strategy 2 – using secondary storage.

[00173] Assume that the adversary has modified some “ k ” words of data in RAM, and that an image of the “intended” contents of this area is kept in secondary storage. When F_{ϵ} would have attempted to access one of the modified cells, F' instead accesses the image in secondary storage.

[00174] This can be done in three ways: (a) The entire RAM can be imaged in secondary storage, and each access is made there; (b) Each memory access of F_{ϵ} is replaced with a comparison of the address to a range, with select access attempts only being rerouted to secondary storage; and (c) the Translation Lookaside Buffer (TLB) is used to automatically reroute accesses with selected ranges.

[00175] In the first case, each memory access will suffer the delay associated with accessing secondary storage instead of RAM. In the second case, only rerouted accesses suffer this delay, but for each memory access, a comparison has to be made. In the third case, there is only a delay imposed for rerouted accesses. However, due to the somewhat coarse granularity of the TLB, there is a minimum number of cells for which accesses can be rerouted.

[00176] **Defending against adversarial strategy 3 – using RAM.**

[00177] Due to the construction of the pseudorandom string that is used to fill free RAM, computing any one cell will require the computation of at least 512 hash function evaluations, since the value in each cell (in some embodiments) is a portion of the second-layer hash function output, whose input is a function of 512 first-level hash function evaluations.

[00178] Suppose there is a setup cost for each hash function computation, and that this is at least 1156 clock cycles. At the same time, the hash computation for MD6-512 – once the setup is completed – takes 155 cycles per byte, or 9920 cycles for a 64 bit input. Both of these assume optimized code and for a typical 32 bit processor. The total cost per hash function evaluation in this setting is therefore 11076 cycles. Computation of the contents of only one cell in RAM will take at least $(512+1)*11076$ cycles.

[00179] **Using Proxies For Timing**

[00180] As mentioned above, in some embodiments an external verifier is used to distribute new seeds and to time the arrival of values from timing checkpoints. To example models of communication between the audited device and the external verifier are: (a) a physical connection; and (b) a VPN between the external verifier and a SIM card. Data down is encrypted, data up is authenticated and/or bi-directional authentication occurs.

[00181] When a physical connection is not used between the device and the verifier, variances in latency can frustrate the timing of the computation. Accordingly, in some embodiments a proxy (with less latency variance) is used to support the timing. One example piece of hardware that can be used as a proxy is a SIM card.

[00182] A proxy can be used to reduce the variance associated with the starting of the computation. This applies both to the entire computational task, and to subtasks. Both of these can be initiated by sending a key or a seed that is needed for the computation, guaranteeing that the computation does not start until the key or seed is received.

[00183] The SIM card receives encrypted seeds, decrypts them, and provides device 902 with the values. This dispensal time can be made relative to other events that are observed on the device. Consider a triplet (location,data,seed), and assume the verifier produces such triplets, and sends them in an encrypted and authenticated fashion to the SIM card. Here, location describes a computational step (or stage) within a program, and data describes some state associated with this location. Assume that the computation can be anticipated by the external verifier, or that a small number of very likely computational paths can be anticipated.

[00184] This makes it possible for the external verifier to compute these triplets, predicting what data (or state) will be observable on the device at a given location (or stage). The third element of the triple, seed, denotes the value to be dispensed to the device if it reaches a given state, associated with location and data.

[00185] The SIM card will send the “value” location to the device, unless it can be computed or predicted by the device, or received from elsewhere. When the computation has reached a stage associated with the value location (this can be a loop iteration value), the software on the device sends the most recently computed value of some predetermined type – or parts thereof – to the SIM card. This can be the most recently computed value, the contents of a given register, or any other appropriate value. The SIM card compares this value to the value data associated with location, and if they are the same or equivalent, then responds with the value seed. The device replaces its current seed value with this new value for seed. This allows the device to replace the seed at a time when it has computed some value that vouches for it having reached a given computational stage. The new seed is not disclosed before this, which has security benefits. Alternatively, new seeds can be disclosed as soon as a new result is reported, independently of whether this reported value was correct or not; any reported value would be recorded, along with when it was received, allowing a verifier later to determine if incorrect results were received, or whether there were any noteworthy delays. The reports sent by the proxy can be compressed or digested in order to save bandwidth, as applicable.

[00186] The computation has reached a given checkpoint associated with a triple (location, data, seed) once a match is made and the associated seed value is dispensed.

[00187] In some embodiments the SIM card has a long vector of such triples, and selects either the first one (in the case of straight computation) in the line, or any one that matches the input (in the case where there are branches not foreseeable by the external verifier unit). It is possible to leave out the value location if a predictable location is used. The SIM card can obtain a long list of triples without it being possible for any of these to be intercepted by the handset or other parties; this is made possible by the use of point-to-point encryption between the external verifier (or associated proxy) and the SIM card.

[00188] It is also possible to use a verifier-initiated event instead of the value data. Thus, once that event is observed by the SIM card, the associated value seed is released, and the device computation is started with this value. This can be used to start the clocking. This can also be implemented by the device having to present a given piece of data that may, for example, be delivered by the external verifier to the handset in order to start the computation.

[00189] It is possible that some of the seed values are empty, in which case no seed value is sent to the device. This could be used to perform a check as to what point of the communication the handset has reached.

[00190] The SIM card does not have to be clocked (and therefore be aware of time), and it can operate in a hostile environment (as long as the values “data” are sufficiently long to preclude exhaustive search by malicious code on the handset, with the goal of obtaining seed values in a premature manner.

[00191] The end of a computation can be identified by having a checkpoint, as above, after which the SIM card reports the time of the computation to the external verifier, or reports information which allows the external verifier to compute an estimate of the time between the beginning and end of the computation. In the former case, a collection of local time values can be recorded by the SIM card, one for each checkpoint that is reached. This can be recorded as a list of pairs (location, time), where the value time can be a local time that does not need to be synchronized with the external verifier’s notion of time. Alternatively, pairs of reported values and the associated times can be recorded. The SIM card can authenticate this list and transmit it to the external verifier, who would verify that the authentication was correct, after which it can make a determination of how long various

computational tasks took. This can be used to infer the device's security posture. Alternatively, the SIM card can make this security determination and send an authenticated result of the determination to the external verifier.

[00192] Even without the ability to tell time, the SIM card can order events. As it receives packets from the verifier (or any entity collaborating with the verifier, and not with the malware), it will determine what device-reported event most recently took place. This corresponds to a value "data" that was received and verified as being correct. Thus, the values data can be generated by the device (given some computational task set up by or known by the external verifier) or can be a result of communication from the external verifier.

[00193] To deal with intentionally delayed reporting to the SIM card, it is possible to require immediate acknowledgment by the SIM card to the external verifier, once an event packet is received. These messages are authenticated and/or have a format that cannot be anticipated by a malware agent. An example way to construct this by letting the value "data" above correspond to the event from the external verifier, and let the device report the associated value "seed" to the external verifier once it is received from the SIM card. The external verifier can then time the arrival of this acknowledgment after a preceding packet is sent. Alternatively, the device can be allowed to report values corresponding to "data" directly to the external verifier. This can be done for some portion of all the checkpoints, or for additional checkpoints.

[00194] Assume that a packet identified by a string S was received by the SIM card after an event E took place, but before any other event took place (and was correctly verified by the SIM card). Then, the pair (S,E) is added to a log L. Alternatively, identifying portions of the strings S and E are added to the log. The log is communicated to the verifier at the end of the communication. Alternatively, ordering information is communicated to the external verifier.

[00195] In some embodiments, triplets (location,data,report) are added to the vector, where "report" is a value that signifies a report being sent to the verifier. This can also be achieved using regular seed values, where the last seed value released is a value that the device communicates to the verifier in order to stop the timing. The timing can also be stopped by having the device perform some additional keyed task and report back values in

exchange for seeds, in which case the arrival of the next packet with a string S would identify when the timing stopped (extrapolated from the number of cycles that apparently were computed after the intended computation ended.)

[00196] Once the timing has ended, the SIM card encrypts (and potentially authenticates) the log, and passes this to the device for transmission to the verifier. The verifier determines whether it was correctly authenticated, decrypts it, and then determines from the log what the completion times of partial steps of the communication were, given information about the transmission times of the packets containing values such as S. It may know the times when these were transmitted by the base station closest to the handset, or may know when they were processed by a network proxy, or by the originator of the packet itself.

[00197] In some embodiments the values S are generated by a trusted beacon and potentially not known by the external verifier by the time it starts the computation on the device.

[00198] Some types of packets, like the packets containing S, need not be encrypted before they are transmitted over the network. In some embodiments they are authenticated. However, the verification by the SIM card of whether a packet has a correct authentication need not be done before a preliminary log entry is compiled. If the authentication verification fails for an already compiled log entry, this log entry can be erased.

[00199] SIM cards are half duplexed, i.e., cannot receive and send data at same time. SIM cards operate as slaves, i.e., will only (with some special exceptions, such as when they are powered up) send data to the attached device (our handset) after having been requested to do so. However, some smart cards can operate independently between queries from their associated handsets, and not only as a reaction to being queried.

[00200] If a SIM card supports multi-threading, it is possible to let one thread perform a simple count (started when the timing starts), and to provide this counter to the other thread, which records this value each time a correct data value is received. The counter can be stored along with the data value, the associated location, or an index that indicates what data value is associated with the counter. In some situations, such as if it is guaranteed that one data value cannot be accepted more than once, and there is a deterministic order of values to be received, then we may record the counter values, and not the data values or other state information.

[00201] Some SIM cards, such as typical Java Cards, only support computation after a message is received from the SIM card interfacing device (CAD). The computation typically ends when the response from the SIM card is generated and transmitted to the CAD.

[00202] If a SIM card allows a counter to be increased for each clock cycle (or another deterministic periodicity) even before messages received from the handset, and after responses being transmitted, then it is possible to maintain the proper count, even without support for multi-threading.

[00203] It is possible to keep time-based state in the SIM card; it is also possible for the SIM card to authenticate events (including the time at which they took place) and export lists of such authenticated events.

[00204] Figure 14 illustrates an embodiment of an implementation of a process for timing a portion of a device audit. In the example shown, a modified Java Card is used as a proxy. The modification permits the process to remain active after responding to a request. The proxy receives a vector of values input and output from the external verifier, and produces a vector duration that is transmitted to the external verifier at the completion of the execution. (All the communication between the proxy and the external verifier is assumed to be authenticated and encrypted.) In the example shown, the value “∞” corresponds to an error message indicative of an attempt of a malware agent on the client to cheat. Upon receipt, the external verifier will determine if the values in the vector duration all fall within the tight bounds suggesting successful completion, and will conclude that the client is in a safe state only if that holds.

[00205] Additional Ways Of Using Proxies For Timing

[00206] In addition to using a SIM or similar piece of hardware as a proxy, it is also possible to use as a proxy another device that is believed to be in a secure state. It is possible to bootstrap security by first having one device (such as a phone) be verified as being secure, and then use that device to perform local timing and other verification tasks in the process of verifying the security of a second device.

[00207] The devices can be of different types. For example, one may use a SIM card, local cell phone tower, or a local computer to assist the security verification of a handset, on the request of an external verifier. It is possible for this external verifier to have prepared

parts of the data used in the verification, such as the pairs of (data,seed) described above. It is also possible for a third party to perform this pre-computation. In various embodiments, a seed is generated by the proxy. One example is for the proxy and the device to both contain a sensor such as an accelerometer or photovoltaic sensor. Both the device and the proxy observe the same phenomenon (e.g., by being held together and shaken together) and calculate a seed the manner. In this case, the device uses the observed seed, and the proxy sends the seed (which it independently experienced) to the external verifier.

[00208] As soon as a first device is determined to be in a secure state, it can be used to time or otherwise assist in the security assessment of a second device, such as the infotainment system of a car, another handset, a netbook, a laptop or desktop, or other device. This security assessment can be of the same type (e.g., based on the time to perform computational tasks), or it can be an alternative security assessment method that is bootstrapped on the first. Likewise, the first security assessment can be of a different type, while later security assessment can be bootstrapped on the first one, and may use our timing-based approach.

[00209] In some embodiments, a first known “safe” device is used to produce a collection of sets of seed values, to be consumed by other devices at a later point. It is possible to have such sets certified (e.g., using a PKI or by peer-based authentication of values). Further, it is possible for the external verifier to operate in an off-line manner if it has the support of an online proxy. For example, the external verifier can send encrypted and authenticated data at a time far ahead of the auditing process; several such transcripts can be sent at a time. They can either be sent to the proxy, where they can be kept; or they can be sent to the audited device, or a proxy thereof, where they would be kept until needed. The transcripts that are generated by the proxy as a result of the memory-printing may also be buffered by a device and later sent to the external verifier, whether when they are requested, or when there is an availability of a communication channel. In some embodiments all records are marked up with time-stamps and serial numbers before they are authenticated and possibly encrypted.

[00210] It is possible to implement this in a network of small nodes, wherein some of the nodes either are trusted a priori, or are assessed to be secure; after which these nodes are used to assist the security assessment of other nodes. This is a potentially recursive approach, and can be circular, i.e., a previous trusted device used to assess the security of other devices

may later on be verified by some of these devices, or devices assessed to be secure in other ways. The external verifier can still be included in the environment, and may help start up the chain of verification events, and help schedule what nodes should be verified when, and by whom.

[00211] Using Compressed Access Tables

[00212] In some embodiments the position of a memory access is determined by the contents of a vector “location” whose contents correspond to a permutation of all cells of free RAM. This vector can take up all of free RAM, if kept there. In some embodiments it is stored in secondary storage (such as flash), and portions are swapped in as needed. An alternative approach that maintains a pseudorandom access order, but which minimizes the computational effort during the main loop will now be described.

[00213] Consider two vectors, locationH and locationL, where both are vectors, each one containing a permutation of partial memory access positions. Here, the actual memory access position is the combination of two partial positions, e.g., the bits of one locationH element concatenated with the bits of one locationL element. Here, the locationH element is assumed to contain the higher order bits and the locationL element the lower order bits. These elements can be of the same or different size, but will, when combined, be the size that addresses one memory location. If each contains all possible elements in the range, then the collection of all combinations will correspond to a collection of all memory addresses. (From this, one can remove those that are not in free RAM by comparing the combined result to a threshold, and trash the result if it falls below this threshold.) This representation takes only the square root of the size of the space addressed to store. It is possible to use three components, in which case they take the third root of the space addressed. It is possible to use a large number of components as well. One example combination function is concatenation.

[00214] In some embodiments the order of access of the elements of the vector follows a geometric pattern that guarantees that all combinations are going to be used with an overwhelming likelihood. It can be beneficial not to have several accesses of one and the same item within one vector, as this reduces the degree of unpredictability to an adversary, given the increased patterns in memory access. It is possible to cover one combination more than once, although it is beneficial to limit the total number of accesses to memory at the same time as guaranteeing that all accesses are made, but for a negligible probability.

[00215] It is possible to access the locationH vector at a position x and the locationL vector at a position y, and to access the x-y positions along the diagonals. Here, the first sequence can be started at position (x,y)=(0,0), after which x and y are both and simultaneously increased by one for each iteration of the loop. When one coordinate is increased beyond the size of the vector, the coordinate is set to 0 again. Then, when the position again becomes (0,0), it can be modified to start at position (x,y)=(1,0), after which the sequence of increments is repeated until it comes back to (1,0), at which time it is changed to (2,0). This is not the location of the memory access: it is the position in the vectors that describe where to make memory accesses.

[00216] It is also possible to otherwise compress the description of what cells to access by having a vector of locations elements, where each such location only describes part of an address, and the remaining bits of the address are computed in another fashion, or inferred from program state at the time of the computation. Moreover, these two approaches can be combined, and combined with yet other related descriptions of access locations that are at least partially pregenerated.

[00217] Additional Information On Timing

[00218] In various computations, timing of computation, and partial computation occurs as follows. (A) The timer is started once the auditor is provided with all necessary values, whether from the external verifier or a proxy thereof. These values typically include the value seed. (B) The timer is stopped (and the time since it was started recorded) when the auditor submits a correct value “state” to the external verifier or a proxy thereof.

[00219] It is possible to immediately start a new time interval when an old one has ended (where the start is signified by step A and the end is signified by step B above). It is also possible to implement “recesses” between these intervals; during these recesses, the computation may not be timed, and the algorithm may perform routine maintenance, such as communicating with external parties, reading or writing to secondary storage, or other functionality. The recess can be ended when the algorithm requests that the next timing interval is started (e.g., step A); one way this can be done is by signaling to the external verifier or proxy thereof to start the next interval; or it can be done by the external verifier or proxy thereof selecting to start the new interval.

[00220] It is also possible to implement recesses as standard timing intervals, whose length is not critical to the final determination of the security posture of the audited device.

[00221] Pseudorandom Access

[00222] In some embodiments, the selective reading performed as part of the audit process is accomplished through access in a pseudo-random order, with a sequence of reads and writes to the accessed positions. An alternate embodiment using pseudorandom access now be described. First, a description of an example of memory filling will be provided. Then a description of an example of periodic timing will be provided.

[00223] Filling Fast Memory

[00224] The following memory-printing function can be used to fill free RAM. It can also be used to fill other types of fast memory, in the event that such other types of memory are comparable with RAM in terms of access times. A pseudo-random sequence is XORed in to free RAM in a pseudo-random order; later, a keyed hash of the entire contents of RAM is computed. Even though RAM does not use blocks and pages, it can nonetheless be divided into “virtual” blocks and pages, corresponding to those of flash. Consecutive chunks of flash are not accessed in a page or block. This makes the access slow in flash, but still fast in RAM.

[00225] In order to fill free RAM with a pseudo-random string, there are two main steps. First, a setup function is run. This determines the random order of memory accesses to be made by the memory-printing function, using a seed obtained from the verifier to generate pseudorandom values. The table is stored in flash, and the program space used by the setup function is cleared after the setup completes. Second, a memory-printing function is used to fill all free RAM. Its execution is timed, both from beginning to end and in shorter intervals.

[00226] Handling Network Delays

[00227] Delays caused by infection can be measured from a device connected to the client device by internal wiring; standard network port, such as USB; over a wired interface; over a WiFi network; over a LAN; over the Internet; over a packet-switched network; over a communication network; or a combination of these. Some of these communication media

may introduce delays and variance, which can be separated from the measurement using statistical methods.

[00228] The verification is made by a device that is connected to the audited device using a cable, a LAN, a WAN, Bluetooth, Wifi, the Internet, another network, or a combination of networks. The verification is made by comparing a received result with a computed result, and to verify that it (and the sequence before it) was received within the proper time bounds. All of these communication media may incur latencies, and some may drop packets.

[00229] Assume for a moment that a “good” event takes 10 units of time, plus between 1 and 5 (for typical network variance).

[00230] Then, assume that a “bad” event takes 15 units of time, plus 1 to 5 for network variance.

[00231] Consider the receipt of partial results at these times:

[00232] Sequence a: 0, 12, 25, (missing packet), 50 – this sequence is likely to be good, in spite of the missing packet, since the last partial result “vouches for” the lost packet.

[00233] Sequence b: 0, 11, 30, 35, 50 – this sequence is likely to be good, in spite of the long delay between the second and third packet, since the fourth packet was received “too early”.

[00234] Sequence c: 0, 11, 30, 45, 57 – this sequence is likely to be bad due to the long delay after the second packet, and no event that explains the delay.

[00235] Although the foregoing embodiments have been described in some detail for purposes of clarity of understanding, the invention is not limited to the details provided. There are many alternative ways of implementing the invention. The disclosed embodiments are illustrative and not restrictive.

[00236] WHAT IS CLAIMED IS:

CLAIMS

1. A system, comprising:
 - a physical memory; and
 - a processor configured to:
 - 5 receive one or more hardware parameters that correspond to a hardware configuration and receive initialization information;
 - selectively read the physical memory and determine at least one result; and
 - provide the result to a verifier.
2. The system of claim 1 wherein the initialization information comprises a seed value.
- 10 3. The system of claim 1 wherein the initialization information comprises a step value.
4. The system of claim 1 wherein the initialization information is received from a proxy.
5. The system of claim 1 wherein the initialization information is received from an external verifier.
- 15 6. The system of claim 1 wherein the initialization information is a representation of a phenomenon.
7. The system of claim 6 wherein the initialization information is received from an accelerometer included in the system.
8. The system of claim 6 wherein the processor is configured to receive the initialization information at least in part by measuring a phenomenon.
- 20 9. The system of claim 1 wherein the processor is further configured to determine a portion of the physical memory to be overwritten in accordance with a function.
10. The system of claim 1 wherein the processor is further configured to write a portion of the physical memory in accordance with a function.
11. The system of claim 1 wherein the processor is configured to selectively read the physical memory based at least in part on a step counter.
- 25 12. The system of claim 11 wherein the step counter is relatively prime to a measurement of the size of the physical memory.

13. The system of claim 11 further comprising a plurality of processors and wherein the step counter is relatively prime to a measurement of the size of the physical memory multiplied by the number of processors.

14. The system of claim 1 wherein the processor is configured to determine the result at least in part by using a function that is not the identity function and takes as input a portion of the physical memory.

15. The system of claim 1 wherein the processor is configured to determine the result at least in part by using an accumulator.

16. The system of claim 15 wherein the accumulator includes a nonlinear function.

17. The system of claim 15 wherein the accumulator uses an XOR and a rotate.

18. The system of claim 1 wherein the processor is configured to provide the result to the verifier at least in part by providing information to a proxy.

19. A method, comprising:

receiving one or more hardware parameters that correspond to a hardware

15 configuration and receiving initialization information;

selectively reading a physical memory and determining at least one result; and

providing the result to a verifier.

20. A computer program product embodied in a computer readable storage medium and comprising computer instructions for:

receiving one or more hardware parameters that correspond to a hardware

configuration and receiving initialization information;

selectively reading a physical memory and determining at least one result; and

providing the result to a verifier.

21. A system, comprising:

25 a physical memory; and

a processor configured to:

receive one or more hardware parameters that correspond to a hardware configuration;

perform a sequence of modifications to the physical memory; and

30 provide results to a verifier.

22. The system of claim 21 wherein the verifier is configured to determine, based at least in part on the received results, whether an expected sequence of physical memory modifications was made by the processor.
23. The system of claim 22 wherein the determination is based at least in part on timing information associated with the performance of the sequence of modifications.
24. The system of claim 22 wherein the determination is based at least in part on the correctness of a result received from the device.
25. The system of claim 21 wherein the verifier is collocated on a device with the processor.
- 10 26. The system of claim 25 wherein the verifier is included in a subscriber identity module.
27. The system of claim 21 wherein the verifier is in communication with the processor via a network connection.
- 15 28. The system of claim 21 wherein the processor is further configured to obtain a seed from the verifier.
29. The system of claim 21 wherein the sequence of modifications depends at least in part on an input obtained from the verifier.
30. The system of claim 21 wherein the one or more hardware parameters is an amount of memory.
- 20 31. The system of claim 21 wherein the one or more hardware parameters is a memory speed.
32. The system of claim 21 wherein the processor is configured to perform a scan once it is determined that no evasive software is active in the physical memory.
33. The system of claim 32 wherein the scan comprises a scan for unlicensed software.
- 25 34. The system of claim 21 wherein the memory includes secret data and wherein providing results to the verifier does not compromise the secret data.
35. The system of claim 21 wherein the processor is included in a mobile phone.
36. A method, comprising:

receiving one or more hardware parameters that correspond to a hardware configuration;

performing a sequence of modifications to the physical memory; and
providing the results to a verifier.

5 37. The method of claim 36 wherein the verifier is configured to determine, based at least in part on the received results, whether an expected sequence of physical memory modifications was made.

38. A computer program product embodied in a computer readable storage medium and comprising computer instructions for:

10 receiving one or more hardware parameters that correspond to a hardware configuration;

performing a sequence of modifications to the physical memory; and
providing the results to a verifier.

39. A system, comprising:

15 a processor configured to:

receive from a device one or more results associated with the performance by the device of a sequence of modifications to a memory included on the device;

determine that the one or more results indicates that the memory included on the device is free of evasive software; and

20 initiate a scan of the device after the determination is made;

a memory coupled to the processor and configured to provide the processor with instructions.

40. A system, comprising:

a processor configured to:


25 initiate a request that a device determine a result;


receive the result from the device; and


determine, based at least in part on the result and a length of time that it takes the device to provide the result, a security state of the device; and


a memory coupled to the processor and configured to provide the processor with



30 instructions.


**FIG. 1**

**FIG. 2**

**FIG. 3**


**FIG. 4**


**FIG. 5A****FIG. 5B**

**FIG. 6**

```
for i:=1 to rounds
  for k:=0 to chunks_per_block -1
    permutation[1 .. number_blocks] := get_permutation
    for j:= 1 to number_blocks
      modify_memory(permutation[j]*chunks_per_block+k, next_string_chunk)
```

FIG. 7

**FIG. 8**

**FIG. 9**




FIG. 10




FIG. 11




FIG. 12

```

state ← 0
location ← 0
repeat ( $s_{RAM}/period/\alpha$ ) times:
    receive value key
    state ← state XOR key
    repeat period times:
        data ← RAM[location]
        state ← ROR(state XOR data)
        data ← RAM[location + 1]
        state ← ROR(state XOR data)
        ...
        data ← RAM[location +  $\alpha - 1$ ]
        state ← ROR(state XOR data)
        location ← (location + step) mod sRAM % next location?
    report value state

```

FIG. 13

```
counter ← 0
time ← 0
repeat period times:
    counter ← counter + 1
    repeat until value state is received:
        time ← time + 1
        if time = maxtime then
            duration(counter) ← ⊥
        counter ← counter + 1
        if state = input(counter) then
            duration(counter) ← time
            respond with value output(counter)
        else
            duration(counter) ← ⊥
```

FIG. 14

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 10/02214

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G06F 17/30, 21/00, 9/455 (2010.01)
USPC - 726/22; 703/13

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC(8): G06F 17/30, 21/00, 9/455 (2010.01)

USPC: 726/22; 703/13

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 703/23, 707/E17.01, 707/E17.032 (keyword limited; terms below)Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
PubWest; Google Scholar; Google Patents; FreePatentsOnline. Search terms used:
malware malicious-software unauthorized-program blacklist, memory, processor microprocessor CPU SPU, hardware-parameter hardware-configuration, initialize, modify change update revise, verify, seed, step ...

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2008/0034406 A1 (GINTER et al.) 07 February 2008 (07.02.2008) entire document, especially Abstract; para [0063], [0157], [0158], [0165], [0166], [0257], [0419], [0422], [0453], [0473], [0490], [0512], [0520], [0535], [0587], [0688], [0827], [0940], [0944], [0955], [1013], [1063], [1092], [1130], [1516], [1629], [1847], [2042], [2112], [2135], [2153]	39, 40 ----- 1 - 38
Y	US 2008/0229415 A1 (KAPOOR et al.) 18 September 2008 (18.09.2008) entire document, especially Abstract; para [0013], [0063], [0129], [0144], [0153], [0257], [0297], [0436], [0455]	1 - 38
Y	US 2005/0097094 A1 (SILVERBROOK et al.) 05 May 2005 (05.05.2005) entire document, especially Abstract; para [0009], [0499]-[0502]	7
Y	US 2005/0188218 A1 (WALMSLEY et al.) 25 August 2005 (25.08.2005) entire document, especially Abstract; para [0045]-[0047], [5179], [5186], [5191], [7774]	15-17
Y	US 7,315,826 B1 (GUHEEN et al.) 01 January 2008 (01.01.2008) entire document, especially Abstract; para col 2, ln 3-12; col 161, ln 44-45	33
A	US 2009/0109959 A1 (ELLIOTT et al.) 30 April 2009 (30.04.2009) entire document	1 - 40

 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	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"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
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 11 October 2010 (11.10.2010)	Date of mailing of the international search report 15 OCT 2010
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774