



- (51) International Patent Classification:
G06F 21/31 (2013.01)
- (21) International Application Number:
PCT/US2014/029444
- (22) International Filing Date:
14 March 2014 (14.03.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
61/792,313 15 March 2013 (15.03.2013) US
- (71) Applicant: **POWER FINGERPRINTING INC.**
[US/US]; 2200 Kraft Drive, Suite 1200 R, Blacksburg,
Virginia 24060 (US).
- (72) Inventors: **AGUAYO GONZALEZ, Carlos R.**; 1228
Wild Hawthorn Way, Reston, Virginia 20194 (US).
REED, Jeffrey H.; 1105 Eheart Street, Blacksburg, Vir-
ginia 24060 (US). **CHEN, Steven C.**; 9844 Avenel Farm
Drive, Potomac, Maryland 20854 (US).
- (74) Agents: **HUTTER, Christopher R.** et al.; Cooley LLP,
1299 Pennsylvania Avenue, NW Suite 700, Washington,
District of Columbia 20004-2400 (US).
- (81) Designated States (*unless otherwise indicated, for every
kind of national protection available*): AE, AG, AL, AM,

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

- (84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

(54) Title: SYSTEMS, METHODS, AND APPARATUS TO ENHANCE THE INTEGRITY ASSESSMENT WHEN USING POWER FINGERPRINTING SYSTEMS FOR COMPUTER-BASED SYSTEMS




Figure 42B

(57) Abstract: A power fingerprinting system is adopted for assessing integrity of a target computer-based system. In one implementation, the power fingerprinting system may receive, at a first module, side-channel information of a first target component of a system, the first module being collocated with the first target component; obtain a power fingerprint for the first target component based on the side-channel information for the first target component, the power fingerprint for the first target component representing a plurality of execution statuses of the first target component; receive, at a second module, side-channel information of a second target component of the system, the second module being collocated with the second target component, the power fingerprint for the second target component representing a plurality of execution statuses of the second target component; and obtain a power fingerprint for the second target component based on the side-channel information for the second target component.

WO 2014/144857 A2

**SYSTEMS, METHODS, AND APPARATUS TO ENHANCE THE
INTEGRITY ASSESSMENT WHEN USING POWER FINGERPRINTING
SYSTEMS FOR COMPUTER-BASED SYSTEMS**

Cross References to Related Applications

[1001] This application is a non-provisional of and claims priority under 35 U.S.C. §119 to U.S. provisional application serial no. 61/792,313, filed March 15, 2013, titled “Systems, Methods, And Apparatus To Enhance The Integrity Assessment When Using Power Fingerprinting Systems For Computer-Based Systems.”

[1002] This application is related to Patent Cooperation Treaty international application serial no. PCT/US11/59244, filed March 11, 2011, titled “Using Power Fingerprinting (PFP) To Monitor The Integrity And Enhance Security Of Computer.”

[1003] The aforementioned applications are all herein expressly incorporated by reference.

Background

[1004] A processor-based device or system uses electrical power for data exchange, storage, and computation at circuit units. For example, in Complementary Metal-Oxide-Semiconductor (CMOS) digital circuits, with every bit transition, a transient current drain results from a brief short circuit in the gates and the charge and discharge of parasitic capacitance at the outputs of the circuit. Hence, the total power consumed in a specific clock cycle, is determined by the total number of bit transitions that take place in that cycle. The number of bit transitions is determined by the specific instruction sequence executed, as well as their addresses and parameters. When a disruption of the execution of the specific instruction sequence occurred, such disruption is reflected in the transient currents at the digital circuit level, and thus the consumption of the electric power.

Summary

[1005] In some embodiments, a power fingerprinting (PFP) system provides an integrity assessment and intrusion detection solution for critical cyber systems based on taking fine-grained measurement of a processor's side channels (e.g. power consumption) and comparing them against trusted baseline references such as power signatures or fingerprints (e.g., patterns that result from the specific sequence of bit transitions during execution) for anomaly

detection. In one implementation, the PFP system may characterize the execution of trusted software and extract its power signatures and use them as reference to compare test traces to determine whether the same code is being executed. In this way, the PFP system may obtain information to the execution status of a target device and may detect any anomaly of the target device.

[1006] In some embodiments, the PFP system may perform a method, comprising: receiving, at a first module, side-channel information of a first target component of a system, the first module being collocated with the first target component, the side-channel information of the first target component being associated with the plurality of authorized execution statuses of the first target component and an execution status of the first target component; obtaining a power fingerprint for the first target component based on the side-channel information for the first target component, the power fingerprint for the first target component representing a plurality of authorized execution statuses of the first target component; receiving, at a second module, side-channel information of a second target component of the system, the second module being collocated with the second target component, the side-channel information of the second target component being associated with the plurality of authorized execution statuses of the second target component and an execution status of the second target component; obtaining a power fingerprint for the second target component based on the side-channel information for the second target component, the power fingerprint for the second target component representing a plurality of execution statuses of the second target component; and sending, from a processor module, a reporting signal based on at least one of (1) the power fingerprint for the first target component and an execution status of the first target component, or (2) the power fingerprint for the second target component and an execution status of the second target component, the reporting signal associated with at least one of the execution status of the first target component or the execution status of the second target component.

[1007] In some embodiments, the PFP system may perform a method, comprising: sending a predefined input to an untrusted electronic device, the predefined input being defined based on a functionality of the untrusted electronic device; receiving side-channel information of the untrusted electronic device in response to the predefined input; obtaining a power fingerprint for the untrusted electronic device based on the side-channel information; and assessing an integrity of the untrusted electronic device based on the power fingerprint.

[1008] In some embodiments, the PFP system may comprise: a power fingerprint monitor module configured to receive a signal representing an unauthorized access of a device; and a response analysis module operatively coupled to the power fingerprint monitor module, the response analysis module configured to select a response module from a plurality of response modules in response to detection of the unauthorized access, a first response module from the plurality of response modules configured to collect information associated with the device during the unauthorized access, a second response module from the plurality of response modules configured to prevent access to the device in response to the unauthorized access, a third response module from the plurality of response modules configured to disable at least a portion of the device subjected to the unauthorized access.

[1009] *Brief Description of the Drawings*

[1010] Figure 1 is a general block diagram of power fingerprinting.

[1011] Figure 2 is a diagram showing ideal sensor locations for a PFP monitor.

[1012] Figure 3 is a diagram showing ideal sensor locations for multi-processor boards.

[1013] Figure 4 is a diagram showing an example of triggering with a physical signal.

[1014] Figure 5 is a schematic diagram showing PFP strategic instruction insertion for synchronization and triggering.

[1015] Figure 6 is a schematic diagram showing indirect access to physical resources in the Linux device driver paradigm.

[1016] Figure 7 is a flow chart showing the trusted code characterization process.

[1017] Figure 8 is a graph showing an example of preprocessing of traces by calculating their power spectral density.

[1018] Figure 9 is a graph showing sample preprocessing of traces in the time domain .

[1019] Figure 10 is a graph showing PSD difference from the execution of test traces against, a stored signature.

[1020] Figure 11 is a flow chart showing the process of detector design.

[1021] Figure 12 is a graph showing sample probability distribution from trusted code execution used for detector design and threshold selection.

[1022] Figure 13 is a flow chart showing the PFP integrity assessment operation process.

[1023] Figure 14 is a schematic diagram showing a sample embodiment setup for the Android platform.

[1024] Figure 15 is a graphical representation of sample distributions resulting from execution of the original untampered routine.

[1025] Figure 16 is a graphical representation of sample distributions resulting from execution of the tampered routine.

[1026] Figure 17 is a graph of sample trace detail showing different sections of the traces containing different levels of discriminatory information.

[1027] Figure 18 is a schematic representation of a platform characterization using a linear projection from the most informative perspective,

[1028] Figure 19 is a schematic diagram showing a reference measurement setup for platform power consumption characterization and monitoring using PFP.

[1029] Figure 20 is a graph showing a sample trace from baseline code execution to evaluate ability to detect minimum power consumption change.

[1030] Figure 21 is a graphical representation of an average signature from, baseline code execution, where each point represents a projection on an n-dimensional Euclidean space.

[1031] Figure 22 is a graph showing sample distribution of Euclidean distances from the average signature extracted from the execution of the baseline code.

[1032] Figure 23 is a graph showing sample distribution of Euclidean distances from the baseline signature in the transformed space obtained using PCA.

[1033] Figure 24 is a graph showing centroids of traces from profiling instructions for LDA.

[1034] Figure 25 is a graph showing sample distribution of Euclidean distances from the baseline signature in the transformed space obtained using LDA.

[1035] Figure 26 is schematic block diagram of an exemplar target platform for detecting deviations from authorized software execution.

[1036] Figure 27 is schematic diagram showing different layers in a defense-in-depth approach to cyber security.

[1037] Figure 28 is a schematic diagram showing PFP monitoring scope within a layered defense-in-depth security solution,

[1038] Figure 29 is a schematic diagram showing PFP Honeypot node for monitoring and intelligence gathering.

[1039] Figure 30 is a schematic diagram showing a PFP Honeypot network.

[1040] Figure 31 is a flow diagram of a supply-chain trust analysis using PFP.

[1041] Figure 32 is a schematic diagram showing potential sources of reference signatures for supply-chain trust analysis using PFP.

[1042] Figure 33 is a diagram showing use of an I/O register to provide synchronization and identification signaling to a PFP monitor.

[1043] Figure 34 is a diagram showing embedding of PFP synchronization and identification signaling into power traces.

[1044] Figure 35 is a diagram showing a sample setup for combining multiple signals for PFP integrity assessment.

[1045] Figure 36 is a flow chart of a process to extract features that, preclude device failure for PFP.

[1046] Figure 37 is a schematic diagram showing relationships between the different system elements interacting for automatic characterization and feature extraction.

[1047] Figure 38 is a diagram of a structure for preventing side-channel attacks from exploiting an embedded PFP monitor.

[1048] Figure 39 is a diagram of a structure for preventing side-channel attacks from exploiting an external PFP monitor.

[1049] Figure 40 provides an example logic flow diagram illustrating aspects of a cycle-accurate power consumption simulator estimating the power fingerprints of a target device, according to an embodiment.

[1050] Figures 41A-B provide data plots illustrating example sample power fingerprinting distributions showing outliers or anomalies with different underlying distributions.

[1051] Figures 42A-B provide example logic flow diagrams illustrating aspects of reverse engineering to validate fingerprints from hardware, according to an embodiment.

[1052] Figures 43A-C provide example logic flow diagrams illustrating aspects of enhanced response to integrity violations, according to an embodiment.

[1053] Figures 44A provides an example of logic flow diagrams illustrating aspects of processing integrity violations event, according to an embodiment.

[1054] Figure 44B provides an example diagram illustrating aspects of temporarily disabling a target device upon detecting anomalies, according to an embodiment.

[1055] Figure 45A provides an example diagram illustrating aspects of permanently disabling a target device upon detecting anomalies, according to an embodiment.

[1056] Figure 45B provide example logic flow diagrams illustrating aspects of using PFP integrity assessment as a gate keeper before execution external actions, according to an embodiment.

[1057] Figures 46A-C provide example logic flow diagrams illustrating aspects of Supervisory, Control, and Data Acquisition (SCADA) systems, according to an embodiment.

[1058] Figure 46D provides an example diagram illustrating aspects of an embedded board interfacing with a PFP monitor, according to an embodiment.

[1059] Figure 47 provides an example block diagram illustrating aspects of integrating PFP into a device charger or sleeve, according to an embodiment.

Detailed Description

[1060] Systems, methods and apparatus disclosed herein relate to integrity assessment and intrusion detection approaches that can detect malicious intrusions, unauthorized modifications, and tampering in digital circuits and computer-based systems, including critical embedded systems, coprocessors, and field-programmable gate arrays (FPGAs). A physical side channel (e.g. power consumption) can be used to observe the execution status in a processor or digital circuit using an external monitor and detect, with extreme accuracy, when an unauthorized execution has managed to disrupt the normal operation of digital systems. This is referred to herein as Power Fingerprinting (PFP); methods and apparatus that implement PFP can be referred to as a PFP system, which can be embodied in a number of different ways and forms. Physical characteristics of a processor, such as power consumption that can be measured externally from a processor and contain information related to the execution status of the processor, are referred to herein as side channels.

[1061] In PFP integrity assessment, an independent monitor can observe a side channel, such as the power consumption of the target system, during operation using a physical sensor. The captured traces can be processed to extract unique fine-grained patterns or features (fingerprints) and compared against trusted stored baseline references, which are used to perform anomaly detection by applying signal detection and pattern recognition techniques. Because the actual monitoring is performed by an external device, the memory and processing overhead on the target systems is greatly reduced.

[1062] This disclosure discusses embodiments of the PFP monitoring system, including example circuit deployment of a PFP system (e.g., see Figures 2-6), statistical comparison of PFP patterns (e.g., see Figures 7-13), and various example PFP platform integration and performance analytics (e.g., see Figures 14-39). This disclosure further discusses embodiments of enhanced integrity assessment using PFP, including using cycle-accurate power consumption simulator to generate estimated power fingerprints (e.g., see Figures 40-41), various circuit deployment of enhanced PFP systems and work flows for detecting anomalies with a target device and remedial measures to handle a detected anomaly (e.g., see Figures 42-45), example system integration with SCADA systems (e.g., see Figures 46A-D), and example apparatuses of PFP monitor integration with handheld devices (e.g., see Figure 47).

[1063] Various PFP modules and/or components for PFP sensors, verification, response analysis, response execution, statistical pattern analysis and/or the like, and/or other modules may be employed. As used in this specification, a module and/or component can be, for example, any assembly and/or set of operatively-coupled electrical components associated with performing a specific function, and can include, for example, a memory, a processor, electrical traces, optical connectors, software (stored in memory and/or executing in hardware) and/or the like.

Power Fingerprinting Monitoring

[1064] Some embodiments provide procedures for enhancing target system execution integrity assessment determined by power fingerprinting (PFP): by integrating PFP into the detection phase of comprehensive defense-in-depth security; by deploying a network of PFP enabled nodes; by executing untrusted devices with predefined inputs forcing a specific state sequence and specific software execution; by embedding module identification information into synchronization signaling; by combining signals from different board elements; by using malware signatures to enhance PFP performance; by automatic characterization and signature extraction; by providing secure signature updates; by protecting against side-channel attacks; performing real-time integrity assessment in embedded platform by monitoring their dynamic power consumption and comparing it against signatures from trusted code; by pre-characterizing power consumption of the platform and concentrating on trace sections carrying the most information about the internal execution status; by improving the ability of PFP to detect deviations from authorized execution in commercial embedded platforms.

[1065] A method performs real-time integrity assessment of execution of a routine in a computer processing platform. This is accomplished by monitoring execution of the routine by tracing power consumption of a processor, by sampling the processor during execution of the routine. A platform characterization technique is employed that detects sections of the traces, that is, those sections that display the largest dependence on state transitions in the processor. These sections are used to select features carrying the most information. This platform characterization applies to the platform and can be used for all routines run on the platform. The next step is to obtain, from a characterization of selected features of the routine, as contained in the sections identified in the platform characterization, a set of trusted power fingerprints of the routine. Then there is established a threshold for a specific false alarm rate based on the probability distribution of distance from a signature comprised of the

trusted fingerprints. A library of the trusted fingerprints is then compared to features extracted from traces from the execution of untrusted code, and there is then determined a distance between the fingerprints in the library and the features extracted from execution of the untrusted code. An exception is reported if the distance exceeds the threshold.

[1066] Various procedures for improving the operation, effectiveness, usability, and performance of integrity assessment and intrusion detection systems based on power fingerprinting (PFP) are described. The different procedures include: embedding module identification information into synchronization signaling; improved PFP monitoring by combining signals from, different board elements; using malware signatures to enhance PFP performance, generalizing on existing battery monitoring technology; automatic characterization and signature extraction; secure signature updates; response to integrity violations and layered security; protection against side-channel attacks.

[1067] Also described are methods and apparatus for: distributed PFP monitor network to monitor malware dynamics and behavior; application of PFP to Supply-Chain Trust Analysis; digital rights management and execution limited leases; failure prediction based on PFP.

[1068] Power fingerprinting (PFP) is a technique that enables an external monitor to assess the execution integrity of a cyber system. PFP relies on the execution status information carried by the dynamic power consumption of a processor. Using this information, along with pre-characterized signatures from trusted references, PFP is able to determine execution integrity in target systems. For practical application of PFP, it is necessary to implement specific apparatus and follow specific procedures to provide an effective monitoring solution. In this work, we describe various procedures for improving the operation, effectiveness, usability, and performance of a PFP monitoring solution.

Applying PFP to detect software modifications in smart phones and other embedded devices,

[1069] Cyber security has become a critical element for national security. Microprocessors are ubiquitous in almost every aspect of modern life. Technology developments in the information technology areas are moving forward at a faster pace than the security solutions necessary to protect them. The threat of cyber attacks remains constant with potential devastating consequences to critical infrastructure and national security. Cyber infrastructure has become so important that cyber space is now considered a new warfare domain and an

element critical for national security that needs to be protected from all lands of threats, including state-sponsored adversaries,

[1070] In one embodiment, a technique to perform real-time integrity assessment is provided in smart phones and other embedded platforms by monitoring their dynamic power consumption and comparing it against signatures from trusted code. The method and technology described build from the general concept of power fingerprinting and provide enhancements for general application on complex commercial devices. We present examples of preferred embodiments of the general techniques to be used as references and examples. The techniques, however, are general and can be adapted to any cyber platform,

[1071] As part of the approach, we also describe a methodology to pre-characterize the way a specific platform and processor consume power to improve the performance of the approach by concentrating classification efforts on the sections of the traces that carry the most information about the internal execution status of the processor and ignore redundant or extremely noisy features that can hamper performance.

[1072] The goal is to enhance the general power fingerprinting (PFP) approach to define a reliable technique to detect, unauthorized software modifications in smart phones, embedded systems, and general information systems. The general prior art approach is depicted in Figure 1.

[1073] The general PFP method begins by collecting fine-grained measurements from the power consumption during the execution of trusted code. The sensor 110 needs to collect a direct or indirect metric representation of the dynamic power consumption or instantaneous current drain of the processor. The sensor 110 can be implemented by means of a commercial current probe, a Hall effect sensor, piezoelectric/magnetostrictive, composite magnetic field sensor, Rogowski coil, a high-bandwidth current mirror, or a simple low-resistance precision shunt resistor. Notice that the sensors need to meet the requirements set by the specific feature extraction techniques selected.

[1074] The physical location of the sensor is a critical element for the success of this approach. The ideal location 210 is shown in Figure 2 at the V_{DD} signal of the processor 205. If this location is not feasible, or introduces excessive power supply noise, then the second best location 220 is also shown. If the sensor 220 is placed in the second location the copper traces with their parasitic capacitance and inductance along with the decoupling capacitors

215 create a low-pass (LP) RLC filter that affects the current traces. For PFP it is beneficial to pre-characterize this hardware effect by identifying the transfer function, H , of the LP filter using a commercial Network Analyzer or another system identification technique. The effect of the inherent LP filter can be minimized by passing the traces through another filter with the inverse transfer function, H_{inv} . It is recommended to implement the inverse filter digitally. Since the direct inversion of H can lead to a unstable H filter, it is necessary to select an the closest stable approximation of H_{inv} .

[1075] In Figure 2, V_{DD_core} 225 can be provided by different sources. For simple processors, it comes directly from the voltage regulators. For more sophisticated platforms, it can come from a power and peripheral management system, which is a complex circuit that provides a wide array of services including delivering different voltage levels required, reset and interrupt handling, and other peripheral management. Power managers are complex systems that merge different signals and add interference from the PFP perspective and tend to hide the power signatures. For system with a power management circuit, it is recommended to design the system board with the necessary provisions to place the current sensor after the power management system to avoid the extra interference and facilitate signature extraction. In a. best case scenario, the power sensor would be included in the power management system as another service provided, facilitating the integration of PFP.

[1076] In the case of multiple processors in the board, the same principle can be repeated for each processor, as shown in Figure 3, where the n th processor 206 is preferably monitored at 21 1 or at second best location 221 after decoupling capacitor 216. In this case, the detector must be designed to combine and consider traces from both sensors. For multi-core processors in the same package, the same principles apply as in the multi-processor example, but the location and feasibility will depend on the architecture of the processor, the number of cores powered by each rail, and decoupling requirements.

[1077] With the sensor in place, the next step is to characterize trusted code. This process is accomplished by repeatedly executing the target trusted code in a controlled environment (including isolation of the target software, set inputs used during execution, and insertion of specific markers that, help synchronizing traces). Markers can be of different nature and help with triggering and synchronization. Potential markers include physical signals (as changing the voltage level of a pin) or a specific sequence of instructions that yields a known power consumption sequence. An example of a physical trigger signal 410 is shown in Figure 4. The

concept of instruction insertion for triggering is depicted in Figure 5. In this case the extra assembly instructions 515 are chosen to yield a known pattern 510 in the traces, usually strong variation in the current drain for a short period of time to help indicate when a specific code 510 is executed.

[1078] When the target application 610 is running on the User Space in a platform that implements the Linux device driver paradigm, or in any other operating system with indirect access to physical signals, as described in Figure 6, it is necessary to account for the inherent uncertainties in execution and timing caused by the indirect access. In this case, the trigger instructions 515 will be executed in the User Space 610 which has no direct access to Physical Memory 640, and can only access the registers 632 necessary to create the physical signal 650 by means of a device driver 631 located in the Kernel Space 620. Uncertainties in execution and timing exist because file access requires the process to wait (block execution) for the appropriate synchronization signaling during which the kernel 620 schedules other process to run.

[1079] Even though the markers 630 are not required to remain in the final code, the process of run-time assessment is facilitated if they remain in place. In the case when the markers are left on the deployed version, it is necessary to ensure that the facilities or services used for the markers will still remain in the deployed platform (e.g. if the marker is supposed to turn on a LED 640, that LED 640 must exist on the deployed platform).

[1080] It is important to note that during characterization the exact code that will be deployed needs to be used. This includes using the exact same tools to build the software, with the same level of optimization, etc.

[1081] For better performance, the characterization should be an iterative, interdependent process, during which the source code structure along with the respective markers are co-developed to yield the strongest signatures with the smallest variance across different, execution instances.

[1082] Several traces from the execution of the trusted code may need to be collected in order to average them and reduce the impact of random noise inherent to any physical system. The characterization process is depicted in Figure 7. After inserting 710 markers into the code, the trusted software is executed and the resulting power traces are captured 720. This is done for all the significant execution paths 730, using predefined input 735 if

necessary. The variations due to random parameters are removed using PCA (principal component analysis) 740. Discriminatory features are extracted 750 and statistical analysis, averaging and clustering 760 is done to generate a set of authorized signatures 770.

[1083] The signatures can be extracted from different, signal domains and be multidimensional. Furthermore, multiple signatures can be used to identify a single piece of code.

Trace Processing And Feature Extraction

[1084] The process of preparing test traces to be compared against the stored signature is known as preprocessing and feature extraction. Trace preprocessing involves general tasks to condition the traces to extract the selected discriminator}' features, e.g. converting the traces to the appropriate domain or aligning the traces in reference to a specific marker. An example of trace preprocessing is shown in Figure 8, in which time-domain traces from the execution of test, software in a BeagleBoard with an OMAP3 processor are first, converted to the frequency domain by calculating their power spectral density.

[1085] Another example of basic preprocessing is to align time-domain traces, as shown by the alignment of the base execution and alternate (-1 bit transition) traces in Figure 9, before being passed to a correlation detector. In this example, each trace of N samples is considered as a point in a multidimensional Euclidean space.

[1086] Feature extraction is the process of calculating the final test, statistic (from new traces) which is passed to the detectors and used to determine integrity. This process is unique to each selected feature. For example, in basic time domain correlation analysis, preprocessing could include coarse synchronization and compensation for specific platform power consumption patterns, while feature extraction involves comparing against the stored signature by calculating the correlation factor or the Euclidean distance. An example of feature extraction is shown in Figure 10, which shows the PSD error in dBs of test traces corresponding to the execution of the trusted code and tampered code in the BeagleBoard 's OMAP3 processor following the PSD example in Figure 8, Using this difference vector, the final test statistic or discriminatory feature passed to the detector can be represented by the mean squared error or any other distance or error metric.

Detector Design

[1087] Once the signatures have been extracted and the discriminator}' features have been selected, the next step in the PFP process is to design optimal detectors to perform the final integrity assessment. These detectors will make the final decision of whether a test trace should be considered an intrusion during monitoring operation. The process of detector design and normal monitoring operation are very similar. In detector design, test traces from the execution of trusted soft ware are captured and processed to extract the selected discriminatory features and compared against the stored signatures. Several traces are collected and processed and their statistical sample distributions are used to identify a threshold that yields the expected performance targets. The process of detector design is shown in Figure 11. Random or predefined input 1110 is provided to trusted software 1120 and fresh test traces are captured from its execution. The results are aligned and synchronized 1130, and the traces are preprocessed and conditioned 1140. Using authorized signatures 770 for comparison, the selected discriminatory features are extracted and a distance metric is generated 1150. Then statistical analysis and distribution fitting is done 1160 on the resulting metrics. Finally, the Neyman-Pearson criterion is applied 1170 to determine a threshold that meets expected performance targets.

[1088] A common approach to create optimal detectors involves the application of the Neyman-Pearson criterion to maximize the probability of detection for a given probability of false alarm. As a brief reminder of this criterion, which is spawned from basic hypothesis testing theory, a target probability of false alarm is set based on the tolerance and estimated cost of making a mistake in the final decision. Using an estimate of the probability distribution of the discriminatory features from the trusted code, a distance threshold is calculated that yields the expected probability of false alarm while maximizing the probability of correct, detection. An example of this process is shown in Figure 12, in which a distance threshold 1220 is calculated for a probability distribution 1210 that yields an expected probability of false alarms 1230.

[1089] It is important, to note, however, that there are different techniques that can yield improved results depending on the nature of the selected discriminatory features. Other techniques for detector design and machine training include: Neural Networks, Support Vector Machines, and Hidden Markov Models.

Monitoring Operation

[1090] Once signatures have been extracted from the execution of trusted code, discriminatory features have been selected, and optimal detectors have been designed, the PFP monitor is ready to assess the integrity of test software. As mentioned before, the normal integrity assessment process is very similar to the detector design process. During normal operation, the monitor also extracts the selected discriminatory features from power traces after the necessary preprocessing, but, instead of collecting the statistics from several traces as was done for detector design, they are passed through the appropriate detector to compare against the respective thresholds and determine the integrity status of the test code execution. The detector compares the test traces against all known signatures and, if no single test statistic is enough to determine that authorized code has executed, then an intrusion is reported. This process is depicted in the diagram shown in Figure 13. The target software is executed 1310 during normal operation or using predefined input to capture test traces 1320, which are then aligned and synchronized 1330, and then preprocessed and conditioned 1340. The detector then compares 1350 the extracted features against the known signatures 1370 to determine a distance, using the predefined threshold 1220 to make an integrity assessment decision 1360.

Sample Results

[1091] In order to exemplify the PFP process on smart phones and other embedded platforms, we describe a reference implementation of this technique using a BeagleBoard revision C4 with the ARM processor (OMAP3 @ 720 MHz) running the Android platform. The BeagleBoard 1410 is slightly modified by cutting the main traces 1420 providing power to the core power rail in order to connect a current probe 1430. The capture system is implemented using a commercial real-time oscilloscope 1440 and current probe 1430. The oscilloscope is configured to a sampling rate of 2.5 GSPs and a total of 30K samples are collected in every trace initiated by the trigger 1450. The setup is described in Figure 14.

[1092] A basic test app was developed to demonstrate the process and show feasibility. This basic app consists of a simple counter that displays an increasing integer on the device screen. The operation of the app is described in LISTING 1 and consists of a typical Android Java application structure with an initialization routine that prepares the screen for showing a text box and sets an integer variable used as a counter. There is also a routine called DisplayCounter in charge of incrementing the value of the counter and displaying it on the screen. This routine is configured as a recurrent task that is called every second.

- i. LISTING 1. Pseudo-code of Android Test App
- ii. Initialize
- iii. DisplayCounter ()
- iv. {
 - 1. counter = IncrementValue (counter);
 - 2. Display data
- v. }
- vi. Sleep for one sec
- vii. DisplayCounter at wake up

[1093] The critical Increment Value routine was implemented in native C code and included as an external library by Android's NDK toolset, instead of the traditional Java implementation. Before the critical section, a physical trigger 1450 is set to signal 1460 the capture systems to start collecting power traces.

[1094] We only characterize and monitor the critical Increments/value routine shown in Listing 2.

- i. LISTING 2. Pseudo-code of monitored native routine in C
- ii. /*Critical native routine*/
- iii. int incrementValue(int Val)
- iv. {
 - v. /*trigger LED usr1*/
 - vi. Open device driver control file
 - vii. Write 1 into file
 - viii. /*Increment Val*/
 - ix. Val++;
 - x. /*General extra processing*/
 - xi. i = 1000;
 - xii. while(i)i--;
 - xiii. /*Reset LED usr1 Trigger*/
 - xiv. Write 0 into file
 - xv. Close driver control file
 - xvi. return Val;
 - xvii. }

[1095] Signature extraction is performed in the frequency domain by simply averaging the PSD of several traces from trusted code execution. The phase information of the trace is ignored. The PSD of two hundred traces is averaged together to yield the signature.

[1096] The discriminatory features are extracted in the frequency domain as well by a mean-squared error between the signature and the PSD of test traces (in dBs). The PSD of the latest

three test traces are averaged together before calculating the MSE. Only the first 200 MHz of the PSD are used in the MSE calculation.

[1097] This process for signature extraction yields a mono-dimensional discriminatory feature.

[1098] Detector design was performed using the Neyman-Pearson criterion described earlier using a target probability of false alarm, P_{FA} , of 1%. The sample statistics of the trace are extracted from a sample of 200 traces from the execution of the trusted code.

[1099] The sample distribution was fit to a Rayleigh distribution with mean and variance equal to the training sample distribution mean and variance. Using this distribution, the inverse probability distribution is calculated to find the threshold that yields the target 1% P_{FA} .

[1100] In order to test the ability to detect execution deviations from trusted code, we test the previously designed monitor using a slightly tampered version of the app. The tampered app, shown in LISTING 3, is designed to emulate a covert attack in which the intrusion remains inactive until a specific condition is met. The intrusion consists of a very simple modification in which a file is written only when the value of a counter reaches a specific value (the condition).

```

i. LISTING 3. Pseudo code of tampered critical native routine
ii. /*Critical native routine*/
iii. Int incrementValue(int Val)
iv. {
v. /*trigger LED usr1*/
vi. Open device driver control file
vii. Write 1 into file
viii. /* Tamper */
ix. if(Val == 1){
x. //open temporary file
xi. //Write Val into file
xii. //Close file
xiii. }
xiv. /*Increment Val*/
xv. Val++;
xvi. /*General extra processing*/
xvii. i = 1000;
xviii. while(i--);
xix. /*Reset LED usr1 Trigger*/
xx. Write 0 into file

```

```
xxi. Close driver control file  
xxii. return Val;  
xxiii. }
```

[1101] It is important to note that the file writing taking place in the tamper only occurs once during execution (i.e. when the counter is 1). The rest of the time, when the routine is called, the condition is not met and the extra file is not written. Hence, for the majority of the time the routine is called, the only modification from a logic standpoint is an extra evaluation of a given condition.

Operational results

[1102] The results from running the monitor when the original untampered version of the routine is executing are shown in Figure 15.

[1103] We can see that for the duration of the test, we had only a few instances that went past the threshold 1510, which is consistent with the designed probability of false alarm.

[1104] The results from running the monitor with the tampered version of the app are shown in Figure 16. Notice that no instance is misclassified as authorized execution and every single execution of the tampered app would be flagged as an intrusion above threshold 1610. It is also important, to note that due to the conditioned execution of the tamper, only once during the execution instances used in these results was the file actually written. The rest of the time, only a condition was checked, and when it was not met, normal execution resumed.

Platform characterization and evaluation of minimum sensitivity

[1105] The fine-grained measurements of the power consumption can lead to redundant information that adds very little discriminatory information, but that, can add significant noise and uncertainty to the signatures. In the time domain this looks like Figure 17. In this case, we would like to focus our attention on the sections of the traces (dimensions) that have the largest variance 1710 between the two executions, in contrast to the sections, e.g. 1720, which show little variance between the two executions. On the other hand, when characterizing a specific software routine that takes random parameters, the effect of these random parameters is to introduce noise to the signatures, which ends up reducing performance and increasing the probability of false alarm. In this case, we would like to focus our attention to the dimensions (e.g. 1720) that remain constant during the execution of the

target software, while ignoring the ones that add noise. In this case, we would like to ignore the dimensions that display large variance (e.g. 1750).

[1106] In order to improve the performance of PFP, it is necessary to reduce the number of features analyzed by concentrating on only the ones that carry the most information. This is accomplished by pre-characterizing the features that carry the most information for a given platform as part of the training and then eliminate redundant, information during preprocessing before passing the traces to the detectors.

Technical Background

[1107] In traditional pattern recognition systems, the process of selecting a subset of features that maximizes a specific criterion (in the case of PFP we want to maximize discriminator) information PFP), is known as optimal feature selection. In clustering systems, this is normally accomplished by projecting the traces, x , to a transformed space with fewer dimensions from the most useful (or informational perspective) by means of a linear transformation.

[1108] This transformation is described as

$$y = Wx$$

[1109] Where W is a carefully designed linear transformation matrix that when applied to test traces, yields a transformed trace with lower dimensionality that maximizes a particular criteria. There are different criteria to identify the optimal transformation. Because we are trying to optimize feature selection in terms of discriminatory information, it is natural to follow an information theoretical approach. This optimization has been performed before and can be found in several sources in the pattern recognition literature, for example see J. T. Ton and R. C. Gonzalez. "Pattern Recognition Principles," Addison- Wesley Publishing Company, 1974; which is incorporated herein by reference.

Principal component analysis

[1110] A well know approach to determine the appropriate W that optimizes the entropy (or information) in the traces is known as Principal Component Analysis (PC A). We assume that the covariance matrices of the different classes, C_i , are normally distributed and identical $C_i=C$. Hence, the eigenvectors can be considered as the information bearers for the traces

under consideration. Some of these vectors carry more discriminatory information in the classification sense than others, which can be safely eliminated without much performance penalty. It should be no surprise that the optimal feature vectors are tied to these eigenvectors and are used to create the transformation matrix W by aggregating eigenvectors in descending order according to the corresponding eigenvalue. Because in PFP we only need a single point per clock cycle, the transformation matrix W given by the eigenvector of the covariance matrix associated to the largest eigenvalue.

[1111] The linear transformation can be interpreted as a projection of the test trace into a transformed space of lower dimensionality from the most informative perspective. PCA can be applied to in different ways, depending of the specific objective. From a clustering perspective, it is preferred to construct W using the eigenvectors associated with the smallest eigenvalues, as this would yield a tighter cluster in the transformed space. On the other hand, it is also possible to use the eigenvectors associated with the largest eigenvalues when traces from different executions are used. When applied this way, PCA will select the features that present the largest variance between classes. With the assumption that the covariance matrices are identical, these eigenvectors will represent the features that contain the maximum discriminatory information between the specific traces used for PCA.

Linear Discriminant Analysis (LDA)

[1112] PCA selects a feature subset in ascending, or descending, order in terms of variance to optimize trace entropy. It does not consider, however, the specific differences between classes to select an optimal set of features that maximizes the distance between distributions. Linear Discriminant Analysis (LDA) maximizes the divergence between distributions, which is a measure of distance between probability distributions. Divergence is closely related to the concept of relative entropy in information theory.

[1113] Using specific information from different classes and divergence as optimization criterion, LDA identifies the optimal transformation matrix to project the traces from the unique perspective that yields the maximum separation between them. This is because the transformation vector W is normal to the optimal discriminating hyper-plane between both distributions.

[1114] Following the assumption that traces are normally distributed, it can be shown that the transformation matrix that yields a divergence extremum is given by the only eigenvector of $C^{-1}\delta\delta^T$ associated with a non-zero eigenvalue. This vector is given by

$$W_0 = C^{-1}(\mu_1 - \mu_0)$$

[1115] where W_0 provides the optimal projection to separate both classes while μ_0 and μ_1 , are the respective centroids for the two training classes. LDA can be extended to M discriminating classes. In this case, there will be M-1 eigenvectors associated with non-zero eigenvalues.

Platform power consumption characterization

[1116] As mentioned before, not all samples in a test trace are equally important to determine whether an execution deviation has taken place or not. Because of the large oversampling ratio and the nature of the power traces, there are certain sections of the traces that carry more discriminatory information than others. For PFP, the goal is to identify a linear transformation that reduces the dimensionality of the traces by removing redundancy while emphasizing the dimensions that carry the most information.

[1117] The idea is to transform discriminatory features to reduce dimensions by using a linear projection of the traces using an optimal transformation matrix. In the time domain, trace sections corresponding to a full clock cycle 1810 are reduced to a single point 1820 in the transformed space, as depicted in Figure 18. Classifiers also have to be designed to operate in the transformed space, reducing the number of dimensions that need to be considered during normal monitoring operation.

[1118] Characterization is performed under controlled conditions in the lab and is required only once per platform. As described in the previous sections, there are two general approaches to identify the optimal transformation matrix: PCA and LDA.

Platform characterization using PCA

[1119] In order to create a transformation matrix using PCA, it is necessary to observe the power consumption of the processor during random clock cycles. The traces are aligned for every clock cycle to clearly show the sections of the traces that are affected the most by the

dynamic behavior of processor execution. Once the traces are aligned, PCA is used to identify the transformation vector that accounts for the most variance in the traces,

[1120] Performing platform characterization using PCA is relatively easy to implement and well suited for complex platforms in which controlling the contents in the pipeline results too difficult.

Platform characterization using LDA

[1121] Performing platform power consumption characterization using LDA requires the development of two carefully tailored routines. These routines must execute the specific instructions with specific addresses and parameters in the right sequence to create two sets of traces that show predetermined differences during a specific clock cycle. Training traces from the execution of both routines provide the two classes for which LDA will find the optimal discriminating hyperplane, which will in turn become the optimal transformation vector.

[1122] The objective of the special characterization routine is to execute a carefully crafted sequence of instructions to properly load the pipeline such that in a specific clock cycle there is a known change during each execution stage (fetching, latching, execution, etc). The changes should be relatively small, preferably due to a few changing bits in the respective registers. The characterization routine is not unique, but it is platform specific as it depends on the architecture, instruction set, etc. of the platform being characterized. Different processors will likely require a different sequence.

[1123] Once the traces from the execution of both sequences are captured and synchronized, LDA is used to find the optimal transformation vector W . It is expected that platform characterization using LDA will provide the best performance, given the availability of two known classes, but its implementations is more complex than PCA.

Platform Power Consumption Characterization Reference Implementation Results

[1124] For this reference implementation, we use a motherboard with a PIC1 8LF4620 8-bit microcontroller from Microchip Technology Inc., similar to the ones used in the PICDEM Z Demonstration Kit, intended as an evaluation and development platform for IEEE 802.15.4. This is a popular embedded microcontroller without a memory management unit.

[1125] The processor motherboard is slightly modified in order to enhance power consumption features. From the board, a total of six decoupling capacitors are removed totaling a cumulative 6 microF. The function of these capacitors is to mitigate the stress placed on the power supplies by the strong current peaks caused by digital processors. It is important to note that removing decoupling capacitors would not be necessary if the current sensor is placed closer to the processor power pins, or if the effect of the resulting LP filter is cancelled using signal processing,

[1126] Trace collection is performed using a Tektronix TDS 649C real-time oscilloscope 1910 and a Tektronix TC-6 current probe 1920. The probe is connected right passed the voltage regulators on the mother board. The oscilloscope is configured to 500 MS/s and 10 mV. The trigger is driven by LED1 1930, and configured for falling-edge, 40 mV level, and no pre-trigger samples are kept. A total of $L = 30,000$ samples are collected after every trigger event. The measurement setup is depicted in Figure 19. Traces are captured and transferred to a host computer using GPIB for their posterior analysis.

[1127] A sample routine is developed for this experiment with a dual purpose 1) to provide training routines to perform the platform characterization and 2) to provide a reference change to measure the performance of the approach. We start by describing the evaluation usage of the routine and provide baseline performance for comparison. The test routine is shown in LISTING 4 and executed in an infinite loop. In this routine, the contents of Register W are toggled from 00 to 0f using different instructions. Notice that the actual logic in the routine has no impact on the performance of power fingerprinting. This routine was chosen because it is easy to control the number of bit transitions that happen. The results, however, do not depend on the specific software being executed. Therefore, this routine provides a representative example.

```

i. LISTING 4.
ii. BYTE i; //addr 00
iii. BYTE j; //addr 01
iv. BYTE k; //addr 10
v. BYTE l; //addr 11
vi. // Initialize the system
vii. BoardInit();
viii. // Initialize data variables
ix. __asm
x. movlw 0x07
xi. movwf i, 0 //addr 0x00
xii. movlw 0x0f
xiii. movwf j, 0 //addr 0x01
xiv. movlw 0x0f //Set for minimum change
xv. movwf k, 0 //addr 0x10
xvi. movlw 0x1f
xvii. movwf l, 0 //addr 0x11
xviii. movlw 0x00
xix. __endasm
xx. //Target code infinite loop
xxi. while(1){
xxii. TMR0H = 0x00; //Restart TIM0
xxiii. TMR0L = 0x00;
xxiv. LED_2 = 1; //Trigger
xxv. LED_2 = 0;
xxvi. __asm
xxvii. nop
xxviii. iorwf j, 0, 0 //w = 0f
xxix. andlw 0x00 //w = 00
xxx. movf j, 0, 0 //w = 0f
xxxi. andlw 0x00 //w = 00
xxxii. movf k, 0, 0 //w = 0f Change in k (one bit)
xxxiii. movlw 0x00 //w = 00
xxxiv. xorwf j, 0, 0 //w = 0f
xxxv. movlw 0x00 //w = 00
xxxvi. iorwf j, 0, 0 //w = 0f
xxxvii. xorlw 0x00 //w = 00
xxxviii. nop
xxxix. ... x 10
xl. nop
xli. __endasm}

```

[1128] The routine, as shown in LISTING 4 represents the base execution. Preceding the target code, we create a trigger using an LED on the board. The trigger is used to synchronize trace capture with the oscilloscope. The "NOP" instruction between the trigger and the target code is included as a buffer to isolate the target traces from any residual effects from the trigger. Once inside the main loop, Register W is toggled from 00 to 0 f creating four bit

transitions in that register every instruction. The alternative, or modified, code has one fewer bit transition. In Line 15, we change the contents of variable j from $0f$ to 07 . This way, when the target code is executing, in Line 35, the parameter k is loaded onto Register W which goes from 00 to 07 , with only three bits transitioning in the register for that instruction. Note that there is just one bit difference between this modified code and the base execution which loads Register W with a $0f$ and that everything else in the execution is kept the same, including instructions, parameters, and addresses. Note that this one-bit change actually affects two clock cycles, as there is one less transition coming into that instruction and one less coming out of it. Trailing the target code there is a string of "NOP" instructions before the loop is repeated.

[1129] A detail of a typical trace is shown in Figure 20. In this figure we capture one full execution cycle of the target code. The trigger effects on the power traces are clearly visible as two square steps 2010 and 2020. Individual instruction cycles are also noticeable. They can be identified as groups of four spikes that repeat every 125 samples. Using timing information from the processor's documentation, we can determine the section of the trace that corresponds to the execution of the target code. In Figure 20, this section is highlighted as a solid line 2030 that spans ten instruction cycles. This is in agreement with the actual code, which consists often of assembly instruction, each taking one bus cycle to execute.

[1130] Several traces are captured from each of the base and alternative executions and the traces from each execution are averaged together to provide a clean picture of both executions showing the total effect of one less bit transition. The averaged traces are shown in Figure 21. In this picture, the ten clock cycles corresponding to the execution of the base code are shown and it appears that the traces from each execution are aligned. Around sample index 650, however, there can be seen a small difference between the two traces. The difference (at 2110) is more noticeable in the top of Figure 21, which provides a closer look. Along with the closeness of the centroids from both scenarios, it is also evident that traces are largely correlated due to over sampling and also that only certain sections of the traces carry useful discriminatory information.

[1131] For comparison purposes, we provide the results of a naive classification approach in the time domain without platform pre-characterization. We use a basic minimum-distance classifier. In this approach each captured trace of length $L = 1250$ (the length of the target code) represents a point in an L -dimensional Euclidean space. The Euclidean distance is

taken from the base execution centroid to each incoming test trace. For classification purposes, the base centroid and test traces represent a single point, or vector, in a multidimensional Euclidean space with 1250 dimensions. The test traces are different from the training ones used to obtain the base centroid. This is to avoid a bias in the evaluations of a minimum-distance classifier to accurately discriminate between the different scenarios.

[1132] Test traces from the execution of both routines have the Euclidean distance distributions shown in Figure 22. In this naive example, the performance of power fingerprinting is not encouraging, as there is barely any difference between the distributions, which substantially overlap. This poor performance is expected considering the small differences in power consumption between the base and alternative scenarios.

[1133] The first results for platform characterization are obtained from, applying PCA. For this process we use all clock cycles corresponding to the execution of our target code in the routine shown in LISTING 4. The trace corresponding to the full execution of the trace is divided into different sections corresponding to a single clock cycle execution. The subsections are then aligned and PCA is used to find the transformation vector W corresponding to the eigenvector that accounts for the most variance. In this case, as explained before, we take the oversampled trace for one clock cycle and we reduce it to a single point.

[1134] After performing the platform characterization using PCA the test traces from the evaluation routine are processed again to demonstrate the performance improvements of platform pre-characterization. The minimum distance distributions from the transformed test traces to the signature in the new PCA transformed space are shown in Figure 23.

[1135] A clear separation is seen between the bulk of the distributions, which represents a clear improvement with respect to the naive classification performance shown in Figure 22. Results with platform characterization using LDA.

[1136] In order to obtain the training traces necessary to apply LDA, we execute the base routine and a slightly modified version. We obtain the special platform characterization traces by comparing two sets of traces: from the base execution, which is once again the code in LISTING 4 and a slightly modified version of it shown in LISTING 5. The changes in execution are carefully selected to cause one less bit transition on each execution stage compared to the base execution. In this modified version, the instruction in Line 36 is

changed from xorwf with opcode 0001 10da to iorwf with opcode 0001 00da (the optional arguments d and a, control the destination and RAM access bit, respectively, and are kept with the same value in both cases). During execution the difference in the opcodes will cause one-less bit transition when latching the instruction word. The parameter in the instruction changed from j, located at address 0x01, to i, located at address 0x00 in Access RAM. Once again, the change will create one-less bit transition when executed. Furthermore, notice that the contents of the j and i also differ in one bit. This will also translate into one less bit transition when parsing the parameter, when executing the instruction and when writing the results.

i. LISTING 5. Modified routine for platform characterization
 ii. ...
 iii. 35 movlw 0x00 //w = 00
 iv. 36 iorwfi, 0, 0 //w = 07
 v. 37 movlw 0x00 //w = 00
 vi. ...

[1137] For platform characterization we use only traces corresponding to the execution of Line 36 in Listing 5, The average of these traces (for each execution, the base execution and the one-less bit, transition) is shown in Figure 24,

[1138] Using these traces, we perform LDA to identify the optimal discriminating hyperplane and the linear transformation that project our traces from the most informative perspective. The test traces from the evaluation routine are processed again to demonstrate the performance improvements of platform pre-characterization. The minimum distance distributions from the transformed test traces to the signature in the new LDA transformed space are shown in Figure 25.

Detecting Deviations from Authorized Software Execution in software-controlled radio platforms mid other embedded systems.

[1139] Dynamic power consumption of a processor can be monitored to determine whether it corresponds to the expected execution or a deviation has occurred.

Platform Description

[1140] The exemplar target platform for illustrating this use of power fingerprinting (PFP) is a software-controlled radio, in which the specific configuration of the radio behavior is

controlled by software. A generic block diagram of the intended platform is shown in Figure 26.

[1141] In this platform, the behavior and configuration of the RF Transceiver 2610 is controlled by the processor 2620. The application 2626 represents the highest layer and implements the intended functionality for the processor. In order to interact efficiently with the RF Transceiver 2610, there is a set of Application Program Interfaces (APIs) 2624 that abstract the complexity of the interactions with the hardware to the main application. These APIs, together with the required drivers and protocol stack implementation 2622, provide a board support package for the specific transceiver. The protocol stack 2622 arranges the data to be transmitted in the prearranged format, adding required headers and preparing the payload so the intended receiver can extract the information. It is also in charge of extracting the information received from remote devices and presenting it to the application layer 2626. The cryptographic module 2612 can be implemented in the RF transceiver or in software as part of the protocol stack. Figure 26 shows it as part of the transceiver. The location of the cryptographic module 2652 presents no practical difference in the approach. The MAC 2614 and PHY 2616 layers of the RF Transceiver 2610 are in charge of medium access and physical transmission and reception of the information.

[1142] The described approach characterizes the execution of the Application software 2626, in particular the execution of the API calls 2624 that impact the behavior of the cryptographic module 2612. In this approach, the specific code executed as a result of an API call is used to determine whether encryption was used and the type of encryption used. For example, if the application calls for a specific kind of encrypted transmission, this approach asserts the execution of the code that invokes the encryption. In case of malicious or accidental tampering, this approach provides a reliable indicator of the modification.

Signature Extraction

[1143] Traces of length L captured during the i -th execution of authorized code a are represented by:

$$r_a^{(i)}[n]; \quad n=0, \dots, L-1$$

[1144] In order to avoid potential low- frequency interference from other board components, a basic no-multiplier high-pass filter is introduced by calculating the difference between trace samples

$$d_a^{(i)}[n] = r_a^{(i)}[n] - r_a^{(i)}[n-1]$$

[1145] Several captured traces from the execution of the authorized code are used to create a signature, our target fingerprint. N traces are averaged to form the target signature and reduce the effects of random noise in the individual traces

$$s_a[n] = \frac{1}{N} \sum_{i=0}^{N-1} d_a^{(i)}[n]; n = 0, \dots, L - 1$$

Feature Extraction

[1146] The process of extracting discriminatory features consists of simple time-domain correlation against the target signature. The correlation, however, is performed on $j > 0$ partial sections of the signature and the trace, each section has a length $w = \text{floor}\{L\}$. This partial correlation is performed to avoid spreading potential differences in the power traces across a full-trace correlation.

[1147] The cross correlation for different sample lags, $0 < k < w$, of section j of the traces is given by:

$$\rho_{s_{ad_b(i)}}(j, k) = \frac{1}{(w-1)\sigma_s\sigma_d} \sum_{n=(j-1)w}^{jw} s_a[n]d_b^{(i)}[k+n] - w\bar{s}\bar{d}$$

[1148] where \bar{s} and σ_s , are the sample mean and standard deviation of the corresponding section in s_a . and \bar{d} and σ_d are the sample mean and standard deviation of the corresponding section in $d_b^{(i)}$.

[1149] In order to compensate for any clock drifts, we keep the maximum correlation values for different lags. This action reduces the dimensionality of our traces to only a sequence of j peak correlation values for every trace:

$$\hat{\rho}_{s_{ad_{b(i)}}}(j) = \max_k \left\{ \rho_{s_{ad_{b(i)}}}(j, k) \right\}$$

[1150] Under ideal conditions and with $b=a$, $\hat{\rho}_{s_{ad_{b(i)}}}(j) = 1$ for every section j . Any deviation from the power consumption characteristics would be reflected by a reduced correlation factor.

[1151] The actual discriminatory feature or test statistic used in this work to evaluate traces is the minimum peak correlation value for that specific trace.

$$x_b^{(i)} = \min_j \left\{ \hat{\rho}_{s_{ad_{b(i)}}}(j) \right\}$$

$$X_b = x_b^{(i)} ; i=0, \dots, N-1$$

[1152] The random variable $x_b^{(i)}$ indicates the maximum deviation from the signature of instance i of code b . Using X_b we can design appropriate detectors using different criteria depending on the statistical information we can gather from the system a priori.

Response to integrity violations and layered security

[1153] PFP is a very effective approach to detect execution deviations in cyber systems. In order to have a complete solution, however, it is necessary to have a structured policy to handle integrity violations when the PFP monitor detects a deviation from expected execution.

[1154] There are three clearly defined phases in computer security: Prevention: Includes active mechanisms to deter, discourage and prevent attackers from carrying on attacks to disrupt the system, disclose information, etc.

[1155] Detection: Because perfect absolute prevention is not feasible, it is necessary to perform constant monitoring of the integrity of the system.

[1156] Response: The set of policies set in place to react to successful attacks.

[1157] The architecture for integrating PFP into a comprehensive defense-in-depth security approach will now be described. In this approach PFP provides a robust, solution for the "Detection" phase to complement a number of different techniques to prevent and deter

potential attacks. The appropriate reaction to different successful attacks is defined in the "Response" phase and is described according to the security policy described below.

[1158] While achieving system security requires a process and not just isolated mechanisms or technologies we will concentrate on describing the areas where PFP can complement traditional security mechanisms to provide continuous or intermittent monitoring for integrity assessment and intrusion detection. Before describing the role of PFP it is important to mention that the security process involves several steps including:

[1159] Design. Follow sound design approaches and design the systems to facilitate security enforcement, reduce vulnerabilities, enforce access control, etc. A typical example is to design the system such that security functions are isolated from the rest of the functionality and where access control features are inherently enforced.

[1160] Development. Follow best development practices to produce maintainable products with reduced vulnerabilities.

[1161] Deployment. Make sure only authorized modules are deployed. This requires strong authentication and non-repudiation approaches.

[1162] Operation. Maintain a secure environment by enforcing strong access control and other security policies.

[1163] Monitoring. Constantly assess the integrity of the system. PFP, antivirus, and Network intrusion detection systems.

[1164] Response. Define the policies and procedures to be followed when an attack is successful. Policies should be developed considering the criticality of the systems and should be strictly enforced.

[1165] This section describes an architecture to integrate a PFP monitor into a comprehensive security solution that includes complementary security mechanisms where the vulnerabilities of one layer are covered by the next. The approaches and technologies included in the different layers include: encryption of data at rest, strong authentication, access control, tamper-resistance, firewalls, sandboxes, virtualization, and physical protection. The architecture also provides a mechanism to define and enforce security policies to react to integrity violations detected by PFP.

[1166] The architecture defines a layered security solution where a PFP monitor provides a last line of defense by detecting when an intruder manages to get through all other defense mechanisms. Figure 27 shows the different layers 2700 in a defense-in-depth approach. The different layers are intended to slow an adversary and make it progressively more difficult to breach a defensive layer without being noticed. In the outer layers are the external] defense mechanisms, such as internet firewalls, physical protection of the equipment and passwords and security policies (i.e. to prevent social engineering attacks). The inner layers correspond to different defenses that reside within the host 2750. They start with access control 2740 and encryption of data at rest. They continue with different security mechanisms intended to protect the applications 2760 and the operating system 2770. At the core 2780, there are controls for the most basic kernel and safety operations.

[1167] PFP can effectively monitor the integrity of different layers. At the core level 2780, PFP can assess the integrity of kernel and safety operations that all other mechanisms depend on. It can also be expanded to monitoring the integrity of core applications in the operating system 2770, as well as the integrity of critical user-level applications 2760. Notice that PFP can monitor the integrity of all the modules that reside within the processor scope, including antivirus modules and encryption modules, as shown in Figure 28.

[1168] Integrating PFP into a defense-in-depth approach for cyber security allows faster handling of potential incidents before they can achieve their goals and cause damage. The power signatures from the execution of other security modules, such as encryption and antivirus, are extracted and evaluated at runtime. From the PFP perspective, signatures from a kernel module and an antivirus program are extracted the same way and using the same techniques.

[1169] It is important to note that PFP can be extended to any other device to monitor its execution integrity. This includes devices that can be used to implement different security layers such as firewalls, digital security locks, etc.

[1170] The last stage of cyber defense with PFP is to define the appropriate policies to handle different anomalies detected by the monitoring and assessment operations. Because the appropriate reaction to successful intrusions depends on several factors unique to each platform and application, we cannot generalize the response to different intrusions detected by the PFP monitor. Therefore, it is necessary to follow an architecture that accepts and

enforces different security- policy definitions which can be adjusted to different systems while maintaining and reusing the basic operational principles and structures.

Distributed PFP monitor network to monitor malware dynamics and behavior

[1171] This section describes the operation of a wide network of nodes with PFP capabilities that are deployed across different geographical or logical regions to monitor the spread of malware, detect targeted attacks, and discover the potential intentions of malicious adversaries. This approach is applicable to discovering furtive remote attacks on specific logical or geographical areas.

[1172] One of the main advantages of using PFP for this application is its stealth, which prevents adversaries from detecting the monitoring activities themselves, giving them a false sense of furtiveness (believing they have not been detected) and tricking them into carrying on with their activities, disclosing intentions and capabilities. This application of PFP is a powerful tool for intelligence gathering.

Operation

[1173] Stealth monitoring is achieved thanks to the low footprint of PFP and negligible impact on memory and latency on the target system. The distributed network of PFP nodes is implemented using the following steps:

[1174] Enable representative nodes with PFP (Fit them with a PFP monitor and extract trusted signatures from their target components). The monitor can be rack mounted and resourceful, as the target nodes only act as a honey pot.

[1175] Deploy a network of the PFP enable nodes on the target geographical or logical areas of interest.

[1176] Monitor each node individually for integrity violations and intrusions as depicted in Figure 29.

[1177] Periodically submit the integrity results to a centric location for logging and analysis.

[1178] In the event of an integrity violation, the report should include:

[1179] a copy of the power traces that experienced the violation

[1180] the ordered sequence of untampered module execution that executed prior to the violation

[1181] the ordered sequence of modules that execute after the violation

[1182] This application of PFP is depicted in Figure 30. The figure shows PFP honey pots in different geographical networks. It is important to note, however, that network separation can be logical, as in different sections of the same network, or socio/political, as in networks for different government agencies or corporate divisions.

[1183] The links between the honey pots and the centralized analysis location, represented on Figure 30 as dotted lines, can be implemented as a separate network (e.g. dedicated wireless links) or made using available wide-area networks, such as the public switched telephone network (PSTN) or the Internet. In any case, strong non-repudiation mechanisms, in charge of providing proof (high assurance authentication) of the origin and integrity of the traces, must be placed to maintain trustworthiness in the system as a whole.

Application of PFP to Supply-Chain Trust Analysis

[1184] Outsourcing device manufacturing and production to foreign and untrusted foundries and manufacturers opens the door to potential security breaches and tampering. Even with trusted providers, there is the potential for foreign or disgruntled personnel to try to disrupt the operation and functionality of critical systems.

[1185] PFP provides a mechanism to detect unauthorized modifications and other tampering in software, firmware, and hardware introduced by untrusted links in the supply chain throughout the system's lifecycle. Integrity assessment of new shipments and untrusted devices using PFP requires the steps shown in Figure 31. Input vector generator 3110 is used to provide the necessary inputs for the execution of the target device under a controlled environment 3120 during which power traces are collected 3130. The parameters of the individual characteristics of the power consumption are compensated 3140 before performing feature extraction 3150. The resulting features are compared 3160 against stored reference signatures 3170 and from this comparison the final assessment result is obtained.

[1186] Detecting the integrity of digital devices using PFP is not a destructive process and requires only minimal collaboration from the device being assessed. Furthermore, the fine

grained measurements of the power consumption provide significant visibility into the internal execution status of the device, making it extremely difficult for a modification to go unnoticed. For instance, PFP can detect tampering that activates only under certain conditions (also known as logic and time bombs) due to partial activation of the extra functionality or execution flow during condition checks. The ability of PFP to detect extra or missing functionality does not depend on the purpose or intentions of the insertions themselves.

[1187] Another advantage of PFP is that a specific execution path verified with PFP can be trusted even if malicious activity is not triggered. In other words, if PFP does not detect a significant deviation from the signatures, it means that no tampering or extra functionality has happened in that particular execution path.

[1188] A key element in performing supply-chain trust Analysis with PFP is executing the untrusted device under a controlled environment 3120. This controlled environment includes predefined inputs 3110 that force a specific state sequence and, for programmable devices, the specific software to be executed. For some systems it may be necessary to develop support scaffolding to control and isolate the execution of specific components. The specific input vectors depend on the functionality of the device or software module and they are expected to exercise the critical execution paths for the operation of the device.

[1189] The same input vectors used to extract the signatures need to be used to assess the integrity of the untrusted devices.

[1190] Because of slight process variations during manufacturing, different devices will show different power consumption characteristics. These variations in power consumption need to be compensated 3140 before feature extraction 3150 to avoid erroneous assessments. This compensation is performed by means of an adaptive filter whose taps are dynamically modified to match the specific characteristic of the power consumption traces. This adaptive filter allows the PFP monitor to concentrate on the power consumption resulting from bit transitions in the device register during execution and eliminate differences in the traces due to manufacturing variations.

[1191] The most critical aspect for effective supply-chain trust analysis using PFP is the availability of reference signatures 3170. There are different potential sources for such signatures as depicted in Figure 32. The best reference would be provided by an identical trusted implementation (a gold standard) 3230. In many occasions, however, such trusted

implementation is not available. In these cases a reference signature can be extracted using alternative methods with varying degrees of error and trustworthiness. For instance, two relatively simple alternative reference sources include a previous implementation of the device (one that has been time tested) 3250 or an alternative implementation from a different supplier 3260. In these cases, the signatures are extracted from the execution of the alternative implementations, reducing the chances of two identical modifications by different providers. Signatures from the former approach can miss unidentified modifications present in the previous version. In the latter approach, an attacker could create an identical modification in both versions from the different suppliers to avoid detection.

[1192] Using a CAD model 3240 to obtain the signatures requires more effort, but it can be done in-house without relying on foreign foundries. In order to extract the signatures using a CAD model it is necessary to simulate the execution of the device using deterministic input vectors. The simulator needs to be accurate in terms of power consumption to the register transfer level.

[1193] **Digital rights management and execution limited leases**

[1194] Another novel application for PFF is the enforcement of Digital Rights and the creation of execution instance-limited leasing to enable licensing based on the number of executions.

[1195] This approach is implemented by extracting signatures from the execution of protected software and monitor at run-time the power fingerprints to enforce the sole execution of authorized modules. For example, a software system can be licensed to include only a set of functional modules with a subset of the modules reserved for a higher license tier. The fingerprints from all modules are extracted before release. At execution time a PFP monitor matches the execution of different modules with the authorized licenses. When a module not licensed executes, as a result of a stolen password or a breach in protection, the PFP monitor could inform the issuing agency about the violation.

[1196] Furthermore, it is possible to enable a trusted execution-instance-limited leasing approach for protected software. In this case, the PFP monitor keeps count of the number of times the licensed software has been executed and informs the issuing agency when the lease has expired.

[1197] A similar approach can be taken for licensed media content. Using a PFP monitor, it is possible to detect the reproduction of specific files on known media players using PFP. In this case, the protected media data takes the place of the predetermined input during PFP characterization. If the same media is reproduced on the specific player, the power signatures will match. Therefore, PFP can be used to detect the reproduction of unauthorized licensed media.

Failure prediction based on PFP

[1198] Hardware components undergo an inevitable aging process, which is accelerated by operation in harsh environments or when the systems operate under continuous environmental stress. This aging is reflected on the power consumption characteristics of the platform. PFP can be used to monitor not only the correct execution of software but also the integrity of the hardware platforms. A PFP monitor can continuously track the power consumption characteristics of the hardware and predict failure before it actually happens, dictating when a specific system or element must be replaced.

[1199] Tracking the power consumption characteristics in PFP is implemented using an adaptive filter. It is necessary to compensate for differences in power consumption from when the signatures are extracted or due to environmental conditions. The same tracking mechanism can be used to monitor the status of the hardware and compare the power consumption characteristics to predetermined patterns captured in laboratory testing of the devices. The process to identify the failure features is depicted in Figure 36. In this process, accelerated aging 3610 can be achieved by exposing the target device to abrupt temperature changes. The characterization process takes place in intervals, with one round of accelerated aging followed by trace capture 3620 during the execution of a test routine. The traces are collected for posterior analysis and the process is repeated until the device fails. Once the device fails the set of traces is examined to determine the specific features that are exhibited before failure 3630. The features are extracted from other similar devices to provide statistical diversity and isolate the generic features 3640.

Embedding module identification information into synchronization signaling

[1200] The PFP requires proper synchronization with the software being executed in order to provide a correct assessment. There are two levels of synchronization in PFP: clock cycle level and routine level. The former can be easily achieved by tracking the distinct cycles in

the power consumption that occur at the clock cycle rate or, for simple platforms, by probing the clock signal itself. The latter synchronization is more difficult to achieve and the process is facilitated by embedding in the routine itself a trigger, or identifier, that informs the PFP monitor the execution of a specific routine.

[1201] In this section we present a mechanism to embed an identification of the node being executed into the triggering and signaling mechanisms. This mechanism not only helps with informing the PFP monitor which specific routine is about to execute, but also provides robust synchronization signaling to for more accurate anomaly detection and behavioral signature extraction.

[1202] The ultimate goal is to provide an identification code for the different modules being characterized that is inserted into the synchronization and triggering artifacts for PFP. There are two main approaches to provide synchronization and identification signaling for PFP: 1) creating an adjacent physical signal, as shown in Figure 33, and 2) embedding a signal in the power consumption itself, as shown in Figure 34. For the former, a binary identification code is written in a physical 10 register 3324 before the execution of the routine 3322. The register is then transmitted 3335 to the PFP monitor 3340, which captures the power traces 3315 from the sensor 3310, in either a parallel or serial fashion. The length of the code and register depends on the number of routines that need to be monitored. In the simplest sense, a single-bit register, such as an LED, can be used to signal the execution of the target routine. In the case of a separate physical signaling the trigger is encoded as a binary number on the signaling register, as shown in Figure 33.

[1203] The second approach requires the synchronization signaling to be embedded on power consumption itself by inserting a carefully crafted sequence of instructions 3422 that yield a distinctive power consumption pattern. This approach is depicted in Figure 34. The instructions on the synchronization routines are chosen such that the bit transitions in their codeword, addresses, and parameters yield a specific number of bit transitions that ultimately drive the power consumption and signal the PFP monitor 3340 that a specific sequence is about to execute to capture the right set of traces 3415 coming from the sensor. More bit transitions result in higher current drain. When developing the sequence the length and characteristics of the pipeline need to be taken into consideration. Similar to the previous approach, the length of the sequence of instructions (code) depends on the number of critical routines that need to be identified. By creating different distinct power consumption patterns,

the sequence itself is chosen to yield different signatures codes used to identify different modules.

[1204] It is important, to notice that synchronization signaling is a required element for effective PFP, as it allows concentrating the assessment efforts on the sections of the code that matter the most. Embedding an identification code in the signaling facilitates the assessment process but it is not a necessary requirement. This is because using a single trigger will allow the PFP monitor to capture the right set of traces and signal classification techniques can be used to determine which specific routine was executed or whether no reliable match can be established (an anomaly).

Improved PFP monitoring by combining signals from different board elements

[1205] Signals from different elements of the system can be used by a PFP monitor and be combined to provide improved performance and reliability. Sources of multiple signals include multiple processors, co-processors, peripherals, or other special-purpose elements introduced with the sole purpose of enhancing PFP (e.g. the 10 registers used for triggering).

[1206] There are different ways to combine signals from different sources in PFP. One of the main approaches includes capturing power traces from different processors or other digital circuits to perform integrity assessment on multiprocessor and multicore boards. Another approach is to monitor other elements of the systems (power consumption or other side and direct channels) to gather extra context information to be used during integrity assessment. The extra context information can be used to improve synchronization and facilitate behavioral characterization. The context information can be generated as a product of normal system operation or deliberately introduced at design time (e.g. the IO registers used for triggering). A sample setup of a PFP monitor that combines multiple signals is depicted in Figure 35,

[1207] Additional signals can be captured from direct support 10 registers, from the power consumption of different elements, or from other side-channels such as electromagnetic radiation. Combining signals from different sources requires a specially designed detector that can support the different features. The specific combination mechanisms depend on the system functionality and support platform. For example, in a multi-core processor, power traces from each core can be scanned in order to find the traces corresponding to a target routine. Another example, in a software-defined radio, the activation of the power amplifier

(PA) can be detected by monitoring the power consumption and occurs when a radio transmission is taking place. The activation of the PA can be used as a trigger mechanism for the routines involved in preparing the data to be transmitted (notice that in this case, the routines execute before the trigger occurs).

Using malware signatures to enhance PFP performance

[1208] Even though the main application of PFP is anomaly detection, there are important benefits of using available information from known malware to improve assessment performance. When a new trend of malware is identified, it is possible to extract its PFP signature and add it to the library of known signatures. These malware signatures can be used to improve the performance of PFP integrity assessment by providing traditional signature - based detection of installed malware, similar to traditional antivirus software. The monitor would need to be made aware of the individual nature of each signature (white list and black list) in order to avoid incorrect assessments. Malware signatures can also be extracted from behavioral patterns in the execution. For instance, certain malware types, such as exhaustion attacks have very distinct execution patterns that can be easily identified using PFP.

[1209] The process of extracting signatures from malware is similar to the process to extract signatures from trusted software, in which the target modules are executed repeatedly in a controlled environment and different signal processing techniques are applied to the resulting power traces to select the features with the best discriminatory properties. It is important to note that malware characterization is facilitated once the malware has been identified, isolated, and executed in a controlled environment.

Automatic characterization and signature extraction

[1210] In order to efficiently characterize a new software system, or a new version of an existing system, it is necessary to have tools to automatically characterize a trusted reference and extract the PFP signatures that uniquely identify that specific software's execution. In a sense, this process is similar to automated testing because it requires the execution of specific modules under controlled conditions. Different from automatic testing, however, PFP characterization is only concerned with "observing" several execution instances of different modules and does not try to evaluate any requirements or properties.

[1211] The purpose of this section is to describe an approach to facilitate the characterization of complex systems and software architectures and make it feasible to extract signature from realistic implementations of cyber systems of practical complexity. Without this automatic approach, it would take too long to characterize and extract the unique signatures from complex systems (i.e. commercial systems) to be used in power fingerprinting.

[1212] The main goal is to automate the process of characterization for the different modules by using scaffolding similar to what is normally used in software testing, as well as by using a variety of statistical analyses and signal processing to identify the best discriminatory features that form the fingerprints. The process starts when a new software stack needs to be characterized. The tools necessary for this process include: critical module descriptors, signal processing tools for feature extraction, detector design tools, scaffolding for module execution (similar to testing scaffolding), input vector generators, report generation, and signature packaging. In order to facilitate the understanding of the approach, we provide a high-level view of the process that describes the details and interrelationships between the different subsystems. The relationships are depicted in Figure 37.

[1213] Descriptors include required information about the critical modules, including unique identifiers, dependencies, input analysis (breakdown of different input classes), execution mode (dynamically linked, priority, kernel module, etc). Information from the descriptors is used to implement the scaffolds to control the isolated execution of the target modules. The scaffolds allow the system to enter deterministic values as inputs to control the execution of the modules.

[1214] The information in the descriptors about the functionality and the different input, types is used to determine a suitable set of input vectors.

[1215] Coverage analysis is performed to identify the execution paths that have been exercised, yielding a metric of the level of protection for the system.

[1216] Once the system is loaded, the operator (which can be an automatic system) executes the different modules with support from the scaffolds and providing the appropriate input vectors. While the modules are being executed, the PFP monitor captures measurements of the power consumption.

[1217] The power traces captured by the monitor are then processed using different signal processing techniques for extracting discriminatory features. There is a predefined set of features to be extracted for each component in different domains and using different techniques.

[1218] After several traces are captured and the respective features are analyzed, statistical analysis is performed to design optimal detectors to discriminate normal activity from anomalies based on the specific requirements for the application.

[1219] Signatures and detectors are then packaged together to be deployed along with the monitors that will be assessing the integrity of the target systems.

[1220] The following sections include more detailed descriptions necessary to successfully implement the approach described above.

Descriptors

[1221] Descriptors contain meta-information about the specific modules to be characterized. They are used to develop scaffolding artifacts to isolate the execution of individual modules and to provide a controlled environment to exercise the different execution paths.

[1222] The descriptors are expected to be provided in a markup language that is easily read by humans and machines, such as the Extensible Markup Language (XML), but the contents, language, and structure will depend on the specific tools used for automating the characterization process and can be proprietary.

[1223] The minimum required information that needs to be contained in a module descriptor for PFP characterization includes:

[1224] Unique identifiers to describe each module. The unique identifiers should be human readable and provide the necessary information to uniquely locate the module in question. Elements in the human readable part includes Company, Product, Class, Module, and Version,

[1225] Dependencies. The required software and hardware dependencies of the module,

[1226] State Dependencies. The elements of the internal state that affect the behavior of the module and that need to be controlled to provide consistent and deterministic execution.

[1227] Interface analysis. Provides a breakdown of the different input classes and the required input classes to exercise the different execution paths

[1228] Execution mode. Describes in what mode the module will be executed when deployed, i.e. static, for statically linked modules; dynamic, for dynamically linked modules; Kernel or protected mode, for the operation mode that the processor will adopt, when executing the module; and priority level.

Input vector generators

[1229] The function of the input vector generators is similar to their counterparts in software testing, to provide the appropriate inputs to force the component into a specific state sequence that includes the different execution paths. Different from testing, however, the goal for PFP input vectors is not to find implementation errors, but simply to exercise the different execution paths.

[1230] Depending on the nature of the target system, it will sometimes be necessary to store the input vectors and distribute them along with signatures for utilization during assessment (i.e. integrity auditing). The decision whether to keep the input vectors depends on the nature of the features selected and whether traces due to random inputs can be removed.

[1231] The input vectors can be generated using different techniques, including search-based (random search, hill-climbing, genetic algorithm, etc.), partial scan, linear programming, and random and pseudo-random approaches.

[1232] The actual identification of effective test vectors, however, remains a largely heuristic approach that depends on the specific functionality of the target module and its input domain, as well as the information available about the module's structure. There will be some cases when specific knowledge of the execution structure of the module (what execution paths exist and the state sequences necessary to execute them) is necessary to find meaningful input vectors within a reasonable time. Furthermore, sometimes the direct input of an expert human analysis may be required to provide guidance to the automatic tools in order to identify and generate meaningful, effective test vectors.

[1233] One key element of test vector generation for PFP is that the objective is to execute the different paths that are expected to occur once the device is deployed, not to find errors.

[1234] This is a relatively risky approach, because a valid execution state can be reached that has not been characterized and, hence, is flagged as an anomaly. The advantage is that it reduces the search space to only a few states. For most critical systems, the execution space is relatively small and the expected execution states are a subset.

Coverage report

[1235] Using the information from the input vector generator is possible to generate a coverage report based on the execution paths traversed by the specific input vectors. Using structural information from the target modules, it is possible to calculate a PFP coverage metric as a percentage of the existing paths in the module and the ones traversed using the generated input vectors. This report is only an indication of the expected coverage for PFP. The report still needs to be completed by identifying the number of execution paths that actually yielded acceptable PFP signatures.

[1236] The report is provided at the end to provide the users the information about the specific modules that can be monitored using PFP.

Scaffolding

[1237] With the descriptors and the software stack, the process of scaffolding to isolate the execution of the critical modules and their different pieces is performed. This is similar to the process of scaffolding for automated testing. The purpose of scaffolding is to execute the target modules in a controlled environment similar to the one that will be found once the whole system is deployed in order to collect the power traces during its execution. Because the modules are expected to have different execution paths that depend on the inputs, the scaffolds need to facilitate the use of different inputs.

[1238] For the case where physical inputs are needed, the scaffolds need to provide the appropriate physical interfaces to provide the necessary inputs.

[1239] This is a partially manual process and depends on the characteristics of the target modules. Fortunately, most of the elements necessary for the scaffolds overlap in functionality with the traditional scaffolds for automatic testing (e.g. unit, integration, and system testing), adding only a little bit of extra work.

[1240] It is important to note that for non-software implementations, the scaffolding will have similar requirements, although the final implementation will be different. In these cases, the modules will be limited by the sections that can be independently exercised. For highly integrated systems, this may represent a challenge.

Signal processing and feature extraction

[1241] With the power traces corresponding to the execution of the different modules and their individual execution paths captured using the power/instantaneous current sensor, the discriminatory features that uniquely identify the execution of the target module need to be extracted. The exact set of techniques and signal analysis necessary to identify practical signatures depends on the specific characteristics of the target modules.

[1242] We just describe a framework for the parallel execution of a number of different feature extraction and signal processing techniques to reduce the overall time required to characterize a target module.

[1243] There is no known effective procedure to determine the optimal discriminatory features for a given problem. There are, however, several techniques that can be evaluated and from which the best, discriminatory features are selected. The set of discriminatory features that are extracted is determined using a combination of heuristic approaches and experience. Included among these features are: Time-domain correlation, Euclidean distance, cyclostationary analysis, frequency analysis, etc. The process to select the best discriminatory features include calculating all the different features in the set in parallel and ranking them based on intra-class variance. The Mahalanobis distance is a sample metric for such an evaluation of features.

[1244] The process of feature selection and detector design, explained below, are tightly related, as the statistical properties of the feature extraction results determine the necessary analysis to determine an optimal detection threshold.

Statistical analysis and detector design

[1245] Statistical analysis is performed on the different features obtained from the power traces captured during independent execution instances of the target module. The goal of the statistical analysis is to select the features with the best discriminatory qualities and to

determine the threshold levels, or areas within which an observed set of features will be considered as generated by the target module (a detector).

[1246] In PFP, being an anomaly detection approach, the probability of false alarm (PFA) is an important operational metric that determines the performance of the system. PFA is defined as the probability that a normal execution instance of the target module fails outside the acceptance area and is classified as an anomaly. A PFP detector needs to be designed to minimize the PFA while maximizing the probability of correctly identifying the target module. This is a classic hypothesis testing problem and the Neyman-Pearson criterion can be applied to detect a threshold. There are, however, several other approaches that can be applied.

[1247] Given sufficient samples, an arbitrary PFA can be achieved in PFP. In practical systems, however, this is not feasible and a finite, practical PFA level must be determined. The PFA that can be tolerated depends on the specific module and the nature of the application where it is expected to operate.

[1248] Ideally, signatures from different instances from the same execution need to fall within the distance for minimum sensitivity calculated during platform characterization. In the event that this desired feature cannot be achieved, there are several ways to make PFP deliver accurate assessments. One simple approach is to average several traces to get rid of some noise.

Signature Packaging and Encryption

[1249] Once the target modules have been characterized, the resulting signatures, feature extraction techniques, and thresholds are packaged for deployment along with the devices. The packaging and delivering mechanism depends on the characteristics of the device and the application. The complete signatures extracted using the selected features need to be stored and passed to the monitors. For example, in case of simple time-domain correlation, the complete vector needs to be stored.

[1250] In order to protect the signatures at rest or during transportation, it is necessary to encrypt them to avoid giving potential attackers an exact reference of the signatures the monitor is looking for. This encryption can be performed using a variety of mechanisms for private or public key encryption. It is important to note, however, that even if a potential

attacker acquires the signatures, it would be still very difficult to match the signatures perfectly while carrying on malicious behavior.

Secure signature update

[1251] When a deployed system being monitored using PFP is updated, it is also necessary to update the PFP signatures in a reliable and secure way in order to maintain effective integrity assessment. This is a critical step, as the trustworthiness of the assessment depends on appropriate signature management. For this update process to be secure it is necessary to verify the signature's integrity and authenticity. In this section we describe the necessary mechanism to provide secure PFP signature update.

[1252] For widely deployed PFP monitors, signatures should be distributed along with other software updates. For centralized PFP monitors, updates can be delivered separate from the software updates. The main challenge in PFP signature updating is authentication (i.e. making sure the sender is an authorized entity and that the signature itself is correct and has not been tampered with or altered in any way). The challenge is not as difficult in the case of centralized PFP monitors, where signatures can be distributed using physical media or trusted networks and where preparations can be made in anticipation of the transfer of sensitive signatures.

[1253] In the case of widely distributed PFP monitors, where signature exchange cannot be made using physical means or alternative trusted networks, signature update needs to be performed along with the actual software update. In this case, there are several vulnerable points that could be exploited by an attacker with enough knowledge of the PFP system. For example, if not properly authenticated, the update process can be disrupted by a man-in-the-middle attack.

Secure signature update process and operation

[1254] Known approaches for secure content, distribution commonly used in over-the-air programming and distribution of software updates can be adapted to PFP signature updates. Secure signature update can be seen from two different perspectives: the authentic signature originator and the PFP monitor. From the signature generation side, it is necessary to provide effective authentication information along with the signature and encrypt the signature pieces with a rotating key scheme.

[1255] Other techniques that can be applied to enable secure signature update include:

[1256] Scramble both the sample buffer and the signature elements

[1257] Public or symmetric-key encryption

[1258] Change the encryption key to be used for descrambling the signature and traces according to a known sequence (PN Sequence) that updates its index after each signature update.

[1259] Protection against side-channel attacks

[1260] PFP uses the same principles for integrity assessment that are exploited for malicious side-channel attacks. Hence, in order to prevent potential adversaries from exploiting PFP infrastructure to perform side channel attacks, it is necessary to protect the traces provided by the sensor by restricting access to them. This is especially important, when the power traces are transmitted using a wireless connection. This section describes a mechanism to protect the unauthorized access to power traces, which can be misused in side-channel attacks.

Operation

[1261] Protection for trace access is accomplished by encrypting or scrambling the traces using a shared key between the PFP sensor and the PFP monitor. From this perspective, there are two basic modes of operation for PFP: embedded monitor (sensor and digitizer) and external monitor.

[1262] In embedded operations, traces are encrypted or scrambled with a strong private key (symmetrical key encryption). Performing this encryption step is especially important when the power traces are transmitted wirelessly for external processing. The encryption process is described in Figure 38. The analog output of the processor 3810 is monitored by sensor 3820 and converted by an analog-to-digital converter 3830 and fed into the cipher 3850. The cipher 3850 can hide the appropriate information from side-channel attackers in several ways, including block encryption of the bit samples or scrambling them (effectively a transposition cipher where the key is a permutation).

[1263] For external monitors, the physical connection that gives access to the traces is enabled by a digital switch that requires a password. In this case, the contact points for the

external monitor are provided by the power management chip on the platform. The power management chip can be as simple as a voltage regulator, but for most commercial processors used in modern smart-phones, the power management chips are much more complex. When the appropriate monitor is connected, the PFP-enabled power manager reads the password from the external monitor once connected and then reroutes the power supply current to go through the external sensor which allows the external monitor to capture the instantaneous current drain or power consumption. Figure 39 shows a graphical representation of this process.

[1264] It is important to note that the solutions here described are not intended to prevent attackers from carrying on side channel attacks against our target systems. Instead they are intended for preventing the exploitation of PFP monitoring facilities for side-channel attacks. With these measures in place, a potential attacker would have to make the same hardware modifications to a board with PFP monitoring as to one without it.

[1265] While the invention has been described in terms of preferred embodiments, those skilled in the art will recognize that the invention can be practiced with modification within the spirit and scope of the appended claims.

[1266] Further implementations of the PFP system may include: a method for performing real-time integrity assessment of execution of a routine in a computer processing platform, comprising: monitoring execution of the routine by tracing power consumption of a processor by sampling during execution of the routine; using a platform characterization technique that further comprises detecting sections of the traces that display the largest dependence on state transitions in the processor; using said sections to select features carrying the most information; obtaining from a characterization of selected features of the routine contained in said sections a set of trusted power fingerprints of the routine; establishing a threshold for a specific false alarm rate based on the probability distribution of distance from a signature comprised of said trusted fingerprints; comparing a library of said trusted fingerprints to features extracted from traces from the execution of untrusted code determining a distance between said fingerprints and the extracted features; and reporting an exception if the distance exceeds the threshold. In one implementation, the method embodiment further comprising synchronizing said tracing with execution of the routine by embedding module identification information into the routine. In one implementation, the method embodiment further comprises that the module identification information is a binary identification code

written in a register before execution of the routine. In one implementation, the module identification information is a sequence of instructions yielding a distinctive power consumption pattern. In one implementation, the tracing of power consumption combines signals from a plurality of processor circuits. In one implementation, the PFP system further enhances the quality of said exception reporting by adding to said library fingerprint signatures of known malware.

[1267] Further implementations of the PFP system may include a system for performing real-time integrity assessment of execution of a routine in a computer processing platform, comprising: means for monitoring execution of the routine by tracing power consumption of a processor by sampling during execution of the routine; means for using a platform characterization technique that further comprises means for detecting sections of the traces that display the largest dependence on state transitions in the processor; means for using said sections to select features carrying the most information; means for obtaining from a characterization of the selected features contained in said sections a set of trusted power fingerprints of the routine; means for establishing a threshold for a specific false alarm rate based on the probability distribution of distance from a signature comprised of said trusted fingerprints; means for comparing a library of said trusted fingerprints to features extracted from traces from the execution of untrusted code means for determining a distance between said fingerprints and the extracted features; and means for reporting an exception if the distance exceeds the threshold.

[1268] In one implementation, the PFP system may further include a module identification information embedded into the routine, and the module identification information is a binary identification code written in an 10 register before execution of the routine. In one implementation, the module identification information is a sequence of instructions yielding a distinctive power consumption pattern. In one implementation, the means for tracing of power consumption combines signals from a plurality of processor circuits. In one implementation, the PFP system may further comprise means for enhancing the quality of said exception reporting by adding to said library fingerprint signatures of known malware.

[1269] Further implementations of the PFP system may include a method for detecting deviations from authorized execution of software in a digital processor, comprising: observing power consumption of the processor during execution of at least one routine of the software in a known setting; determining a sequence of bit transitions characterizing the

power consumption of said at least one routine; and using the power consumption characterization as a fingerprint for comparison against observed execution of the software to determine deviation.

Enhanced Integrity Assessment Using PFP

[1270] In one embodiment, PFP may be employed to enhance integrity assessment of a target device/system, e.g., industrial control systems, communications systems, a cyber defense system, a cyber supply chain network, and/or the like. In one implementation, the PFP may detect anomalies on the power consumption level of the target system and employ statistical analysis to identify cyber attacks that have affected the software and/or hardware of the target system. For example, when a target entity is being attacked, e.g., a targeting government agency, a financial institution, a military command communicative network, and/or the like, the malware may execute malicious command sequences within the target system, which may be reflected, at the digital circuit level of the target system, as abnormal current patterns. The PFP may detect the current patterns with the target system, and apply statistical methods to analyze the pattern to identify an attack. Within implementations, the PFP may employ mechanisms to disable a target system or some of its parts when an attack is identified, so as to improve the security of the target system.

[1271] In one embodiment, the characterization process to extract the baseline references for use in PFP may be performed by taking measurements from a gold sample of the test platform that is executed in a controlled environment. Different approaches to perform PFP characterization and fingerprint reference extraction and validation tasks that enable more effective characterization of targets or enable characterization in difficult or challenging circumstances are described in Figure 40.

[1272] In one embodiment, a computer model of the target platform can be used to simulate the operation of the target device and estimate the resulting side channels information. For example, a simulation using a cycle-accurate power consumption simulator or a SPICE simulation can be used to estimate the fingerprints of the target device; the process of extracting the fingerprints from a computer simulation may be shown at Figure 40. In one implementation, a cycle-accurate power consumption simulator, e.g., 4002, may be adopted to receive characteristics of a target device model 4001 and input test vectors 4003 (e.g., example testing commands simulating application usage, anomalies, etc.), and generate

estimated power fingerprints, e.g., 4005. In one implementation, the generated estimated power fingerprints 4005, together with statistical models for manufacturing variations 4006 (e.g., standard deviations of power fingerprints of the same target device model, etc.) may be employed for detector design 4007.

[1273] Figure 41 (A-B) provides data plots illustrating example sample power fingerprinting distributions showing outliers or anomalies with different underlying distributions. In one implementation statistical analysis may be used to identify outliers of estimated power fingerprints distribution (e.g., from 4005 in Figure 40). In the absence of a reliable trusted gold reference, the power fingerprints (or baseline references for PFP) can be extracted using statistical analysis for outlier detection. This approach may help when having a relatively large number of test units where the tampering is expected to be present only on a relatively small subset of the samples, as shown in Figures 41A-B. For example, Figure 41A depicts discriminatory features distributions (legitimate distribution 4101 and outlier distribution 4102) in a single dimension 4105; and Figure 41B depicts distribution features (e.g., legitimate distribution 4106 and outlier distribution 4107) in multiple dimensions, e.g., second dimension 4104 versus the first dimension 4103. There are different statistical approaches to detect outliers, including but not limited to the Grubbs method, K-nearest neighbors, local outlier factor, Chauvenet's criterion, and/or the like.

[1274] Figure 42A provides an example logic flow diagram illustrating the process to perform outlier detection in PFP characterization. In one implementation, the process includes collecting side-channel traces (e.g., at 4201) from a complete population of test devices or from a representative sample of the population of test devices, perform PFP feature extraction (e.g., at 4202) on all the collected traces, perform statistical analysis (e.g., at 4203) to further process the traces if necessary, and perform outlier detection (anomaly detection) analysis. The results may indicate whether the captured traces can be used to generate the baseline references or further analysis is required to determine the origin of the outliers.

[1275] Figure 42B provides a logic flow diagram illustrating aspects of the process to use reverse engineering for PFP characterization. In one implementation, reverse engineering may be employed to validate fingerprints from hardware, according to an embodiment. For certain critical systems in which it is suspected that all the samples may contain tampering or in which outlier detection failed to provide conclusive evidence, test target samples may be

selected (e.g., at 4205) to extract characteristics such as PFP references and fingerprints (e.g., at 4206), where the extracted PFP references and fingerprints may be stored at a database (e.g., at 4211). In one implementation, a reverse engineering of target devices may be performed (e.g., at 4207), which is a destructive test to determine whether any intrusion could be found, e.g., at 4208. If no tampering is found, PFP fingerprints can be trusted by being stored as a “whitelist,” at 4210; otherwise, the PFP fingerprints are to be stored as a “blacklist,” at 4209.

[1276] Figure 43 provides example logic flow diagrams illustrating aspects of enhanced response to integrity violations, according to an embodiment. In one embodiment, when a PFP monitor determines an execution violation, intrusion, or anomaly has occurred, the PFP monitor can initiate or perform a response. Different mechanisms in which a PFP monitor can respond in the event an integrity violation or intrusion is detected are described herein. Response mechanisms are intended to prevent, stop, limit, or expose the execution disruption or malicious intrusion in an effort to minimize the damage to the target system. In case of a malicious intrusion the goal is to prevent the attacker from accomplishing his or hers malicious objectives or gather intelligence about the attack’s capabilities, intentions, identity, etc. The appropriate response is application-dependent and may be different between different devices and even different for similar devices used in different applications.

[1277] In one embodiment, a description of the expected response to a given intrusion event is called a policy. A PFP response policy can be described in a number of ways and can be static (hardwired), updateable, or even dynamic. A policy may be stated using any language (formal or not) or machine-readable description mechanism capable of describing event-response relationships, including first-order logic, truth tables, ontology language, etc. As shown in Figure 43A, the PFP monitor 4303 may receive input monitor data from a sensor 4301, where detection data 4304 may be relayed to the response analysis component 4305. For example, after the PFP monitor 4303 has determined an anomaly has occurred, the response policy is analyzed by the response analysis component 4305 to determine the appropriate course of action, the appropriate action is then executed (e.g., at the response execution module 4306), and the action and context and potentially the results are recorded and logged, e.g., at the record and log module 4307.

[1278] In one embodiment, the response analysis module 4305 interprets the type of violation detected by PFP monitor 4303 and, based on the available policy, determines the adequate

response, selects the response module from the pool of available modules, and initiates the response actions at 4306.

[1279] In one embodiment, the type and scope of response may depend on the application itself, as deploying a given response in a target system can have significant consequences to the application or mission of the target system. Different responses are not necessarily mutually exclusive and can be combined when compatible. Different response mechanisms for integrity violations detected by PFP can be implemented as part of the response execution module, including: warning, log, and report.

[1280] For example, warning and logging may be a basic response to an integrity violation detected by PFP. Figure 43B shows an example process to generate an anomaly event report. In one implementation, once the PFP monitor deems a given trace to be an anomaly or violation and triggers the corresponding event (e.g., at 4311), the response mechanism or module gathers all available information (e.g., at 4312) from the PFP monitor, including the offending trace, resulting discriminatory features, confidence levels, context information, and if available the potential module disrupted, and/or the like.

[1281] The response module also collects and records the available system status information (e.g., at 4313), including model and version numbers, time of the day, PFP information, and other useful information, such as environmental context, to build an event report, e.g., at 4314. Once the event report is logged in the system at 4315, the appropriate notifications are sent (at 4316) as specified in the corresponding policy. Different notification mechanisms can be used including: email message, audio/visual alarm indication, voice communications (e.g. cellular telephone call), SMS message, paging, etc.

[1282] Figure 43C provides a diagram depicting an example system to prevent access to critical resources and peripherals, according to one embodiment. In the event the integrity of the execution status cannot be validated or authenticated, this response may prevent access to certain critical resources to prevent the target system from being damaged, exploited, or its information from being leaked.

[1283] In one implementation, the critical resources that can be controlled by this response include, but are not limited to, communications peripherals 4322 (such as network cards, IO ports, Bluetooth ports, near-field communications, USB, etc), physical controls, memory sections, hard-drives, input-output devices, expansion cards, graphics cards, printers, image

scanners, tape drives, microphones, loudspeakers, webcams, digital cameras, and more. For example, the peripherals 4322 may receive monitor data from a processor/target device 4321.

[1284] Controlling access to peripherals 4322 can prevent damage and information leakage, for instance by preventing access to a network interface to an attacker that has breached the security measures and compromised the device attempting to steal sensitive information, or by preventing the attacker from manipulating physical processes under control by the target device 4321.

[1285] For example, when an integrity violation event is detected at PFP, at 4324, to control the access to the different peripherals 4322, the PFP response module 4323 has different options, including temporarily disabling the device by removing the “enable” signal, resetting the peripheral, shutting down the peripheral by cutting power via a switch 4325, or preventing data communications by blocking or buffering data transfers, among others.

[1286] Figure 44 provides example logic flow diagrams illustrating aspects of processing integrity violations event, according to an embodiment. In one embodiment, intelligence and forensic information may be gathered from the attack or attackers. In case of detection of integrity violation event from PFP at 4401, the PFP monitor may record the status (at 4402) of critical components leading to the event, during the event, and after the event, including execution status, power traces, accessible register, memory contents, status of peripherals, physical entities, and/or any other recordable information related to the status of the processor, system, peripherals, and environment.

[1287] In one implementation, the PFP monitor may allocate more resources to the offending device (at 4403) to enhance the PFP capabilities (including more processing resources, extended power fingerprints, extra side-channel measurements, etc.).

[1288] Figure 44B provides a diagram illustrating a system to temporarily disable the affected device for security enhancement, according to one embodiment. For example, affected devices may be temporarily disabled by shutting down, or resetting or preventing operation/computation.

[1289] In case of detecting an integrity violation event from PFP at 4406, the PFP response module 4405 may immediately disable the device by shutting it down (e.g. cutting out the power supply via a switch 4407), forcing a cold reset on it, or stopping the clock to prevent

attackers or unexpected execution from performing actions that can be damaging to the target system, e.g., the processor or the target device 4404.

[1290] In the case of an attacker breaching the system, this response may prevent the device from further operation at the first sign of unauthorized execution and prevent the attacker from installing a permanent presence, such as a back door.

[1291] Figure 45A provides example system using PFP integrity assessment as a gate keeper before execution external actions, according to an embodiment. In one embodiment, the PFP response module 4508 may permanently disable the affected device, upon receiving an integrity violation event 4509 from PFP, e.g., by physically damaging or destroying the device 4501, e.g., kill switch, or erasing critical or sensitive memory contents, or preventing it from booting up. For example, the PFP response module 4508 may use, via switch 4504, a capacitive charge pump 4503 or another circuit capable of delivering a power spike (direct or reverse biased) or electro-magnetic pulse capable of physically damaging or disabling the circuit. The response module 4508 can also use a different circuit, which can also be based on a charge pump 4503, to erase (at 4507) non-volatile memory (e.g., EEPROM/Flash Memory 4506) containing critical or sensitive information.

[1292] To temporarily disable the device 4501, one of the available approaches is to erase boot memory components and then reset the device, which will not be able to start. Without physical damage, the boot up memory can be reprogrammed to make the device operational again.

[1293] Figure 45B provides a block diagram illustrating a system to buffer and access control for enhanced security, according to one embodiment. In one embodiment, a device that verifies the legitimacy and integrity of the code execution that preceded a communication request may be implemented. Illegitimate, malicious, or unauthorized execution/operation can be prevented from extracting information for the target system (e.g., the processor 4511) and communicating the information to the outside world using a communications peripheral 4519 such as (such as network cards, input/output (IO) ports, Bluetooth ports, near-field communications, USB, etc) or impact a physical process using a peripheral. The same principles can be applied to non-communications peripherals that interact with the outside world

[1294] For example, a request 4512 sent by a processor 4511 to access the communications peripherals 4519 may be relayed and/or temporarily stored at a buffer 4516 at a gate keeper 4515. During operation, such requests to the peripherals may be intercepted and buffered, and may be reported to the PFP legitimate execution verification module 4518 that a request has been received, which can be sent to the PFP legitimate execution verification module 4518 for verification. The PFP sensor 4517 may detect the status information of the target device, e.g., the processor 4511, and send a status indication to the PFP legitimate execution verification module 4518, indicating whether the processor 4511 is at a “clear” status. The communications peripheral 4519 can wait for a “clear” signal from the PFP module to verify that the originating routine is a legitimate, or white listed, module before servicing the request.

[1295] Therefore, in this way, when a request is generated by the target device (e.g., the processor 4511), the PFP monitor verifies that the execution leading to the request is legitimate. Only when the legitimacy of the preceding execution is correctly verified by the PFP system, the request is passed to the peripheral 4519. Otherwise, the request is blocked and further actions may be taken. The buffering and access control (e.g., 4515) may be implemented as a separate intermediate device between the requesting device (e.g., 4512) and the peripheral (e.g., 4519), or integrated into the peripheral or the device itself.

[1296] In one embodiment, the PFP system may improve integrity assessment by enhancing PFP with data analytics. For example, data analytics techniques can be applied to enhance PFP integrity assessment, improve detection performance and suggest operational conclusions by leveraging available useful information. In this implementation, the training phase to identify reference power fingerprints is enhanced to include extra information and measurements in addition to power traces. Data analytics may help highlight useful information and finding relationships that provide important clues of the internal operation of the target system.

[1297] In one embodiment, data analytics may be used to perform exploratory analysis, and identify and select the best discriminatory features in PFP including power traces and other measurements and events. Also data analytics can be used during monitoring, performing confirmatory data analysis in support of anomaly detection, integrity assessment, and execution status validation.

[1298] In one embodiment, during characterization, the input from different PFP sensors, along with the inputs from other physical sensors, measurements, signals, events and observations of other parts of the target systems are analyzed using exploratory data analysis techniques to find the discriminatory features that yield the best discriminatory qualities and better identify the unique execution of legitimate or trusted modules in the target systems. Data analytics are used for optimal feature selection and to identify underlying relationships between the internal execution status of the target device and the outside environment that impacts the operation of the target system or is affected or controlled by the target system.

[1299] In one embodiment, data analytics may also be utilized to evaluate the potential consequences of an execution violation and help determine and coordinate the best course of action to respond to the given violation.

[1300] In one embodiment, the PFP system may obtain enhanced performance by integrating alternative external signals and other events into analysis. Other physical or logical signals collected directly or through sensors that relate directly or indirectly to the physical state of the device, the operational environment, the processes controlled, or other events that have a direct or indirect relationship to the execution status of the target device can be included in a PFP analysis to improve the performance of integrity assessment, execution validation, and intrusion detection in PFP.

[1301] In one embodiment, the physical state of the device (e.g., position of a movable part, actuator status, communication protocol status, etc.) can be used to correlate with the power fingerprint and the power trace to provide additional useful information to improve the assessment. The information can be used to establish timing properties and order relationships, or as additional discriminatory information for data fusion, in the execution of the target system and its interaction with other elements of the system and the environment.

[1302] In one embodiment, the extra information can be used in coordination with data analytics. It can also be used to help improving the performance of the PFP monitor by increasing the probability of correctly detecting an intrusion or anomaly, as well as reducing the probability of incorrectly flagging a legitimate operation as tampering or execution violation.

[1303] In one embodiment, the PFP system may include forensics as an output and application and associated processes. The PFP system can provide valuable forensic

information output when applied to digital systems that are suspected of being compromised or attacked, or to malfunctioning systems, as part of a forensic process. Specific information extracted as part of the PFP process or resulting from the PFP process itself can provide useful forensic information. For instance, during a cyber attack to a target system, the attackers may modify critical modules in the target system to allow them access to sensitive information, send it to the attackers, and hide the evidence of the attack in order to avoid detection and prevent the owner from taking corrective measures to fix the security breach. There is much information gained from PFP that can be useful in a forensics process.

[1304] In one embodiment, the PFP system may help determine the modules that are tampered, replaced, modified, or in any way affected by a digital attack or security breach. The PFP system may help identify the presence of stealth and hidden artifacts, including advanced persistent threats, introduced into the target system as part of a malicious attack or unauthorized modification. The PFP system may trigger on execution violations. In other words, the PFP system waits until a crime is committed, but it catches them almost immediately, makes them evident (preventing the attacker from stealth operation), and increases the chance to respond before extensive damage can be done.

[1305] With proper synchronization, the PFP system may target only specific critical modules and reduce the amount of information that needs to be analyzed in forensic analysis. In a forensic investigation and monitoring, the PFP system introduces negligible overhead preventing the attacker and deployed malicious artifacts from being aware of the PFP monitoring operations that would trigger them to abandon their efforts. Thus, the PFP system allows the forensic investigator to see the attackers in their uninhibited behavior and observe their full capabilities.

[1306] In one embodiment, the coverage of the PFP system on the target systems can be expanded by adding processing resources and increased characterization efforts to deliver additional or enhanced power fingerprints. The PFP system may improve assessment performance by including data fusion with other authentication, anomaly detection, or intrusion detection techniques. The intrusion detection capabilities of a PFP monitor can be enhanced when PFP analysis is complemented using data fusion techniques to integrate information obtained as a result from, or during the execution of other authentication, anomaly detection, or intrusion detection techniques. Alternatives authentication, anomaly detection, or intrusion detection techniques that enhance PFP using data fusion include:

network traffic analysis, typing style, user profiling, and other monitoring of system activities.

[1307] In one embodiment, the PFP system may enable intrusion detection based on a supervised learning approach, preventing attackers from training the intrusion detection system (IDS) after deployment to accept their malicious actions, traffic packets or other activities. In supervised learning, the specific features and patterns used as a baseline reference (power fingerprint) to determine when anomalies exist are determined during a pre-characterization stage and are not changed until the system gets updated. This is in contrast to unsupervised learning, which includes a learning mechanism and has the option to update the reference or baseline dynamically in order to compensate for changes on usage or the environment. Unsupervised learning is more vulnerable to evasion attempts and it opens the possibility of the attacker retraining the PFP detectors into thinking the attack itself is part of the normal operation. The supervised approach in PFP, however, requires heavier investments during characterization.

[1308] In the PFP system, the power fingerprint is extracted from the execution of the trusted system, leaving no opportunity to an attacker to re-train a PFP detector after the IDS is deployed.

[1309] The PFP system enables intrusion detection independent from user modeling, which instead relies on actual operation or computation of the target system, such as software execution, which should not change. The PFP system eliminates the need for a user profile or normal traffic profile.

[1310] In one embodiment, the PFP system may generalize PFP by sensing any physical signal related to the instantaneous current drain. The process and techniques of PFP, including the ones here described, can be implemented independently of the type of sensor use to capture the side-channel information. The spectrum of physical sensors that provide input for PFP can be increased to include different types and classes that collect physical observable properties that relate directly or indirectly to the execution status of the target device. In addition to sensors that measure directly or indirectly the instantaneous current drained by a digital device, such as current probes, other sensors than can be integrated into a PFP monitor include: electromagnetic interference (EMI) sensors, near-field sensors, thermometers, etc.

[1311] The PFP system can observe and characterize the fine-grained patterns in the power consumption of digital systems. The patterns are initially captured by a sensor for any observable feature in the environment or the system that is directly or indirectly impacted by the internal execution status of the digital system. Such observable elements are also known as side channels and include: electro-magnetic radiation, temperature, time delays, etc. The PFP system can be applied to a digital circuit to assess its integrity, validate it, or authenticate it, by observing and analyzing any side-channel in the system.

[1312] In one embodiment, sensors may be used by a PFP monitor include, but are not limited to: acoustic and vibration, temperature, electro-magnetic: electric current, magnetic flux, electro-magnetic radiation, near-field radiation, position, distance, angle, speed, acceleration, light and optical, environmental: moisture, humidity, etc, pressure, force, level, and/or the like. These sensors and their signals can be used in PFP in an individual basis or combine them using data fusion or related techniques.

[1313] In one embodiment, the PFP system may enhance PFP performance by integrating multiple power traces from similar execution instances and other readings and measurements to increase diversity. Multiple power traces or readings of the target device can be integrated when executing a specific module or operation and combine them to improve performance metrics such as reducing the probability of false alarm. For certain devices and implementations, collecting multiple traces is straight forward, as the system operation simply has to be observed repeatedly using the same inputs. Examples of such devices include hardware implementations (ASICs) and basic embedded systems commonly used in industrial control systems. For more complex systems, the process to collect multiple traces of the same execution may involve carefully crafted input vectors (to set the target system in the appropriate state) and synchronization mechanisms.

[1314] In one embodiment, because the PFP system uses anomaly detection, its detection effectiveness is also a function of its ability to minimize false positives or false alarms. Because the PFP system approaches intrusion detection from a signal detection perspective, however, it is possible to combine a large number of traces from different execution instances of the critical modules and gradually reduce the variance of the estimates. In other words, a PFP monitor can achieve an arbitrary probability of false alarm (PFA) provided that enough execution instances are captured. To achieve this level of performance in systems that do not implement repetitive deterministic execution, a mechanism can be provided to ensure that

only traces from the execution of a critical module are averaged together. Such mechanism can involve the synchronization and identification of specific execution using intrinsic characteristics of the side channel information or using intentional markers introduced into the target execution with the intent of facilitating detection and identification of specific module execution. Further improvements in synchronization and execution identification can be achieved by PFP if the input to the target system is controlled by the PFP monitor itself to provide even more control over the execution status of the target device.

[1315] In one embodiment, introducing makers and synchronization artifacts into the structure of the target modules can enable monitor synchronization with the execution status and allow proper alignment of traces. These slight modifications are introduced without affecting the functionality or logic of the module itself. There are two basic ways of introducing PFP markers. First, a physical signal using available peripherals can be used. Second, carefully crafted sequences of instructions can be executed before the critical modules. These embedded routines yield known power consumption patterns.

[1316] In one embodiment, the markers and synchronization artifacts serve a dual purpose as they are also used to encode identification numbers for the specific module, which allows proper combining with traces from previous execution instances from the same module. Embedding identification codes into the markers can be accomplished with both, physical and preamble, signaling approaches.

[1317] One example approach to reduce the variance of the estimates, and thus the probabilities of making a classification error, is to increase the time-bandwidth product of the observation. Increasing the time-bandwidth product allows observing the intended signal for a longer time and reduces the accuracy on the estimates. In the case of the PFP system, the performance improvement comes from averaging power traces from multiple execution instances of the selected modules. In the PFP system, encoded identification data embedded into signaling artifacts allow aligned traces from multiple executions of the target modules to be stored and then averaged together before making a decision of whether the target modules have been tampered. By averaging multiple traces together, random noise is averaged out and a cleaner version of the power traced that highlights the identity of the underlying processes (software) is used to make an assessment. The more traces from the execution of a target module that are combined, the smaller the variance of the resulting estimates, hence, the smaller the chances of making a classification error. Using this approach, the sensitivity of

the PFP monitor can be adjusted to average more or less traces depending on the performance requirements of the specific application. To achieve such performance, enough traces from the target modules are to be collected. Using the Neyman-Pearson criterion, detectors can be designed that achieve arbitrary sensitivity in terms of probability of false alarm.

[1318] In one embodiment, the PFP system may apply PFP for supply-chain validation, assessment, and risk management. The PFP system can be used to detect tampering at the supply chain. Malicious intrusions and tampering across all levels of the execution stack, from hardware to application level, introduced throughout the supply chain can be detected. PFP can be directly applied to detect unauthorized modifications and other tampering in software, firmware, and hardware introduced by untrusted links in the supply chain throughout the system's lifecycle. The approach can be implemented in a location designated as an integrity lab, in which new shipments, or deployed devices that need to be analyzed, are inspected for potential tampering. The PFP monitor is operated with help of specialized harnesses and tools that provide convenient access to the hardware to sense the power consumption and control the inputs to the modules of the deployed devices. Connecting the harnesses may require small modifications to the deployed (untrusted) device to expose power rails or access test points. Note that modification is not necessary if the module's design includes these test points. Even when this is not the case, integrity analysis using PFP can be performed by using an external sensor such as electromagnetic or near field sensor.

[1319] One aspect in performing supply-chain analysis with PFP is executing the untrusted device under a controlled environment. This controlled environment includes predefined inputs that force a specific state sequence and, for programmable devices, the specific software to be executed. For most systems, support scaffolding is developed to control and isolate the execution of specific components. The specific input vectors depend on the functionality of the device or software module and they are expected to exercise the critical execution paths for the operation of the device. The same input vectors used to extract the signatures need to be used to assess the integrity of the untrusted devices. Note that the process of generating inputs and harnesses is closely related to automatic unit testing in software development and can be, for the most part, automated.

[1320] In one embodiment, because of slight process variations during manufacturing, different devices will show different power consumption characteristics. These variations in the power consumption are to be compensated before feature extraction to avoid erroneous

assessments. This compensation can be performed by an adaptive filter whose taps are dynamically modified to match the specific characteristic of the power consumption traces. This adaptive filter allows the PFP monitor to concentrate on the power consumption resulting from bit transitions in the device register during execution and eliminate differences in the traces due to manufacturing variations. Note that the power consumption characteristics of a device also depend on the environment (e.g. temperature) and, more importantly, on the aging experienced by the device. This dependency on age can be used as an indicator of counterfeit elements.

[1321] Another aspect for effective supply-chain trust analysis using PFP is the availability of baseline references. There are different potential sources for such signatures (baseline references). The best reference would be provided by an identical trusted implementation (a gold standard). When a trusted implementation is not available, a reference signature can be extracted using alternative methods with varying degrees of effort and trustworthiness. For instance, two relatively simple alternative reference sources include a previous implementation of the device (one that has been time tested) or an alternative implementation from a different supplier. In these cases, the references are extracted from the execution of the alternative implementations, reducing the chances of two identical modifications by different providers. Power fingerprints from the former approach can miss unidentified modifications present in the previous version. In the latter approach, an attacker could create an identical modification in both versions from the different suppliers to avoid detection.

[1322] In one embodiment, it is also possible to use a computer model or computer-aided design (CAD) model to obtain the power fingerprints. Although this approach typically involves more effort, it can be done in-house without relying on foreign foundries. To extract the power fingerprints using a CAD model, the execution of the device can be simulated using deterministic input vectors. The simulator should be accurate in terms of power consumption to the register transfer level. The PFP system can be used to detect the introduction of counterfeit components throughout the supply chain. Following the same principles of characterization to detect execution anomalies and tampering, the way a device consumes power while aging can be characterized and when a device is reaching the end of its usable life can be determined. To achieve this, it is possible to accelerate the aging process of a target device by exposing it, among other means to wide temperature changes. With this

information, a PFP monitor can identify when a specific component is being stressed and likely to fail, which presents a strong indication of counterfeit elements.

[1323] In one embodiment, the PFP system may enhance the IDS using PFP. For example, the PFP system can provide an alternative, complementary intrusion detection approach capable of detecting local tampering and intrusions. The PFP system enables intrusion detection that does not rely on network traffic analysis, but instead on monitoring a physical side-channel to perform integrity assessment of critical modules and use anomaly detection to catch malicious intrusions.

[1324] In one embodiment, the proposed PFP Intrusion Detection System (PFP IDS) is capable of detecting sophisticated Advanced Persistent Threats (APTs) that effectively mask their network traffic or that minimize their network activity to avoid displaying identifiable external behavior. In known systems, detecting APT is extremely difficult when relying solely on network intrusion detection systems. Sophisticated APTs are known to remain hidden inside a compromised system, effectively avoiding signature-based detection, erasing their tracks, and limiting their network traffic. For instance, an APT that only listens for specific inputs or events to trigger the actual attack without producing any network traffic would typically not be detected by traditional network IDS systems until the attack gets triggered even if the intrusion had existed for an extended period of time.

[1325] Under the PFP IDS system described herein, such stealth intrusion that remains hidden and listens but without producing network traffic, will impact the execution of the system to listen to an input. In this case, the change in execution by monitoring an input would be detected by the PFP monitor that is checking for execution integrity.

[1326] In one embodiment, the PFP system can monitor the integrity of critical modules commonly used by APTs to disrupt operations or to hide the presence of malicious content such as Kernel rootkits or hardware trojans. The PFP monitor sensor object can be operated by itself in an independent fashion, or networked to provide integrity status updates to a central server. More importantly, a PFP monitor can be used along traditional IDS systems, including network-based approaches, and other traditional security mechanisms.

[1327] PFP IDS allows the setting of an arbitrary probability of false alarm on detecting execution tampering of specific modules all across the execution stack, from hardware Trojans to Kernel Rootkits. It also allows the detection of zero-day attacks and allows the

detection of threats that use alternative networks, such as SMS, Bluetooth, or other side-channels, to ex-filtrate data or contact a command and control server.

[1328] In one embodiment, the PFP system may apply PFP to industrial control and Supervisory, Control, and Data Acquisition (SCADA) systems. Industrial control systems (ICS), such as Programmable Logic Controllers (PLCs) are computer-controlled systems that monitor and control industrial processes, from manufacturing processes to nuclear reactors. ICS, when controlling large-scale processes, are also known as SCADA systems. ICS and SCADA systems are ubiquitous in critical infrastructure, including water treatment and distribution, transportation systems, oil and gas pipelines, electrical power transmission/distribution, wind farms, defense systems, and large communication systems.

[1329] In one embodiment, the PFP system can be applied to ICS and SCADA systems to assess integrity, detect malicious intrusion, detect unauthorized modifications, detect tampering, etc of individual components of the system. The application of PFP to ICS and SCADA systems can be performed in different configurations:

[1330] Figure 46A provides an example system to illustrate aspects of a PFP solution for spot-testing SCADA deployments, according to an embodiment. For example, as shown at Figure 46A, a SCADA system 4601a may employ a stand-alone PFP monitor (e.g., 4602) that can be brought to the target device for spot testing, which can be performed periodically or at random times.

[1331] Figure 46B provides an example system to illustrate aspects of distributed SCADA systems with round-robin monitoring, according to one embodiment. For the configuration that performs the PFP processing in a central location (e.g., at 4606), multiple PFP sensors 4603a-c may be employed by distributed SCADA systems 4601b-d, and the interface between the PFP sensor (4603a-c) and the PFP central processing module (4606) can be implemented using a variety of communication approaches, including wired networks such as Ethernet or wireless such as WiFi.

[1332] The PFP sensor (4603a-c) and a PFP preprocessing module (4606) are collocated with the target embedded platform. In this case, the PFP preprocessing module 4606 includes certain processing capabilities to digitize the signal, perform preprocessing to reduce communication bandwidth usage, and send the resulting data to the remote PFP monitor using a network interface. For distributed SCADA systems, the PFP central processing

module 4606 may conduct a round-robin monitoring, e.g., each device may be assessed several times a day, or continuous parallel monitoring if enough communications and processing resources are available.

[1333] Figure 46C provides an alternative example system to illustrate aspects of around-the-clock monitoring using embedded PFP monitors, according to one embodiment. An alternative distributed processing approach mimics a technique used in cellular base stations called fiber-to-the-antenna or RF-over-Fiber. The approach involves an optical transceiver coupled to optical fiber to transfer signals captured by the antenna (sensor) to the central processing base station without degradation in signal quality and with minimal signal processing performed at the remote antenna site.

[1334] In one embodiment, a fiber-to-the-sensor (FTTS) approach can be used. The FTTS approach can be very helpful in certain embedded systems deployments and environments. Industrial control and SCADA systems provide a great application for this approach due to the monitoring of multiple processing modules that are distributed in relatively close spatial proximity but in potentially noisy environments. For example, embedded PFP monitors 4605a-c may be employed by distributed SCADA systems 4601e-g, respectively, to perform around-the-clock monitoring using the embedded PFP monitors. Furthermore, factory floors are often wired and configured to provide a central processing station. It is important to note that while we use optical fiber to describe this embodiment, the approach is not limited to optical fiber and can be implemented using alternative technologies.

[1335] Figure 46D provides a diagram illustrating infrastructure of embedded boards and the PFP monitor, according to one embodiment. For example, the PFP monitor 4615 may interface with one or more distributed embedded boards 4611a-b, each contains an embedded processor 4612a-b, a PFP sensor 4614a-b, and PFP signal modulator 4613a-b. In one implementation, the PFP sensors 4614a-b may send detection signals to the PFP signal modulators 4613a-b for signal modulation for communication with the PFP monitor 4615. For critical infrastructure systems that keep spare parts for their deployed elements, those spare parts can be used for characterization of the different elements off-line. Often, critical IDS and SCADA systems have a test setup that mirrors the actual operational systems, with identical equipment and configuration. Such mirror setup should also be used for PFP characterization when available.

[1336] Figure 47 provides an example block diagram illustrating aspects of integrating PFP into a device charger or sleeve, according to an embodiment. In one embodiment, the PFP monitor 4702/4704 may be integrated into a device charger 4703 or sleeve 4701 to perform integrity assessment, intrusion detection, or unauthorized modifications, e.g., in the form of a PFP monitoring sleeve or charger to assess the integrity of portable embedded devices. The PFP attachment for hand-held devices allows for PFP monitoring in platforms that present constrained size implementations and provide little room for a PFP monitor. For PFP monitor 4704 at the charger 4703, the PFP integrity assessment is executed while the device is being charged. Both, the charger or sleeve options, include a stand-alone PFP monitor, including the sensor, digitizer, processor, and storage for PFP references. Both sleeve and charger implementations have the ability to monitor the target device in a passive or active way. In the passive way, the PFP monitor simply assesses the integrity of the execution that is normally executing on the target device, without extra interaction with the device. In the active mode, the PFP monitor interacts with the target device to force it to execute specific routines or take specific actions to assess the integrity of specific actions and tasks.

[1337] In one embodiment, the PFP monitor has multiple options for effective sensing including, for example, using an electro-magnetic (EM) sensor (near or far field). The EM sensor has a configurable location to help locate the specific sensor placement that yields the strongest and cleanest side-channel signal.

[1338] Both the sleeve 4701 and the charger 4703 can be made in different form factors and with different physical specifications to meet the requirements of the target system. The PFP module integrated into the sleeve or charger can include a full PFP solution, including sensor, digitizer, and processing, for local PFP assessment or a PFP subset for distributed PFP where the captured PFP traces, directly or after some processing and conditioning, are sent to a different location using a network for full PFP processing.

[1339] For the sleeve 4701, the PFP monitor 4702 can include an independent battery for independent operation without impacting the battery life of the portable target device. Both, the sleeve and charger include a separate memory for trusted signature storage.

[1340] In a further embodiment, the PFP system may introduce PFP watermarks to enhance PFP assessment performance, which includes a predetermined power consumption execution, done by executing a known set of code or execution of a sequence of hardware components.

The watermark may be used for synchronization for determining what segment of code has executed, is executing or will shortly be executed or as timing markers to determine if added execution time has been inserted into the device due to the insertion of malware or Trojans between two power consumption fingerprint watermarks.

[1341] In one embodiment, PFP watermarks can also be introduced by triggering, initiating, or requesting specific events or actions on the target system. Such watermarks can also appear intrinsically as part of the normal operation of the system. Turning on a given peripheral with known power consumption characteristics can provide a reliable watermark. For instance, erasing an EEPROM memory may introduce a strong watermark that can be used for synchronization; a similar behavior can be observed when turning on a wireless card. When watermarks are introduced by events and tasks from outside the processor, they may also relax the measurement requirements on the PFP system. For example, the watermark could be observed from a test point early on the power conditioning circuitry, on the main battery, or even using an electro-magnetic probe.

[1342] In further embodiments, for example, a method for assessing integrity, detecting malicious intrusions, detecting unauthorized modifications, and detecting tampering in a system that is at least one of a Industrial control system (ICS) or a Supervisory, Control, and Data Acquisition (SCADA) system includes monitoring the execution of individual components of the system by observing power consumption of the individual component and applying PFP to the individual components.

[1343] In some embodiments, a portable PFP monitor can be used to perform spot assessment on the individual components of the system.

[1344] In some embodiments, the PFP monitor can be distributed across the system with a sensor and preprocessing module collocated with the individual components. The method can include sending partially processed traces using a network to a centralized location to complete PFP processing and assessment.

[1345] In some embodiments, the PFP monitor can be distributed across the system with a sensor collocated with the individual elements. The method can include directly sending an output of the sensor, after modulation, using a fiber optics link or network to a centralized location for PFP processing and assessment.

[1346] In some embodiments, the PFP monitor can be embedded with the individual components of the system such that continuous monitoring of the individual components can be performed.

[1347] Further implementations of the PFP system may include: a method for detecting unauthorized modifications and other tampering in software, firmware, and hardware introduced by untrusted links in the supply chain can include the application of a PFP analysis in an integrity lab for incoming devices or on the field for deployed devices; a method for responding to malicious intrusions, tampering or unauthorized modifications detected by a PFP monitor can include at least one of: warning, logging, and reporting the event; preventing access to certain critical resources to prevent the target system being protected from being damaged, exploited, or its information from being leaked; gathering intelligence and forensic information from the attack or attackers; temporarily disabling the affected device by shutting down, or resetting or preventing operation/computation; or permanently disabling the affected device by physically damaging or destroying the device.

[1348] A further implementation of the PFP system may include a method for PFP characterization and fingerprint reference extraction and validation that enables more effective characterization or enables characterization of targets in difficult or challenging circumstances can include at least one: using simulation to get the power consumption of a device; using statistical analysis to identify outliers; using reverse engineering to validate fingerprints from hardware.

[1349] It is intended that the systems and methods described herein can be performed by software (executed on hardware), hardware, or a combination thereof. Hardware modules may include, for example, a general-purpose processor, a field programmable gate array (FPGA), and/or an application specific integrated circuit (ASIC). Software modules (executed on hardware) can be expressed in a variety of software languages (e.g., computer code), including C, C++, Java™, Ruby, Visual Basic™, and other object-oriented, procedural, or other programming language and development tools. Examples of computer code include, but are not limited to, micro-code or micro-instructions, machine instructions, such as produced by a compiler, code used to produce a web service, and files containing higher-level instructions that are executed by a computer using an interpreter. Additional examples of computer code include, but are not limited to, control signals, encrypted code, and compressed code.

[1350] Some embodiments described herein may be implemented by a stored program component that is executed by a CPU. In one embodiment, the stored program component incorporates any and/or all combinations of the aspects of the PFP monitoring and enhanced integrity assessment discussed in the previous figures and/or specification. As such, the stored program component affects accessing, obtaining and the provision of information, services, transactions, and/or the like across various communications networks, which enables access of information between nodes may be developed by employing standard development tools and languages such as, but not limited to: Apache components, Assembly, ActiveX, binary executables, (ANSI) (Objective-) C (++), C# and/or .NET, database adapters, CGI scripts, Java, JavaScript, mapping tools, procedural and object oriented development tools, PERL, PHP, Python, shell scripts, SQL commands, web application server extensions, web development environments and libraries (e.g., Microsoft's ActiveX; Adobe AIR, FLEX & FLASH; AJAX; (D)HTML; Dojo, Java; JavaScript; jQuery(UI); MooTools; Prototype; script.aculo.us; Simple Object Access Protocol (SOAP); SWFObject; Yahoo! User Interface; and/or the like), WebObjects, and/or the like.

[1351] In one embodiment, the PFP system may employ a cryptographic server to encrypt and decrypt communications. The PFP system may communicate to and/or with other components in a component collection, including itself, and/or facilities of the like. Most frequently, the PFP system communicates with a database, operating systems, other program components, and/or the like. The PFP system may contain, communicate, generate, obtain, and/or provide program component, system, user, and/or data communications, requests, and/or responses

[1352] Some embodiments described herein relate to a computer storage product with a non-transitory computer-readable medium (also can be referred to as a non-transitory processor-readable medium) having instructions or computer code thereon for performing various computer-implemented operations. The computer-readable medium (or processor-readable medium) is non-transitory in the sense that it does not include transitory propagating signals per se (e.g., a propagating electromagnetic wave carrying information on a transmission medium such as space or a cable). The media and computer code (also can be referred to as code) may be those designed and constructed for the specific purpose or purposes. Examples of non-transitory computer-readable media include, but are not limited to, magnetic storage media such as hard disks, floppy disks, and magnetic tape; optical storage media such as

Compact Disc/Digital Video Discs (CD/DVDs), Compact Disc-Read Only Memories (CD-ROMs), and holographic devices; magneto-optical storage media such as optical disks; carrier wave signal processing modules; and hardware devices that are specially configured to store and execute program code, such as Application-Specific Integrated Circuits (ASICs), Programmable Logic Devices (PLDs), Read-Only Memory (ROM) and Random-Access Memory (RAM) devices.

[1353] While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Where methods and steps described above indicate certain events occurring in certain order, the ordering of certain steps may be modified. Additionally, certain of the steps may be performed concurrently in a parallel process when possible, as well as performed sequentially as described above. Although various embodiments have been described as having particular features and/or combinations of components, other embodiments are possible having any combination or sub-combination of any features and/or components from any of the embodiments described herein.

Claims

What is claimed is:

1. A method, comprising:

receiving, at a first module, side-channel information of a first target component of a system, the first module being collocated with the first target component, the side-channel information of the first target component being associated with the plurality of authorized execution statuses of the first target component and an execution status of the first target component;

obtaining a power fingerprint for the first target component based on the side-channel information for the first target component, the power fingerprint for the first target component representing a plurality of authorized execution statuses of the first target component;

receiving, at a second module, side-channel information of a second target component of the system, the second module being collocated with the second target component, the side-channel information of the second target component being associated with the plurality of authorized execution statuses of the second target component and an execution status of the second target component;

obtaining a power fingerprint for the second target component based on the side-channel information for the second target component, the power fingerprint for the second target component representing a plurality of execution statuses of the second target component; and

sending, from a processor module, a reporting signal based on at least one of (1) the power fingerprint for the first target component and an execution status of the first target component, or (2) the power fingerprint for the second target component and an execution status of the second target component, the reporting signal associated with at least one of the execution status of the first target component or the execution status of the second target component.

2. The method of claim 1, wherein the processor module is physically separate from the first target component and the second target component.

3. The method of claim 1, wherein:

the power fingerprint for the first target component is obtained by the first module;

the power fingerprint for the second target component is obtained by the second module;
the method further comprising:

receiving, at the processor module from the first module, a signal representing the power fingerprint for the first target component; and

receiving, at the processor module from the second module, a signal representing the power fingerprint for the second target component.

4. The method of claim 1, further comprising:

receiving, at the processor module from the first module, a signal representing the side-channel information of the first target component; and

receiving, at the processor module from the second module, a signal representing the side-channel information of the second target component,

the power fingerprint for the first target component is obtained by the processor module,
the power fingerprint for the second target component is obtained by the processor module.

5. The method of claim 1, wherein the reporting signal represents at least one of an integrity assessment, a malicious intrusion detection, an unauthorized modification detection or a tampering detection, for at least one of the first target component or the second target component.

6. The method of claim 1, wherein the first target component and the second target component are included within at least one of (1) an Industrial Control System (ICS), (2) a Supervisory, Control and Data Acquisition (SCADA) system, or (3) an embedded compute platform.

7. The method of claim 1, wherein the processor module are operatively coupled to the first module and the second module through at least one network.

8. The method of claim 1, wherein the obtaining the power fingerprint for the first target component, the obtaining the power fingerprint for the second target component, and the sending the reporting signal are repeated substantially continuously.

9. A method, comprising:

sending a predefined input to an untrusted electronic device, the predefined input being defined based on a functionality of the untrusted electronic device;

receiving side-channel information of the untrusted electronic device in response to the predefined input;

obtaining a power fingerprint for the untrusted electronic device based on the side-channel information; and

assessing an integrity of the untrusted electronic device based on the power fingerprint.

10. The method of claim 9, wherein the sending the predefined input includes sending the predefined input to the untrusted electronic device such that the untrusted electronic device is put into a specific state sequence or software of the untrusted electronic device is executed, in response to the predefined input.

11. The method of claim 9, wherein the sending the predefined input includes sending the predefined input to the untrusted electronic device such that a first execution path of the untrusted electronic device is operational in response to the predefined input and a second execution path of the untrusted electronic device is not operational in response to the predefined input.

12. The method of claim 9, wherein the untrusted electronic device is a first untrusted electronic device, the method further comprising:

repeating the sending, the receiving, the obtaining and the assessing for each untrusted electronic device from a plurality of electronic devices.

13. The method of claim 9, wherein the untrusted electronic device is a first untrusted electronic device, the method further comprising:

repeating the sending, the receiving, the obtaining and the assessing for each untrusted electronic device from a plurality of electronic devices,

the assessing including identifying an unauthorized modification of the untrusted electronic device based on statistical analysis of at least a portion of the plural of electronic devices.

14. The method of claim 9, wherein the untrusted electronic device is a first untrusted electronic device, the method further comprising:

for each untrusted electronic device from a plurality of electronic devices, performing the following:

sending a predefined input to that untrusted electronic device;

receiving side-channel information of that untrusted electronic device in response to the predefined input;

dynamically modifying the side-channel information of that untrusted electronic device such that manufacturing variability is reduced to produce a modified side-channel information;

obtaining a power fingerprint for that untrusted electronic device based on the modified side-channel information; and

assessing an integrity of that untrusted electronic device based on the power fingerprint.

15. The method of claim 9, wherein the assessing includes identifying an unauthorized modification of the untrusted electronic device when the untrusted electronic device has been subjected to unauthorized modification.

16. An apparatus, comprising:

a power fingerprint monitor module configured to receive a signal representing an unauthorized access of a device; and

a response analysis module operatively coupled to the power fingerprint monitor module, the response analysis module configured to select a response module from a plurality of response modules in response to detection of the unauthorized access, a first response module from the plurality of response modules configured to collect information associated with the device during the unauthorized access, a second response module from the plurality of response modules configured to prevent access to the device in response to the unauthorized access, a third response module from the plurality of response modules configured to disable at least a portion of the device subjected to the unauthorized access.

17. The apparatus of claim 16, wherein the first response module is configured to send a report signal based on information associated with the device during the unauthorized access.

18. The apparatus of claim 16, wherein the plurality of response modules includes a fourth module configured to collect information associated with at least one of the unauthorized access or a device that triggered the unauthorized access.

19. The apparatus of claim 16, wherein the third response module is configured to temporarily disable the portion of the device including at least one powering down the device, resetting the device or preventing operation of the portion of the device.

20. The apparatus of claim 16, wherein the third response module is configured to permanently disable the portion of the device including at least one of (1) physically damaging the portion of the device, (2) destroying the portion of the device, (3) setting the portion of the device into an irrecoverable logic state, or (4) permanently erasing a memory of the device.




Figure 1

PRIOR ART




Figure 2




Figure 3

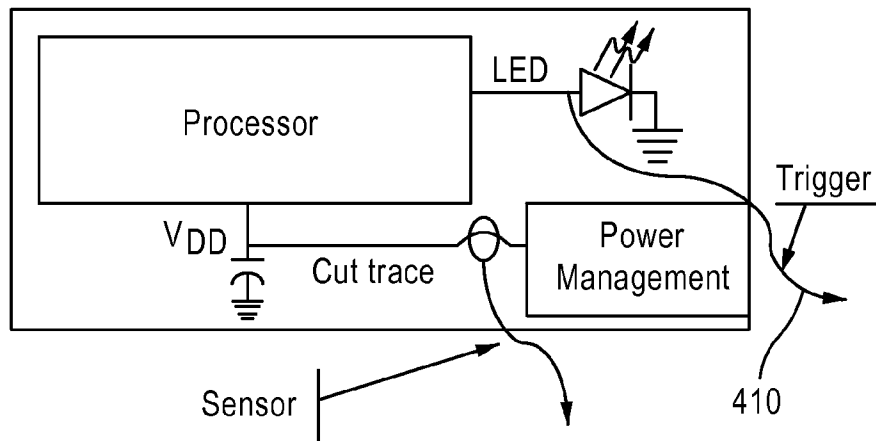


Figure 4

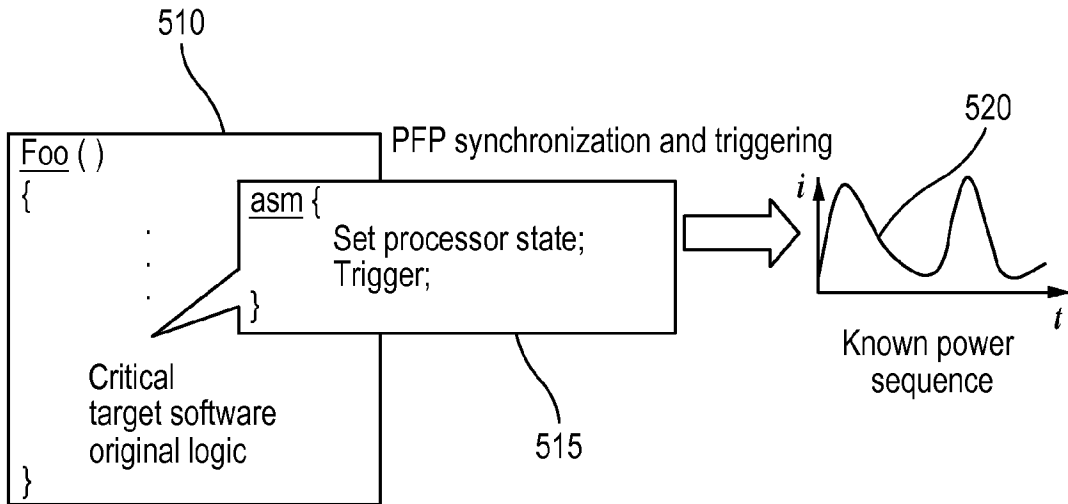


Figure 5




Figure 6

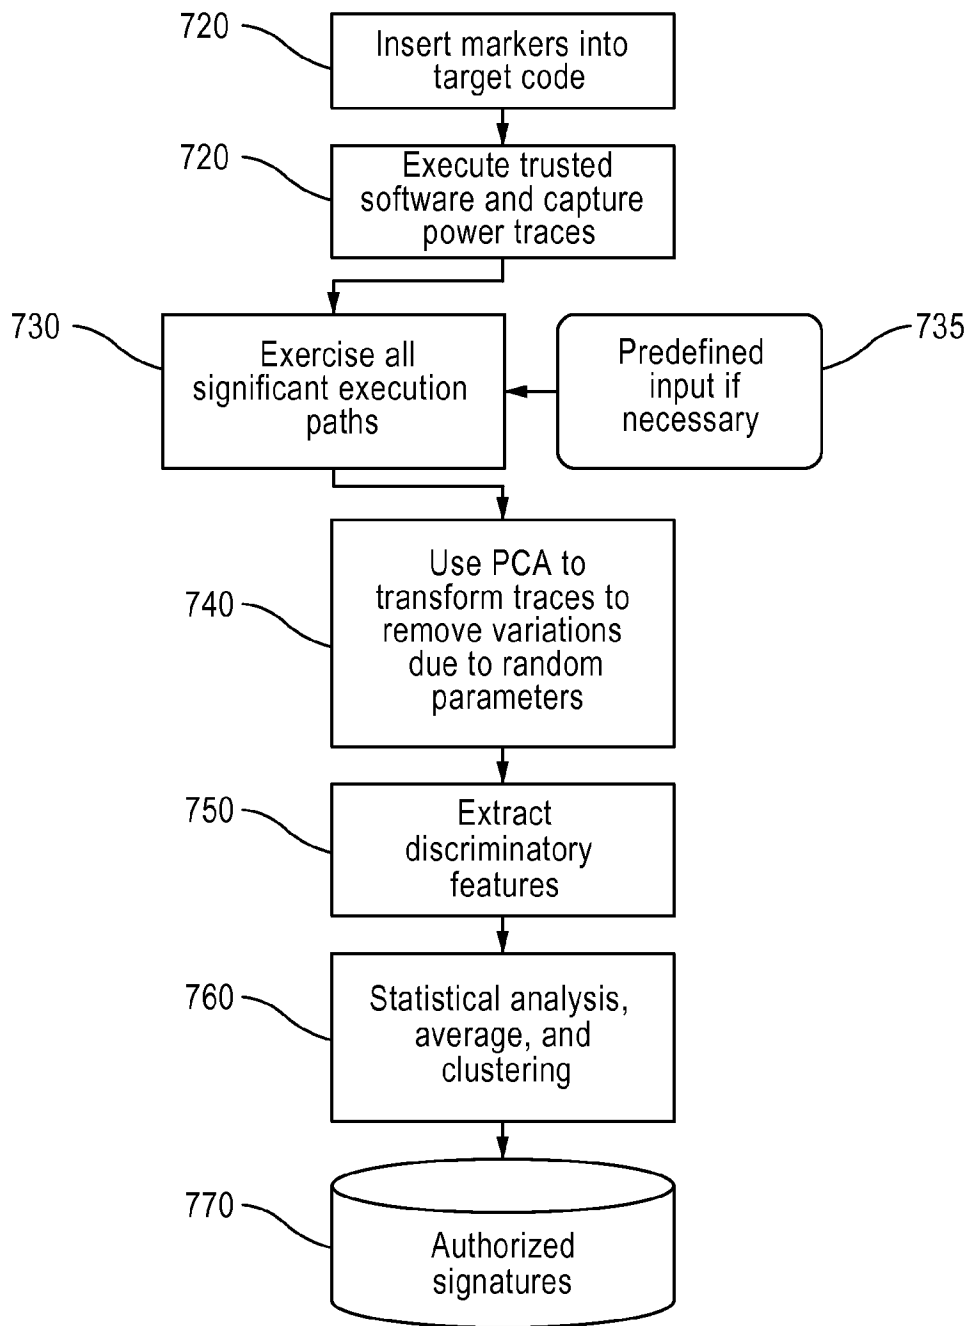


Figure 7




Figure 8




Figure 9

7/34




Figure 10

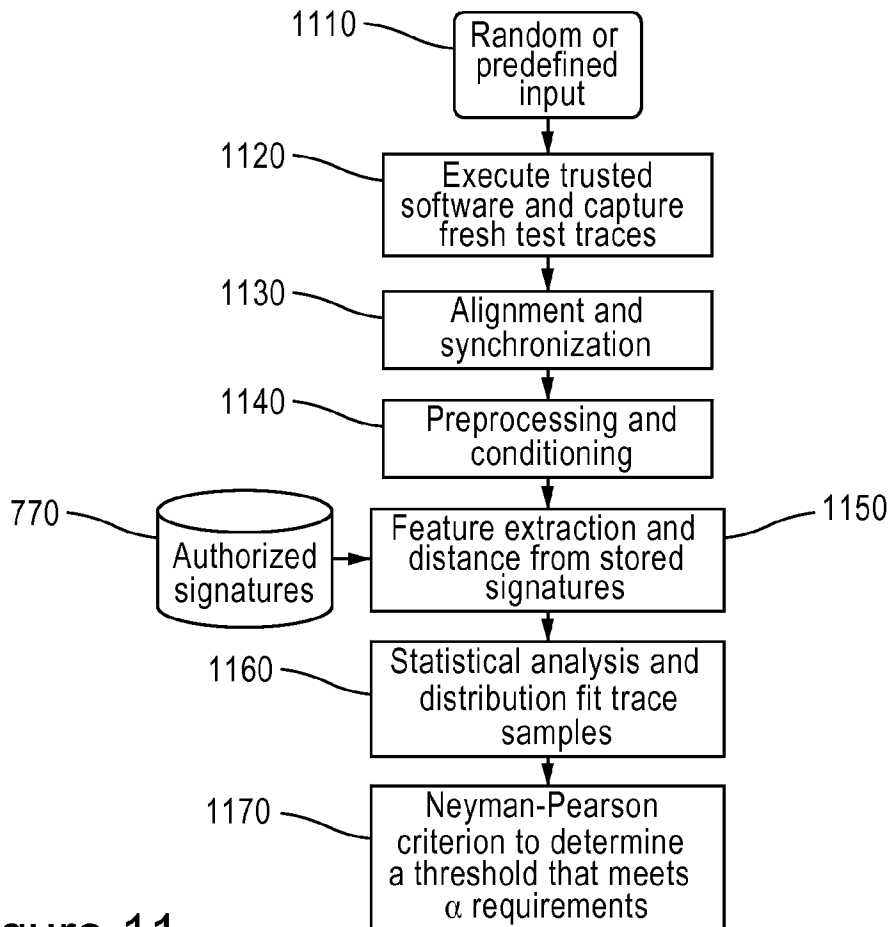


Figure 11




Figure 12

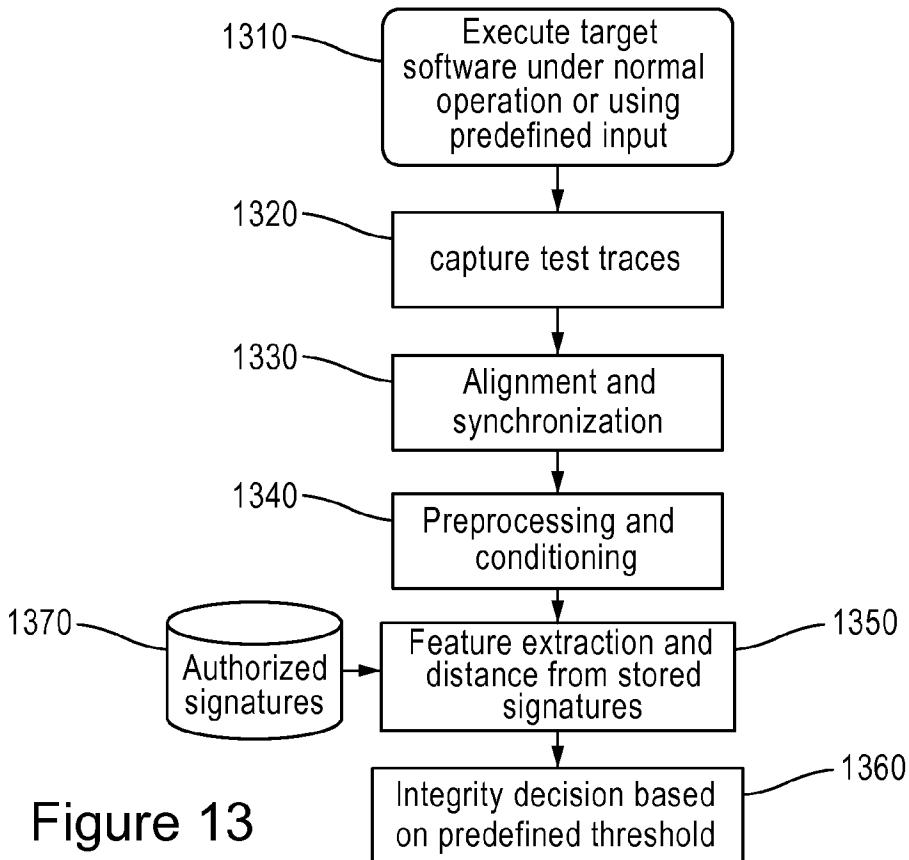


Figure 13




Figure 14

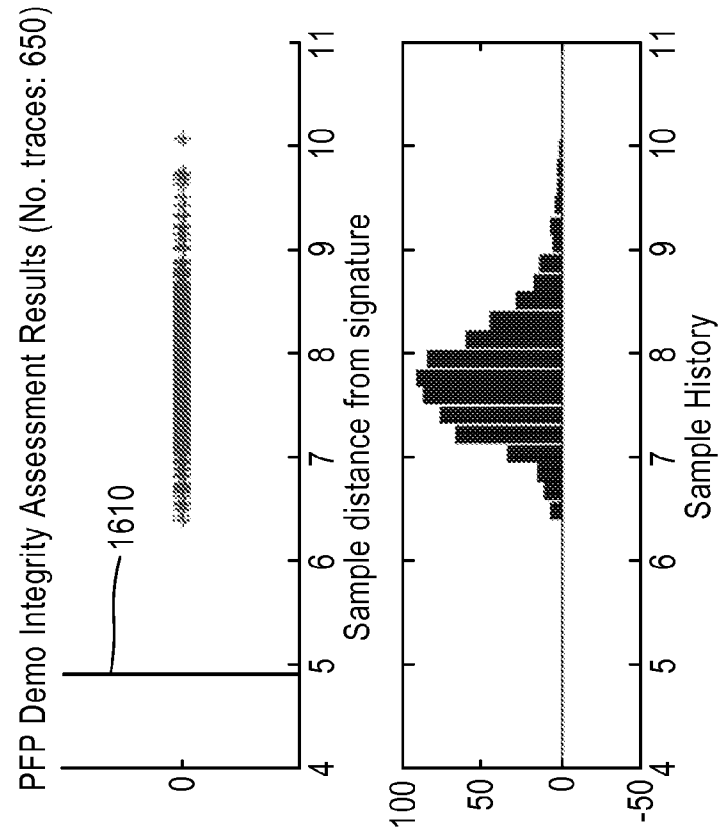


Figure 16




Figure 15




Figure 17

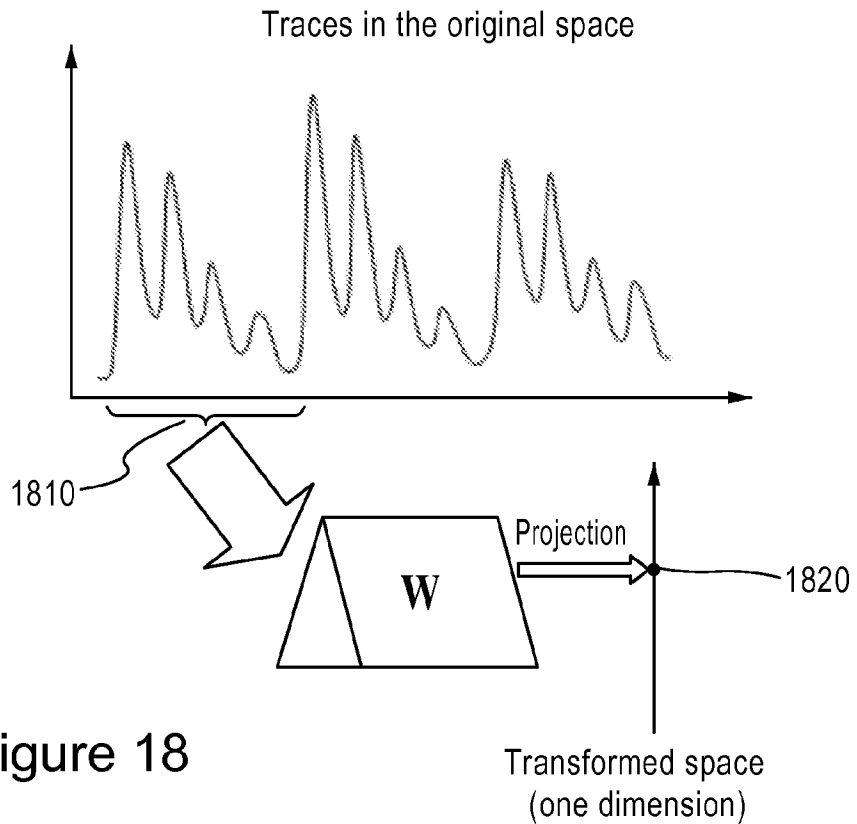


Figure 18

12/34




Figure 19




Figure 20




FIG.21




Figure 22

15/34




Figure 23




Figure 24

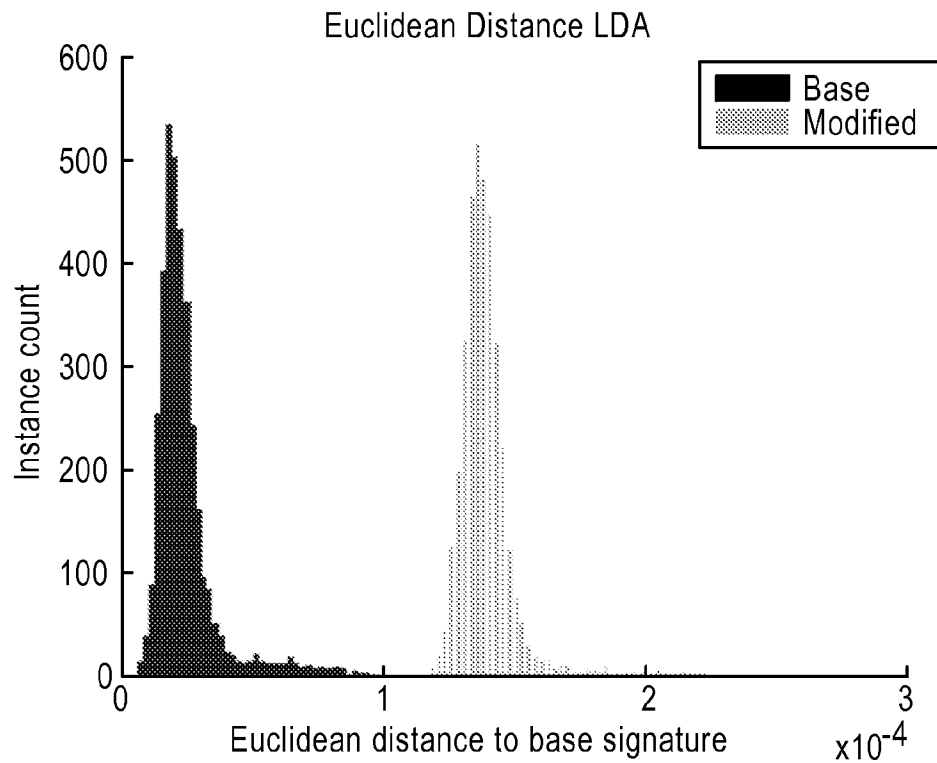


Figure 25




Figure 26

2700




Figure 27

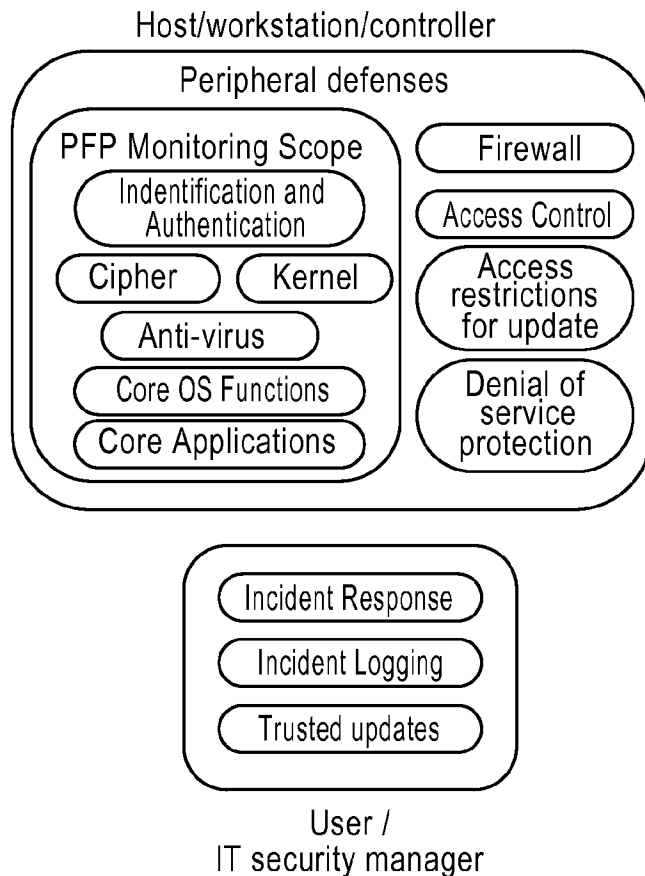


Figure 28




Figure 29

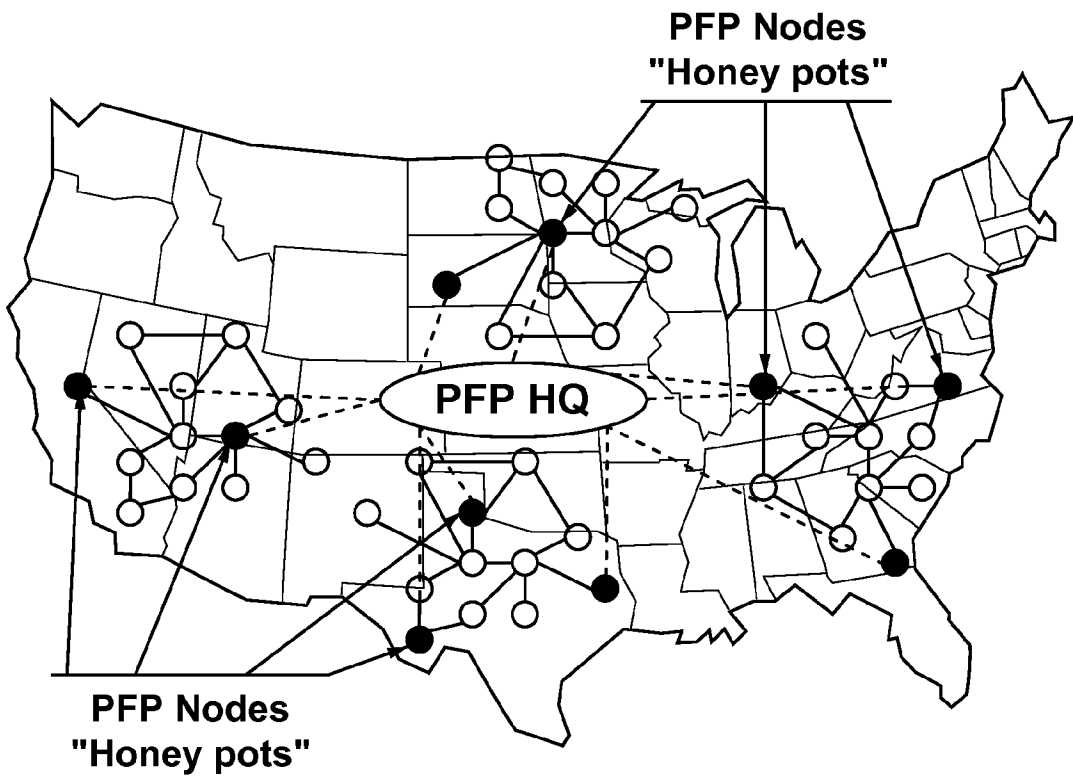


Figure 30




Figure 31

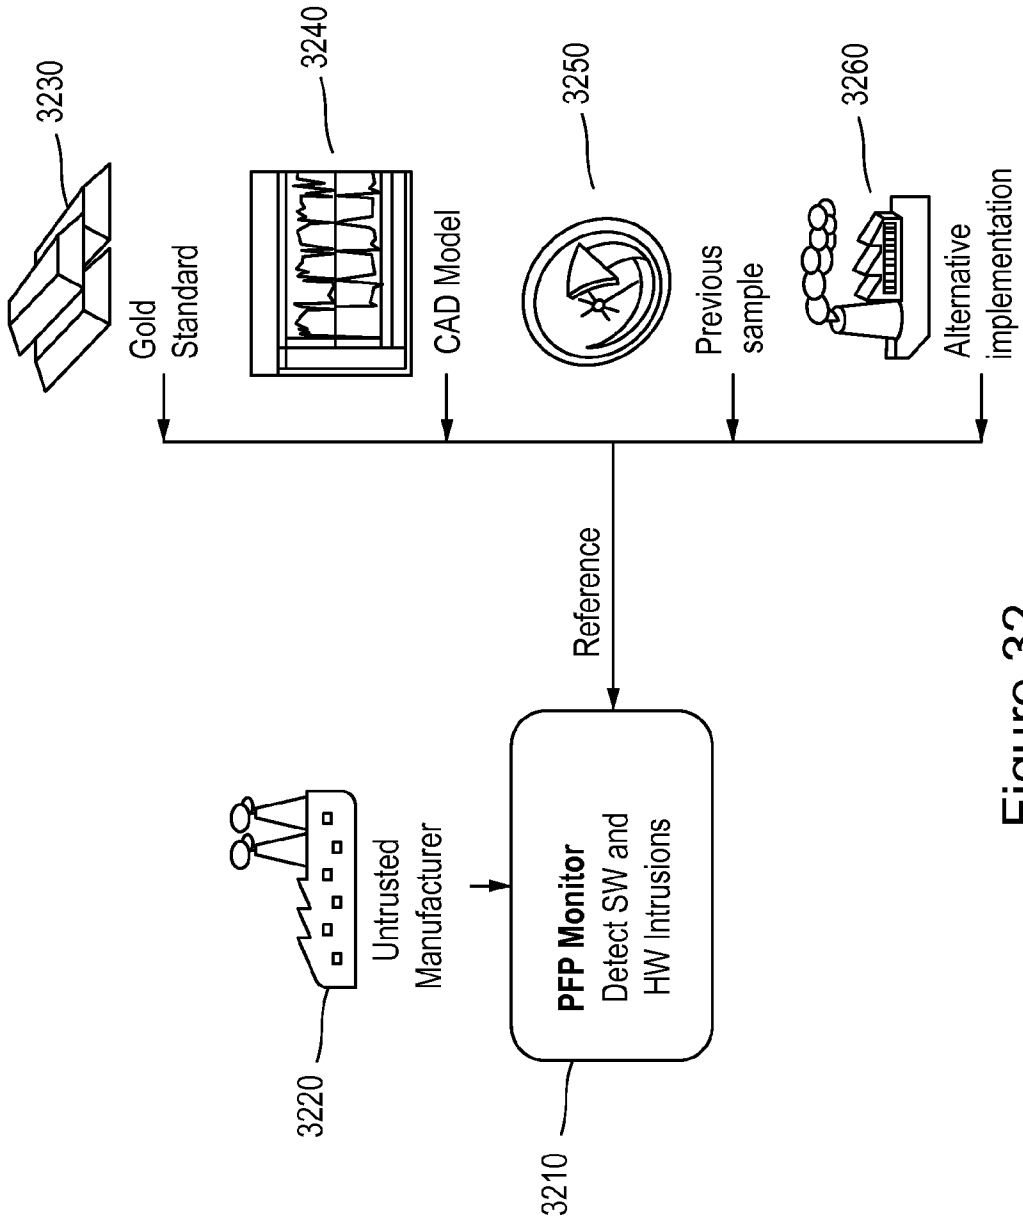


Figure 32




Figure 33




Figure 34

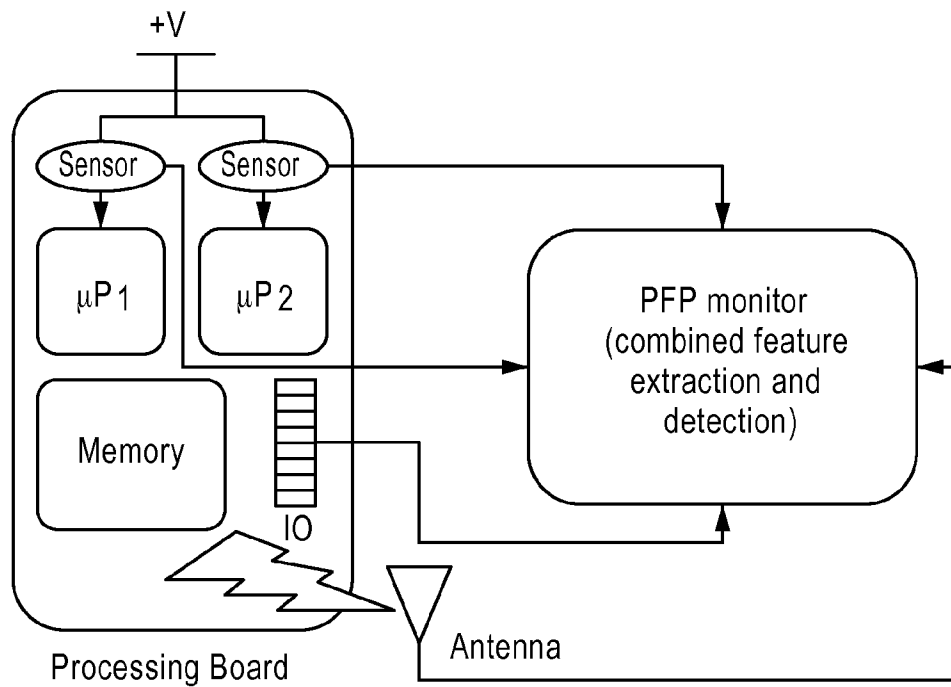


Figure 35




Figure 37




Figure 36




Figure 38




Figure 39

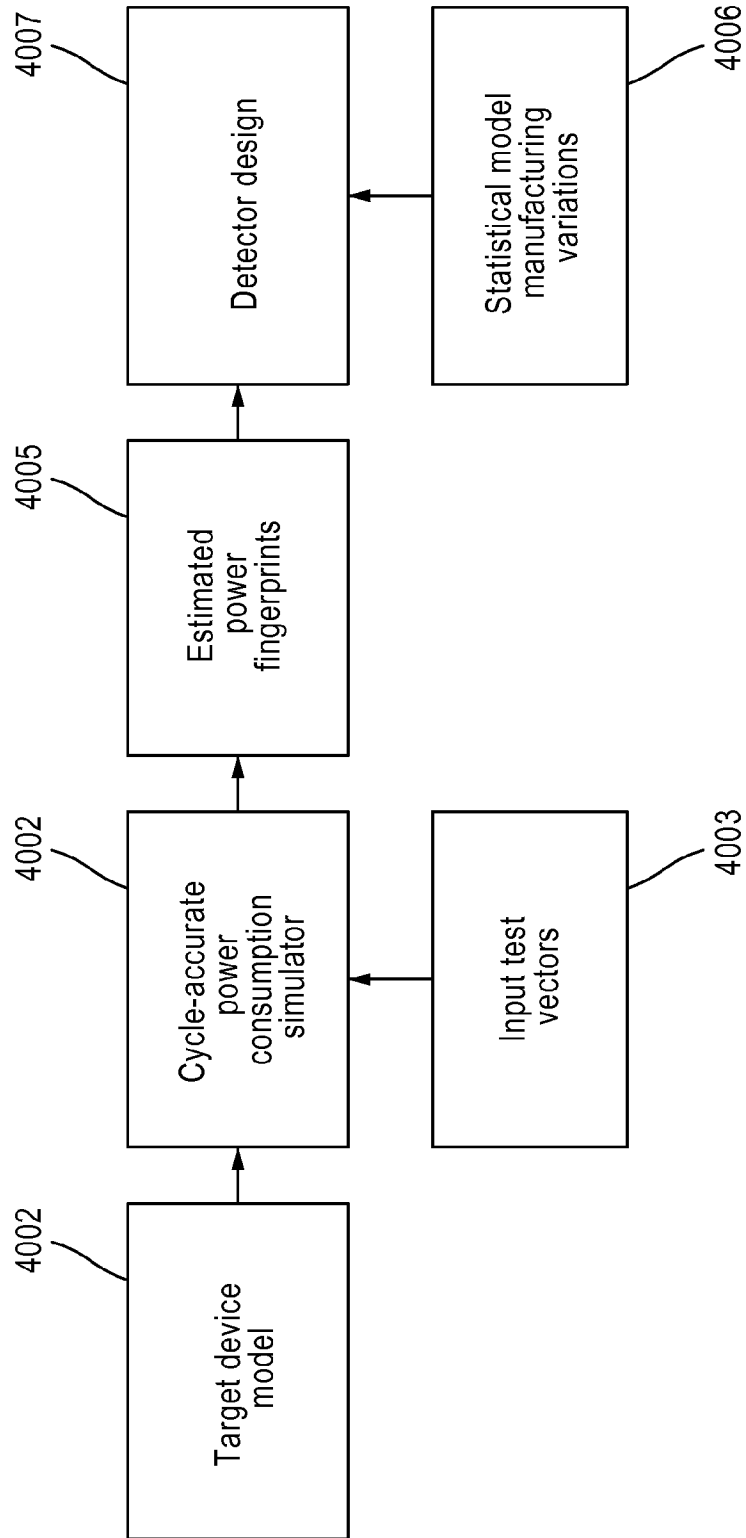


Figure 40




Figure 41A




Figure 41B




Figure 42A




Figure 42B




Figure 43A




Figure 43B




Figure 43C




Figure 44A

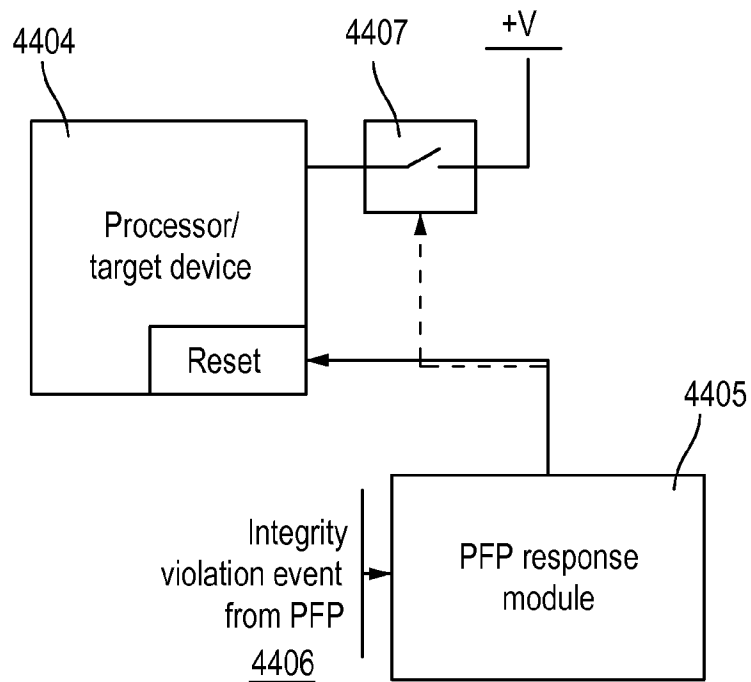


Figure 44B




Figure 45A




Figure 45B




Figure 46A
PFP Solution for spot-testing SCADA deployments




Figure 46C
Around-the-clock monitoring using the embedded PFP monitor

Figure 46B
For distributed SCADA systems, round-robin monitoring. Each device can be assessed several times a day




Figure 46D




Figure 47