

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
8 May 2008 (08.05.2008)

PCT

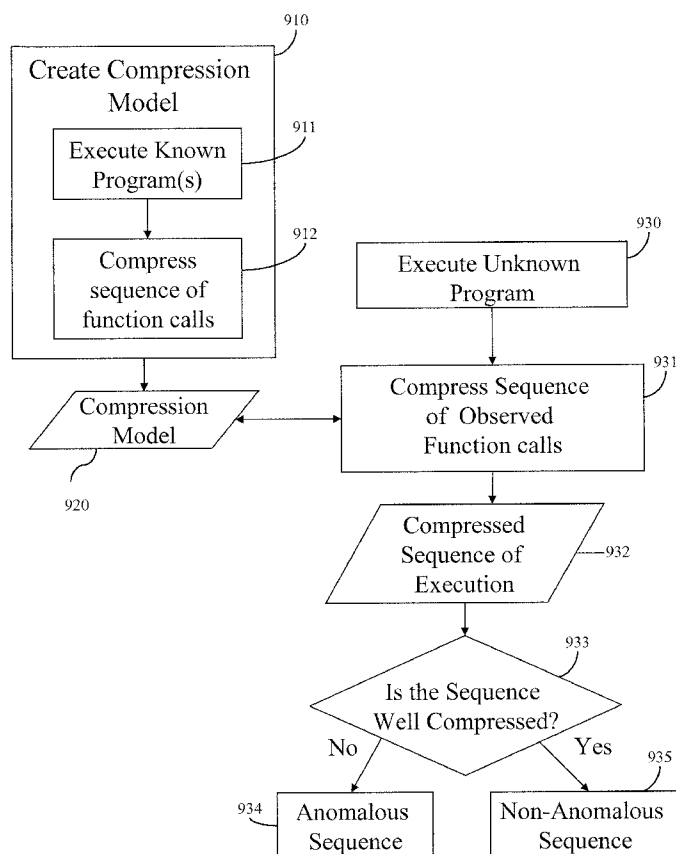
(10) International Publication Number  
WO 2008/055156 A2

- (51) International Patent Classification:  
G06F 11/08 (2006.01)
- (21) International Application Number:  
PCT/US2007/083003
- (22) International Filing Date: 30 October 2007 (30.10.2007)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
60/855,704 30 October 2006 (30.10.2006) US  
60/856,669 3 November 2006 (03.11.2006) US
- (71) Applicant (for all designated States except US): THE TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK [US/US]; 412 Low Memorial Library, 535 West 116th Street, New York, NY 10027 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): KEROMYTIS, Angelos, D. [GR/US]; 423 W. 120th Street, #98, New York,

- NY 10027 (US). STOLFO, Salvatore, J. [US/US]; 80 Kenilworth Road, Ridgewood, NJ 07450 (US).
- (74) Agents: HALAS, Michael, O. et al.; Wilmer Cutler Pickering Hale And Dorr LLP, 399 Park Avenue, New York, NY 10022 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

[Continued on next page]

(54) Title: METHODS, MEDIA, AND SYSTEMS FOR DETECTING AN ANOMALOUS SEQUENCE OF FUNCTION CALLS



(57) Abstract: Methods, media, and systems for detecting an anomalous sequence of function calls are provided. The methods can include compressing a sequence of function calls made by the execution of a program using a compression model; and determining the presence of an anomalous sequence of function calls in the sequence of function calls based on the extent to which the sequence of function calls is compressed. The methods can further include executing at least one known program; observing at least one sequence of function calls made by the execution of the at least one known program; assigning each type of function call in the at least one sequence of function calls made by the at least one known program a unique identifier; and creating at least part of the compression model by recording at least one sequence of unique identifiers.

WO 2008/055156 A2



FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL,  
PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM,  
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

— *without international search report and to be republished  
upon receipt of that report*

## METHODS, MEDIA, AND SYSTEMS FOR DETECTING AN ANOMALOUS SEQUENCE OF FUNCTION CALLS

### Cross Reference to Related Applications

[0001] This application claims the benefit of United States Provisional Patent Application No. 60/855,704 filed October 30, 2006 and United State Provisional Patent Application No. 60/856,669 filed November 3, 2006, each of which is hereby incorporated by reference herein in its entirety.

### Technology Area

[0002] The disclosed subject matter relates to methods, media, and systems for detecting an anomalous sequence of function calls.

### Background

[0003] Applications may terminate due to any number of threats, program errors, software faults, attacks, or any other suitable software failure. Computer viruses, worms, trojans, hackers, key recovery attacks, malicious executables, probes, etc. are a constant menace to users of computers connected to public computer networks (such as the Internet) and/or private networks (such as corporate computer networks). In response to these threats, many computers are protected by antivirus software and firewalls. However, these preventative measures are not always adequate. For example, many services must maintain a high availability when faced by remote attacks, high-volume events (such as fast-spreading worms like Slammer and Blaster), or simple application-level denial of service (DoS) attacks.

[0004] Aside from these threats, applications generally contain errors during operation, which typically result from programmer error. Regardless of whether an application is attacked by one of the above-mentioned threats or contains errors during operation, these software faults and failures result in illegal memory access errors, division by zero errors, buffer overflows attacks, etc. These errors cause an application to terminate its execution or “crash.”

### Summary

[0005] Methods, media, and systems for detecting an anomalous sequence of function calls are provided. In some embodiments, methods for detecting an anomalous sequence of

function calls is provided. The methods including compressing a sequence of function calls made by the execution of a program using a compression model; and determining the presence of an anomalous sequence of function calls in the sequence of function calls based on the extent to which the sequence of function calls is compressed. In some embodiments, the methods further include executing at least one known program; observing at least one sequence of function calls made by the execution of the at least one known program; assigning each type of function call in the at least one sequence of function calls made by the at least one known program a unique identifier; and creating at least part of the compression model by recording at least one sequence of unique identifiers based on the unique identifiers assigned to each type of function call and the observed at least one sequence of function calls.

**[0006]** In some embodiments, computer-readable media containing computer-executable instructions that, when executed by a processor, cause the processor to perform a method for detecting an anomalous sequence of function calls are provided. The methods including compressing a sequence of function calls made by the execution of a program using a compression model; and determining the presence of an anomalous sequence of function calls in the sequence of function calls based on the extent to which the sequence of function calls is compressed. In some embodiments, the methods further include executing at least one known program; observing at least one sequence of function calls made by the execution of the at least one known program; assigning each type of function call in the at least one sequence of function calls made by the at least one known program a unique identifier; and creating at least part of the compression model by recording at least one sequence of unique identifiers based on the unique identifiers assigned to each type of function call and the observed at least one sequence of function calls.

**[0007]** In some embodiments, systems for detecting an anomalous sequence of function calls including a memory; and a processor in communication with the memory are provided. Wherein the processor compresses a sequence of function calls made by the execution of a program using a compression model; and determines the presence of an anomalous sequence of function calls in the sequence of function calls based on the extent to which the sequence of function calls is compressed. In the some embodiments, the processor further executes at least one known program; observes at least one sequence of function calls made by the execution of the at least one known program; assigns each type of function call in the at least one sequence of function calls made by the at least one known program a unique identifier; and creates at least part of the compression model by recording at least one sequence of

unique identifiers based on the unique identifiers assigned to each type of function call and the observed at least one sequence of function calls.

### **Brief Description of the Drawings**

- [0008] FIG. 1 is a schematic diagram of an illustrative system suitable for implementation of an application that monitors other applications and protects these applications against faults in accordance with some embodiments;
- [0009] FIG. 2 is a detailed example of the server and one of the workstations of FIG.1 that may be used in accordance with some embodiments;
- [0010] FIG. 3 shows a simplified diagram illustrating repairing faults in an application and updating the application in accordance with some embodiments;
- [0011] FIG. 4 shows a simplified diagram illustrating detecting and repairing an application in response to a fault occurring in accordance with some embodiments;
- [0012] FIG. 5 shows an illustrative example of emulated code integrated into the code of an existing application in accordance with some embodiments;
- [0013] FIG. 6 shows a simplified diagram illustrating detecting and repairing an application using an application community in accordance with some embodiments of the disclosed subject matter;
- [0014] FIG. 7 shows an illustrative example of a table that may be calculated by a member of the application community for distributed bidding in accordance with some embodiments of the disclosed subject matter; and
- [0015] FIG. 8 shows a simplified diagram illustrating shows identifying a function call as being anomalous in accordance with some embodiments.
- [0016] FIG. 9 shows a simplified illustration of methods for generating a compression model and detecting an anomalous sequence of function calls in accordance with some embodiments.

### **Detailed Description**

[0017] Methods, media, and systems for detecting an anomalous sequence of function calls and/or detecting anomalous program executions are provided. In some embodiments, systems and methods are provided that model application level computations and running programs, and that detect anomalous executions by, for example, instrumenting, monitoring and analyzing application-level program function calls and/or arguments. Such an approach

can be used to detect anomalous program executions that may be indicative of a malicious attack or program fault.

**[0018]** The anomaly detection algorithm being used may be, for example, a probabilistic anomaly detection (PAD) algorithm or a one class support vector machine (OCSVM), which are described below, or any other suitable algorithm.

**[0019]** Anomaly detection may be applied to process execution anomaly detection, file system access anomaly detection, and/or network packet header anomaly detection.

Moreover, as described herein, according to various embodiments, an anomaly detector may be applied to program execution state information. For example, as explained in greater detail below, an anomaly detector may model information on the program stack to detect anomalous program behavior.

**[0020]** In various embodiments, using PAD to model program stack information, such stack information may be extracted using, for example, Selective Transactional EMulation (STEM), which is described below and which permits the selective execution of certain parts, or all, of a program inside an instruction-level emulator, using the Valgrind emulator, by modifying a program's binary or source code to include indicators of what functions calls are being made (and any other suitable related information), or using any other suitable technique. In this manner, it is possible to determine dynamically (and transparently to the monitored program) the necessary information such as stack frames, function-call arguments, etc. For example, one or more of the following may be extracted from the program stack specific information: function name, the argument buffer name it may reference, and other features associated with the data sent to or returned from the called function (e.g., the length in bytes of the data, or the memory location of the data).

**[0021]** For example, as illustrated in FIG. 8, an anomaly detector may be applied, for example, by extracting data pushed onto the stack (e.g., by using an emulator or by modifying a program), and creating a data record provided to the anomaly detector for processing at 802. According to various embodiments, in a first phase, an anomaly detector models normal program execution stack behavior. In the detection mode, after a model has been computed, the anomaly detector can detect stacked function references as anomalous at 806 by comparing those references to the model based on the training data at 804.

**[0022]** Once an anomaly is detected, according to various embodiments, selective transactional emulation (STEM) and error virtualization may be used to reverse (undo) the effects of processing the malicious input (e.g., changes to program variables or the file system) in order to allow the program to recover execution in a graceful manner. In this

manner, the precise location of the failed (or attacked) program at which an anomaly was found may be identified. Also, the application of an anomaly detector to function calls can enable rapid detection of malicious program executions, such that it is possible to mitigate against such faults or attacks (e.g., by using patch generation systems, or content filtering signature generation systems). Moreover, given precise identification of a vulnerable location, the performance impact may be reduced by using STEM for parts or all of a program's execution.

[0023] As explained above, anomaly detection can involve the use of detection models. These models can be used in connection with automatic and unsupervised learning.

[0024] In some embodiments, such models can be created from a training set of function calls and optionally at least some of their arguments. For example, a model can include a compressed series of function calls observed by execution of a training set of known non-anomalous applications. Various compression techniques can be used in various embodiments. For example, in some embodiments, a compression model can be used as a dictionary in a compression technique, such as Lempel-Ziv-Welch (LZW) compression. More particularly, for example, if each function call is assigned an identifier (e.g., a two digit number, a string, a code, etc.), a sequence of function calls can appear as a series of identifiers. Different series of identifiers in a training set can then be used to form a library or table of sequences which can be part of a compression model. A determination can then be made as to whether a sequence of function calls is anomalous based upon how well they can be compressed using a compression model (e.g., a dictionary, a library, a table of sequences) derived from the training set. If a stream of function calls can be compressed well using a model (e.g., created from known non-anomalous applications), then training and test sets have similar function calls and the function calls can be considered non-anomalous. However, if the test set cannot be compressed well, then the test set may contain anomalous function calls. The determination of the extent to which test data needs to be compressed in order to be found to be non-anomalous can be performed using various techniques, such as, for example, techniques based on empirical data and/or user and/or administrator settings.

[0025] FIG. 9 illustrates a method for the creation of a compression model at 910. This can include executing, at 911, one or more known non-anomalous programs and compressing, at 912, sequence of observed function calls, to generate a compression model 920. A test program can be executed, at 930, and a sequence of function calls resulting from the program's execution can be compressed, at 931, using compression model 920, to generate a compressed sequence of function calls 932. It can be determined, at 933, whether

the compressed sequence 932 is well compressed. This can be based on, for example, the percent of function calls that were not compressed, the length of various sequences of uncompressed function calls, the distance between uncompressed function calls, the density and/or distribution of uncompressed function calls, and/or the number of unique uncompressed sequences and/or unique uncompressed function calls. If the compressed sequence 932 is well compressed, it can be considered non-anomalous, at 935. If it is not well compressed, it can be considered anomalous, at 934. If a program execution is considered anomalous, various recovery actions can be performed, such as those described herein. In some embodiments, programs executed at 911 can be known anomalous programs and determinations of a sequence 932 being well compressed can indicate anomalous sequences. Some embodiments can create various models 920, including models of both anomalous and non-anomalous programs. In such embodiments, a compressed sequence 931 can be compared to various models.

**[0026]** In some embodiments, execution, at 930, can include execution of all or part of a program and can be performed, for example, natively or in an emulator. If an anomalous sequence is detected, at 933, one or more of the function calls and/or the order of the function calls can be altered to make the sequence of function calls non-anomalous, execution 930 can be rolled back, and/or execution 930 can be terminated. In various embodiments, a compressed sequence 932 may not be formed and/or stored in memory. Instead compression 931 can feed directly into the determination, at 933, of how well (or not well) a test program is compressed. The determination, at 933, can take place, for example, continuously during execution, at 930, at intervals and/or breaks in execution, at 930, or upon completion of execution, at 930. In various embodiments, differentiation of function calls (e.g. for assigning them a compression value) can be based on, for example, the value of an argument or arguments, the number of arguments, and/or the identity of the function (e.g., based on a memory address). In various embodiments, creation, at 910, execution, at 930, and/or determination, at 933, can take place on the same digital processing device or on different digital processing devices. Compression model 920 can include various compression models formed by executing various training data on various digital processing devices.

**[0027]** A probabilistic anomaly detection (PAD) algorithm can also be used to train a model for detecting anomalies. In various embodiments such a model can be used in place of, in addition to, or in combination with, for example, a compression model (e.g., model 926). This model may be, in essence, a density estimation, where the estimation of a density function  $p(x)$  over normal data allows the definition of anomalies as data elements that occur

with low probability. The detection of low probability data (or events) are represented as consistency checks over the normal data, where a record is labeled anomalous if it fails any one of these tests.

**[0028]** First and second order consistency checks can be applied. First order consistency checks verify that a value is consistent with observed values of that feature in the normal data set. These first order checks compute the likelihood of an observation of a given feature,  $P(X_i)$ , where  $X_i$  are the feature variables. Second order consistency checks determine the conditional probability of a feature value given another feature value, denoted by  $P(X_i|X_j)$ , where  $X_i$  and  $X_j$  are the feature variables.

**[0029]** One way to compute these probabilities is to estimate a multinomial that computes the ratio of the counts of a given element to the total counts. However, this results in a biased estimator when there is a sparse data set. Another approach is to use an estimator to determine these probability distributions. For example, let  $N$  be the total number of observations,  $N_i$  be the number of observations of symbol  $i$ ,  $\alpha$  be the “pseudo count” that is added to the count of each observed symbol,  $k^0$  be the number of observed symbols, and  $L$  be the total number of possible symbols. Using these definitions, the probability for an observed element  $i$  can be given by:

$$P(X = i) = \frac{N_i + \alpha}{k^0 \alpha + N} C \tag{1}$$

and the probability for an unobserved element  $i$  can be:

$$P(X = i) = \frac{1}{L - k^0} (1 - C) \tag{2}$$

where  $C$ , the scaling factor, accounts for the likelihood of observing a previously observed element versus an unobserved element.  $C$  can be computed as:

$$C = (\sum_{k=k^0}^L \frac{k^0 \alpha + N}{k \alpha + N} m_k) (\sum_{k \geq k^0} m_k)^{-1} \tag{3}$$

where

$$m_k = P(S = k) \frac{k!}{k^0!} \frac{\Gamma(k \alpha)}{\Gamma(k \alpha + N)}$$

and  $P(s=k)$  is a prior probability associated with the size of the subset of elements in the alphabet that have non-zero probability.

Because this computation of  $C$  can be time consuming,  $C$  can also be calculated by:

$$C = \frac{N}{N + L - k^0} \quad (4)$$

The consistency check can be normalized to account for the number of possible outcomes of  $L$  by  $\log(P/(1/L)) = \log(P) + \log(L)$ .

**[0030]** Another approach that may be used instead of using PAD for model generation and anomaly detection is a one class SVM (OCSVM) algorithm. The OCSVM algorithm can be used to map input data into a high dimensional feature space (via a kernel) and iteratively find the maximal margin hyperplane which best separates the training data from the origin. The OCSVM may be viewed as a regular two-class SVM where all the training data lies in the first class, and the origin is taken as the only member of the second class. Thus, the hyperplane (or linear decision boundary) can correspond to the classification rule:

$$f(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle + b \quad (5)$$

**[0031]** where  $\mathbf{w}$  is the normal vector and  $b$  is a bias term. The OCSVM can be used to solve an optimization problem to find the rule  $f$  with maximal geometric margin. This classification rule can be used to assign a label to a test example  $\mathbf{x}$ . If  $f(\mathbf{x}) < 0$ ,  $\mathbf{x}$  can be labeled as an anomaly, otherwise it can be labeled as normal. In practice, there is a trade-off between maximizing the distance of the hyperplane from the origin and the number of training data points contained in the region separated from the origin by the hyperplane.

**[0032]** Solving the OCSVM optimization problem can be equivalent to solving the dual quadratic programming problem:

$$\min_{\alpha} \frac{1}{2} \sum_{ij} \alpha_i \alpha_j K(x_i, x_j) \quad (6)$$

subject to the constraints

$$0 \leq \alpha_i \leq \frac{1}{\nu l} \quad (7)$$

and

$$\sum_i \alpha_i = 1 \tag{8}$$

where  $\alpha_i$  is a lagrange multiplier (or “weight” on example  $i$  such that vectors associated with non-zero weights are called “support vectors” and solely determine the optimal hyperplane),  $\nu$  is a parameter that controls the trade-off between maximizing the distance of the hyperplane from the origin and the number of data points contained by the hyperplane,  $l$  is the number of points in the training dataset, and  $K(x_i, x_j)$  is the kernel function. By using the kernel function to project input vectors into a feature space, nonlinear decision boundaries can be allowed for. Given a feature map:

$$\phi : X \rightarrow \mathbb{R}^N \tag{9}$$

where  $\Phi$  maps training vectors from input space  $X$  to a high-dimensional feature space, the kernel function can be defined as:

$$K(x, y) = \langle \phi(x), \phi(y) \rangle \tag{10}$$

**[0033]** Feature vectors need not be computed explicitly, and computational efficiency can be improved by directly computing kernel values  $K(x, y)$ . Three common kernels can be used:

- Linear kernel:  $K(x, y) = (x \cdot y)$
- Polynomial kernel:  $K(x, y) = (x \cdot y + 1)^d$ , where  $d$  is the degree of the polynomial
- Gaussian kernel:  $K(x, y) = e^{-\|x-y\|^2/(2\sigma^2)}$ , where  $\sigma^2$  is the variance

**[0034]** Kernels from binary feature vectors can be obtained by mapping a record into a feature space such that there is one dimension for every unique entry for each record value. A particular record can have the value 1 in the dimensions which correspond to each of its specific record entries, and the value 0 for every other dimension in feature space. Linear kernels, second order polynomial kernels, and gaussian kernels can be calculated using these feature vectors for each record. Kernels can also be calculated from frequency-based feature vectors such that, for any given record, each feature corresponds to the number of occurrences of the corresponding record component in the training set. For example, if the second component of a record occurs three times in the training set, the second feature value for that record is three. These frequency-based feature vectors can be used to compute linear and polynomial kernels.

**[0035]** According to various embodiments, “mimicry attacks” which might otherwise thwart OS system call level anomaly detectors by using normal appearing sequences of system calls can be detected. For example, mimicry attacks are less likely to be detected when the system calls are only modeled as tokens from an alphabet, without any information about arguments. Therefore, according to various embodiments, the models used are enriched with information about the arguments (data) such that it may be easier to detect mimicry attacks.

**[0036]** According to various embodiments, models are shared among many members of a community running the same application (referred to as an “application community”). In particular, some embodiments can share models with each other and/or update each other’s models such that the learning of anomaly detection models is relatively quick. For example, instead of running a particular application for days at a single site, according to various embodiments, thousands of replicated applications can be run for a short period of time (e.g., one hour), and the models created based on the distributed data can be shared. While only a portion of each application instance may be monitored, for example, the entire software body can be monitored across the entire community. This can enable the rapid acquisition of statistics, and relatively fast learning of an application profile by sharing, for example, aggregate information (rather than the actual raw data used to construct the model).

**[0037]** Model sharing can result in one standard model that an attacker could potentially access and use to craft a mimicry attack. Therefore, according to various embodiments, unique and diversified models can be created. For example, such unique and diversified models can be created by randomly choosing particular features from the application execution that is modeled, such that the various application instances compute distinct models. In this manner, attacks may need to avoid detection by multiple models, rather than just a single model. Creating unique and diversified models not only has the advantage of being more resistant to mimicry attacks, but also may be more efficient. For example, if only a portion of an application is modeled by each member of an application community, monitoring will generally be simpler (and cheaper) for each member of the community. In the event that one or more members of an application community are attacked, according to various embodiments, the attack (or fault) will be detected, and patches or a signature can be provided to those community members who are blind to the crafted attack (or fault).

**[0038]** Random (distinct) model building and random probing may be controlled by a software registration key provided by a commercial off-the-shelf (COTS) software vendor or some other data providing “randomization.” For example, for each member of an application

community, some particular randomly chosen function or functions and its associated data may be chosen for modeling, while others may simply be ignored. Moreover, because vendors can generate distinct keys and serial numbers when distributing their software, this feature can be used to create a distinct random subset of functions to be modeled. Also, according to various embodiments, even community members who model the same function or functions may exchange models.

**[0039]** According to various embodiments, when an application execution is being analyzed over many copies distributed among a number of application community members to profile the entire code of an application, it can be determined whether there are any segments of code that are either rarely or never executed, and a map can be provided of the code layout identifying “suspect code segments” for deeper analysis and perhaps deeper monitoring. Those segments identified as rarely or never executed may harbor vulnerabilities not yet executed or exploited. Such segments of code may have been designed to execute only for very special purposes such as error handling, or perhaps even for triggering malicious code embedded in the application. Since they are rarely or never executed, one may presume that such code segments have had less regression testing, and may have a higher likelihood of harboring faulty code.

**[0040]** Rarely or never executed code segments may be identified and may be monitored more thoroughly through, for example, emulation. This deep monitoring may have no discernible overhead since the code in question is rarely or never executed. But such monitoring performed in each community member may prevent future disasters by preventing such code (and its likely vulnerabilities) from being executed in a malicious/faulty manner. Identifying such code may be performed by a sensor that monitors loaded modules into the running application (e.g., DLL loads) as well as addresses (PC values) during code execution and creates a “frequency” map of ranges of the application code. For example, a set of such distributed sensors may communicate with each other (or through some site that correlates their collective information) to create a central, global MAP of the application execution profile. This profile may then be used to identify suspect code segments, and then subsequently, this information may be useful to assign different kinds of sensors/monitors to different code segments. For example, an interrupt service routine (ISR) may be applied to these suspect sections of code.

**[0041]** It is noted that a single application instance may have to be run many times (e.g., thousands of times) in order to compute an application profile or model. However, distributed sensors whose data is correlated among many (e.g., a thousand) application

community members can be used to compute a substantially accurate code profile in a relatively short amount of time. This time may be viewed as a “training period” to create the code map.

**[0042]** According to various embodiments, models may be automatically updated as time progresses. For example, although a single site may learn a particular model over some period of time, application behavior may change over time. In this case, the previously learned model may no longer accurately reflect the application characteristics, resulting in, for example, the generation of an excessive amount of false alarms (and thus an increase in the false positive rate over time). A possible solution to this “concept drift” issue entails at least two possible approaches, both intended to update models over time. A first approach to solving (or at least reducing the effects of) the “concept drift” issue involves the use of “incremental learning algorithms,” which are algorithms that piecemeal update their models with new data, and that may also “expire” parts of the computed model created by older data. This piecemeal incremental approach is intended to result in continuous updating using relatively small amounts of data seen by the learning system.

**[0043]** A second approach to solving (or at least reducing the effect of) the “concept drift” issue involves combining multiple models. For example, presuming that an older model has been computed from older data during some “training epoch,” a new model may be computed concurrently with a new epoch in which the old model is used to detect anomalous behavior. Once a new model is computed, the old model may be retired or expunged, and replaced by the new model. Alternatively, for example, multiple models such as described above may be combined. In this case, according to various embodiments, rather than expunging the old model, a newly created model can be algorithmically combined with the older model using any of a variety of suitable means. In the case of statistical models that are based upon frequency counts of individual data points, for example, an update may consist of an additive update of the frequency count table. For example, PAD may model data by computing the number of occurrences of a particular data item, “X.” Two independently learned PAD models can thus have two different counts for the same value, and a new frequency table can be readily computed by summing the two counts, essentially merging two tables and updating common values with the sum of their respective counts.

**[0044]** According to various embodiments, the concept of model updating that is readily achieved in the case of computed PAD models may be used in connection with model sharing. For example, rather than computing two models by the same device for a distinct application, two distinct models may be computed by two distinct instances of an application

by two distinct devices, as described above. The sharing of models may thus be implemented by the model update process described herein. Hence, a device may continuously learn and update its models either by computing its own new model, or by downloading a model from another application community member (e.g., using the same means involved in the combining of models).

**[0045]** In the manners described above, an application community may be configured to continuously refresh and update all community members, thereby making mimicry attacks far more difficult to achieve.

**[0046]** As mentioned above, it is possible to mitigate against faults or attacks by using patch generation systems. In accordance with various embodiments, when patches are generated, validated, and deployed, the patches and/or the set of all such patches may serve the following.

**[0047]** First, according to various embodiments, each patch may be used as a “pattern” to be used in searching other code for other unknown vulnerabilities. An error (or design flaw) in programming that is made by a programmer and that creates a vulnerability may show up elsewhere in code. Therefore, once a vulnerability is detected, the system may use the detected vulnerability (and patch) to learn about other (e.g., similar) vulnerabilities, which may be patched in advance of those vulnerabilities being exploited. In this manner, over time, a system may automatically reduce (or eliminate) vulnerabilities.

**[0048]** Second, according to various embodiments, previously generated patches may serve as exemplars for generating new patches. For example, over time, a taxonomy of patches may be assembled that are related along various syntactic and semantic dimensions. In this case, the generation of new patches may be aided by prior examples of patch generation.

**[0049]** Additionally, according to various embodiments, generated patches may themselves have direct economic value. For example, once generated, patches may be “sold” back to the vendors of the software that has been patched.

**[0050]** As mentioned above, in order to alleviate monitoring costs, instead of running a particular application for days at a single site, many (e.g., thousands) replicated versions of the application may be run for a shorter period of time (e.g., an hour) to obtain the necessary models. In this case, only a portion of each replicated version of the application may be monitored, although the entire software body is monitored using the community of monitored software applications. Moreover, according to various embodiments, if a software module has been detected as faulty, and a patch has been generated to repair it, that portion of the

software module, or the entire software module, may no longer need to be monitored. In this case, over time, patch generated systems may have fewer audit/monitoring points, and may thus improve in execution speed and performance. Therefore, according to various embodiments, software systems may be improved, where vulnerabilities are removed, and the need for monitoring is reduced (thereby reducing the costs and overheads involved with detecting faults).

**[0051]** It is noted that, although described immediately above with regard to an application community, the notion of automatically identifying faults of an application, improving the application over time by repairing the faults, and eliminating monitoring costs as repairs are deployed may also be applied to a single, standalone instance of an application (without requiring placements as part of a set of monitored application instances).

**[0052]** Selective transactional emulation (STEM) and error virtualization can be beneficial for reacting to detected failures, attacks, and or anomalous behavior in software. According to various embodiments, STEM and error virtualization can be used to provide enhanced detection of some types of attacks, and enhanced reaction mechanisms to some types of attacks/failures.

**[0053]** A learning technique can be applied over multiple executions of a piece of code (e.g., a function or collection of functions) that may previously have been associated with a failure, or that is being proactively monitored. By retaining knowledge on program behavior across multiple executions, certain invariants (or probable invariants) may be learned, whose violation in future executions indicates an attack or imminent software fault.

**[0054]** In the case of control hijacking attacks, certain control data that resides in memory is overwritten through some mechanism by an attacker. That control data is then used by the program for an internal operation, allowing the attacker to subvert the program. Various forms of buffer overflow attacks (stack and heap smashing, jump into libc, etc.) operate in this fashion. Such attacks can be detected when the corrupted control data is about to be used by the program (i.e., after the attack has succeeded). In various embodiments, such control data (e.g., memory locations or registers that hold such data) that is about to be overwritten with “tainted” data, or data provided by the network (which is potentially malicious) can be detected.

**[0055]** In accordance with various embodiments, how data modifications propagate throughout program execution can be monitored by maintaining a memory bit for every byte or word in memory. This bit is set for a memory location when a machine instruction uses as input data that was provided as input to the program (e.g., was received over the network, and

is thus possibly malicious) and produces output that is stored in this memory location. If a control instruction (such as a JUMP or CALL) uses as an argument a value in a memory location in which the bit is set (i.e., the memory location is “tainted”), the program or the supervisory code that monitors program behavior can recognize an anomaly and raises an exception.

**[0056]** Detecting corruption before it happens, rather than later (when the corrupted data is about to be used by a control instruction), makes it possible to stop an operation and to discard its results/output, without other collateral damage. Furthermore, in addition to simply retaining knowledge of what is control and what is non-control data, according to various embodiments, knowledge of which instructions in the monitored piece of code typically modify specific memory locations can also be retained. Therefore, it is possible to detect attacks that compromise data that are used by the program computation itself, and not just for the program control flow management.

**[0057]** According to various embodiments, the inputs to the instruction(s) that can fail (or that can be exploited in an attack) and the outputs (results) of such instructions can be correlated with the inputs to the program at large. Inputs to an instruction are registers or locations in memory that contain values that may have been derived (in full or partially) by the input to the program. By computing a probability distribution model on the program input, alternate inputs may be chosen to give to the instruction or the function (“input rewriting” or “input modification”) when an imminent failure is detected, thereby allowing the program to “sidestep” the failure. However, because doing so may still cause the program to fail, according to various embodiments, micro-speculation (e.g., as implemented by STEM) can optionally be used to verify the effect of taking this course of action. A recovery technique (with different input values or error virtualization, for example) can then be used. Alternatively, for example, the output of the instruction may be caused to be a value/result that is typically seen when executing the program (“output overloading”).

**[0058]** In both cases (input modification or output overloading), the values to use may be selected based on several different criteria, including but not limited to one or more of the following: the similarity of the program input that caused failure to other inputs that have not caused a failure; the most frequently seen input or output value for that instruction, based on contextual information (e.g., when particular sequence of functions are in the program call stack); and most frequently seen input or output value for that instruction across all executions of the instruction (in all contexts seen). For example, if a particular DIVIDE instruction is detected in a function that uses a denominator value of zero, which would cause

a process exception, and subsequently program failure, the DIVIDE instruction can be executed with a different denominator (e.g., based on how similar the program input is to other program inputs seen in the past, and the denominator values that these executions used). Alternatively, the DIVIDE instruction may be treated as though it had given a particular division result. The program may then be allowed to continue executing, while its behavior is being monitored. Should a failure subsequently occur while still under monitoring, a different input or output value for the instruction can be used, for example, or a different repair technique can be used. According to various embodiments, if none of the above strategies is successful, the user or administrator may be notified, program execution may be terminated, a rollback to a known good state (ignoring the current program execution) may take place, and/or some other corrective action may be taken.

**[0059]** According to various embodiments, the techniques used to learn typical data can be implemented as designer choice. For example, if it is assumed that the data modeled is 32-bit words, a probability distribution of this range of values can be estimated by sampling from multiple executions of the program. Alternatively, various cluster-based analyses may partition the space of typical data into clusters that represent groups of similar/related data by some criteria. Vector Quantization techniques representing common and similar data based on some “similarity” measure or criteria may also be compiled and used to guide modeling.

**[0060]** FIG. 1 is a schematic diagram of an illustrative system 100 suitable for implementation of various embodiments. As illustrated in FIG. 1, system 100 may include one or more workstations 102. Workstations 102 can be local to each other or remote from each other, and can be connected by one or more communications links 104 to a communications network 106 that is linked via a communications link 108 to a server 110.

**[0061]** In system 100, server 110 may be any suitable server for executing the application, such as a processor, a computer, a data processing device, or a combination of such devices. Communications network 106 may be any suitable computer network including the Internet, an intranet, a wide-area network (WAN), a local-area network (LAN), a wireless network, a digital subscriber line (DSL) network, a frame relay network, an asynchronous transfer mode (ATM) network, a virtual private network (VPN), or any combination of any of the same. Communications links 104 and 108 may be any communications links suitable for communicating data between workstations 102 and server 110, such as network links, dial-up links, wireless links, hard-wired links, etc. Workstations 102 may be personal computers, laptop computers, mainframe computers, data displays, Internet browsers, personal digital assistants (PDAs), two-way pagers, wireless terminals,

portable telephones, etc., or any combination of the same. Workstations 102 and server 110 may be located at any suitable location. In one embodiment, workstations 102 and server 110 may be located within an organization. Alternatively, workstations 102 and server 110 may be distributed between multiple organizations.

**[0062]** The server and one of the workstations, which are depicted in FIG. 1, are illustrated in more detail in FIG. 2. Referring to FIG. 2, workstation 102 may include digital processing device (such as a processor) 202, display 204, input device 206, and memory 208, which may be interconnected. In a preferred embodiment, memory 208 contains a storage device for storing a workstation program for controlling processor 202. Memory 208 may also contain an application for detecting an anomalous sequence of function calls and/or detecting and repairing applications from faults according to various embodiments. In some embodiments, the application may be resident in the memory of workstation 102 or server 110.

**[0063]** Processor 202 may use the workstation program to present on display 204 the application and the data received through communication link 104 and commands and values transmitted by a user of workstation 102. It should also be noted that data received through communication link 104 or any other communications links may be received from any suitable source, such as web services. Input device 206 may be a computer keyboard, a cursor-controller, a dial, a switchbank, lever, or any other suitable input device as would be used by a designer of input systems or process control systems.

**[0064]** Server 110 may include processor 220, display 222, input device 224, and memory 226, which may be interconnected. In some embodiments, memory 226 contains a storage device for storing data received through communication link 108 or through other links, and also receives commands and values transmitted by one or more users. The storage device can further contain a server program for controlling processor 220.

**[0065]** In accordance with some embodiments, a self-healing system that allows an application to automatically recover from software failures and attacks is provided. By selectively emulating at least a portion or all of the application's code when the system detects that a fault has occurred, the system surrounds the detected fault to validate the operands to machine instructions, as appropriate for the type of fault. The system emulates that portion of the application's code with a fix and updates the application. This increases service availability in the presence of general software bugs, software failures, attacks.

**[0066]** Turning to FIGS. 3 and 4, simplified flowcharts illustrating various steps performed in detecting faults in an application and fixing the application in accordance with

some embodiments are provided. These are generalized flow charts. It will be understood that the steps shown in FIGS. 3 and 4 may be performed in any suitable order, some may be deleted, and others added.

**[0067]** Generally, process 300 begins by detecting various types of failures in one or more applications at 310 (in some embodiments, this detection can occur, for example, at 930 and/or 933 of FIG. 9). In some embodiments, detecting for failures may include monitoring the one or more applications for failures, e.g., by using an anomaly detector as described herein. In some embodiments, the monitoring or detecting of failures may be performed using one or more sensors at 310. Failures include programming errors, exceptions, software faults (e.g., illegal memory accesses, division by zero, buffer overflow attacks, time-of-check-to-time-of-use (TOCTTOU) violations, etc.), threats (e.g., computer viruses, worms, trojans, hackers, key recovery attacks, malicious executables, probes, etc.), and any other suitable fault that may cause abnormal application termination or adversely affect the one or more applications.

**[0068]** Any suitable sensors may be used to detect failures or monitor the one or more applications. For example, in some embodiments, anomaly detectors as described herein can be used.

**[0069]** At 320, feedback from the sensors may be used to predict which parts of a given application's code may be vulnerable to a particular class of attack (e.g., remotely exploitable buffer overflows). In some embodiments, the sensors may also detect that a fault has occurred. Upon predicting that a fault may occur or detecting that a fault has occurred, the portion of the application's code having the faulty instruction or vulnerable function can be isolated, thereby localizing predicted faults at 330.

**[0070]** Alternatively, as shown and discussed in FIG. 4, the one or more sensor may monitor the application until it is caused to abnormally terminate. The system may detect that a fault has occurred, thereby causing the actual application to terminate. As shown in FIG. 4, at 410, the system forces a misbehaving application to abort. In response to the application terminating, the system generates a core dump file or produces other failure-related information, at 420. The core dump file may include, for example, the type of failure and the stack trace when that failure occurred. Based at least in part on the core dump file, the system isolates the portion of the application's code that contains the faulty instruction at 430. Using the core dump file, the system may apply selective emulation to the isolated portion or slice of the application. For example, the system may start with the top-most function in the stack trace.

[0071] Referring back to FIG. 3, in some embodiments, the system may generate an instrumented version of the application (340). For example, an instrumented version of the application may be a copy of a portion of the application's code or all of the application's code. The system may observe instrumented portions of the application. These portions of the application may be selected based on vulnerability to a particular class of attack. The instrumented application may be executed on the server that is currently running the one or more applications, a separate server, a workstation, or any other suitable device.

[0072] Isolating a portion of the application's code and using the emulator on the portion allows the system to reduce and/or minimize the performance impact on the immunized application. However, while this embodiment isolates a portion or a slice of the application's code, the entire application may also be emulated. The emulator may be implemented completely in software, or may take advantage of hardware features of the system processor or architecture, or other facilities offered by the operating system to otherwise reduce and/or minimize the performance impact of monitoring and emulation, and to improve accuracy and effectiveness in handling failures.

[0073] An attempt to exploit such a vulnerability exposes the attack or input vector and other related information (e.g., attacked buffer, vulnerable function, stack trace, etc.). The attack or input vector and other related information can then be used to construct a vaccine. A vaccine can be of various forms, such as, for example, but not limited to, an emulator-based vaccine, source-code changes, binary rewriting, a fix that implements array bounds checking at the machine-instruction level at 350, and/or other fixes based on, for example, the detected type of failure. The vaccine can then be tested in the instrumented application using, for example, an instruction-level emulator (e.g., libtasvm x86 emulator, STEM x86 emulator, etc.) to determine whether the fault was fixed and whether any other functionality (e.g., critical functionality) has been impacted by the fix.

[0074] By continuously testing various vaccines using the instruction-level emulator, the system can verify whether the specific fault has been repaired by running the instrumented application against the event sequence (e.g., input vectors) that caused the specific fault. For example, to verify the effectiveness of a fix, the application may be restarted in a test environment or a sandbox with the instrumentation enabled, and is supplied with the one or more input vectors that caused the failure. A sandbox generally creates an environment in which there are strict limitations on which system resources the instrumented application or a function of the application may request or access.

[0075] At 360, the instruction-level emulator can be selectively invoked for segments of the application's code, thereby allowing the system to mix emulated and non-emulated code within the same code execution. The emulator may be used to, for example, detect and/or monitor for a specific type of failure prior to executing the instruction, record memory modifications during the execution of the instruction (e.g., global variables, library-internal state, libc standard I/O structures, etc.) and the original values, revert the memory stack to its original state, and simulate an error return from a function of the application. That is, upon entering the vulnerable section of the application's code, the instruction-level emulator can capture and store the program state and processes all instructions, including function calls, inside the area designated for emulation. When the program counter references the first instruction outside the bounds of emulation, the virtual processor copies its internal state back to the device processor registers. While registers are updated, memory updates are also applied through the execution of the emulation. The program, unaware of the instructions executed by the virtual processor, continues normal execution on the actual processor.

[0076] In some embodiments, the instruction-level emulator may be linked with the application in advance. Alternatively, in response to a detected failure, the instruction-level emulator may be compiled in the code. In another suitable embodiment, the instruction-level emulator may be invoked in a manner similar to a modern debugger when a particular program instruction is executed. This can take advantage of breakpoint registers and/or other program debugging facilities that the system processor and architecture possess, or it can be a pure-software approach.

[0077] The use of an emulator allows the system to detect and/or monitor a wide array of software failures, such as illegal memory dereferences, buffer overflows, and buffer underflows, and more generic faults, such as divisions by zero. The emulator checks the operands of the instructions it is about to emulate using, at least partially, the vector and related information provided by the one or more sensors that detected the fault. For example, in the case of a division by zero, the emulator checks the value of the operand to the div instruction. In another example, in the case of illegal memory dereferencing, the emulator verifies whether the source and destination address of any memory access (or the program counter for instruction fetches) points to a page that is mapped to the process address space using the mincore() system call, or the appropriate facilities provided by the operating system. In yet another example, in the case of buffer overflow detection, the memory surrounding the vulnerable buffer, as identified by the one or more sensors, is padded by one byte. The emulator then watches for memory writes to these memory locations. This may

require source code availability so as to insert particular variables (e.g., canary variables that launch themselves periodically and perform some typical user transaction to enable transaction-latency evaluation around the clock). The emulator can thus prevent the overflow before it overwrites the remaining locations in the memory stack and recovers the execution. Other approaches for detecting these failures may be incorporated in the system in a modular way, without impacting the high-level operation and characteristics of the system.

**[0078]** For example, the instruction-level emulator may be implemented as a statically-linked C library that defines special tags (e.g., a combination of macros and function calls) that mark the beginning and the end of selective emulation. An example of the tags that are placed around a segment of the application's code for emulation by the instruction-level emulator is shown in FIG. 5. As shown in FIG. 5, the C macro `emulate_init()` moves the program state (general, segment, eflags, and FPU registers) into an emulator-accessible global data structure to capture state immediately before the emulator takes control. The data structure can be used to initialize the virtual registers. `emulate_begin()` obtains the memory location of the first instruction following the call to itself. The instruction address may be the same as the return address and can be found in the activation record of `emulate_begin()`, four bytes above its base stack pointer. The fetch/decode/execute/retire cycle of instructions can continue until either `emulate_end()` is reached or when the emulator detects that control is returning to the parent function. If the emulator does not encounter an error during its execution, the emulator's instruction pointer references the `emulate_term()` macro at completion. To enable the instrumented application to continue execution at this address, the return address of the `emulate_begin()` activation record can be replaced with the current value of the instruction pointer. By executing `emulate_term()`, the emulator's environment can be copied to the program registers and execution continues under normal conditions.

**[0079]** Although the emulator can be linked with the vulnerable application when the source code of the vulnerable application is available, in some embodiments the processor's programmable breakpoint register can be used to invoke the emulator without the running process even being able to detect that it is now running under an emulator.

**[0080]** In addition to monitoring for failures prior to executing instructions and reverting memory changes made by a particular function when a failure occurs (e.g., by having the emulator store memory modifications made during its execution), the emulator can also simulate an error return from the function. For example, some embodiments may generate a map between a set of errors that may occur during an application's execution and a limited set of errors that are explicitly handled by the application's code (sometimes referred to herein as

“error virtualization”). As described below, the error virtualization features may be based on heuristics. However, any suitable approach for determining the return values for a function may be used. For example, aggressive source code analysis techniques to determine the return values that are appropriate for a function may be used. In another example, portions of code of specific functions can be marked as fail-safe and a specific value may be returned when an error return is forced (e.g., for code that checks user permissions). In yet another example, the error value returned for a function that has failed can be determined using information provided by a programmer, system administrator, or any other suitable user.

**[0081]** These error virtualization features allow an application to continue execution even though a boundary condition that was not originally predicted by a programmer allowed a fault to occur. In particular, error virtualization features allows for the application’s code to be retrofitted with an exception catching mechanism, for faults that were unanticipated by the programmer. It should be noted that error virtualization is different from traditional exception handling as implemented by some programming languages, where the programmer must deliberately create exceptions in the program code and also add code to handle these exceptions. Under error virtualization, failures and exceptions that were unanticipated by, for example, the programmer can be caught, and existing application code can be used to handle them. In some embodiments, error virtualization can be implemented through the instruction-level emulator. Alternatively, error virtualization may be implemented through additional source code that is inserted in the application’s source code directly. This insertion of such additional source code can be performed automatically, following the detection of a failure or following the prediction of a failure as described above, or it may be done under the direction of a programmer, system operator, or other suitable user having access to the application’s source code.

**[0082]** Using error virtualization, when an exception occurs during the emulation or if the system detects that a fault has occurred, the system may return the program state to its original settings and force an error return from the currently executing function. To determine the appropriate error value, the system analyzes the declared type of function. In some embodiments, the system may analyze the declared type of function using, for example, a TXL script. Generally, TXL is a hybrid function and rule-based language that may be used for performing source-to-source transformation and for rapidly prototyping new languages and language processors. Based on the declared type of function, the system determines the appropriate error value and places it in the stack frame of the returning function. The appropriate error value may be determined based at least in part on heuristics. For example,

if the return type is an int, a value of -1 is returned. If the return type is an unsigned int, the system returns a 0. If the function returns a pointer, the system determines whether the returned pointer is further dereferenced by the parent function. If the returned pointer is further dereferenced, the system expands the scope of the emulation to include the parent function. In another example, the return error code may be determined using information embedded in the source code of the application, or through additional information provided to the system by the application programmer, system administrator or third party.

**[0083]** In some embodiments, the `emulate_end()` is located and the emulation terminates. Because the emulator saved the state of the application before starting and kept track of memory modification during the application's execution, the system is capable of reversing any memory changes made by the code function inside which the fault occurred by returning it to its original setting, thereby nullifying the effect of the instructions processed through emulation. That is, the emulated portion of the code is sliced off and the execution of the code along with its side effects in terms of changes to memory have been rolled back.

**[0084]** For example, the emulator may not be able to perform system calls directly without kernel-level permissions. Therefore, when the emulator decodes an interruption with an intermediate value of 0x80, the emulator releases control to the kernel. However, before the kernel executes the system call, the emulator can back-up the real registers and replace them with its own values. An INT 0x80 can be issued by the emulator and the kernel processes the system call. Once control returns to the emulator, the emulator can update its registers and restore the original values in the application's registers.

**[0085]** If the instrumented application does not crash after the forced return, the system has successfully found a vaccine for the specific fault, which may be used on the actual application running on the server. At 370, the system can then update the application based at least in part on the emulation.

**[0086]** In accordance with some embodiments, artificial diversity features may be provided to mitigate the security risks of software monoculture.

**[0087]** FIG. 6 is a simplified flowchart illustrating the various steps performed in using an application community to monitor an application for faults and repair the application in accordance with some embodiments. This is a generalized flow chart. It will be understood that the steps shown in FIG. 6 may be performed in any suitable order, some may be deleted, and others added.

**[0088]** Generally, the system may divide an application's code into portions of code at 610. Each portion or slice of the application's code may, for example, be assigned to one of

the members of the application community (e.g., workstation, server, etc.). Each member of the application community may monitor the portion of the code for various types of failures at 620. As described previously, failures include programming errors, exceptions, software faults (e.g., illegal memory accesses, division by zero, buffer overflow attacks, TOCTTOU violations, etc.), threats (e.g., computer viruses, worms, trojans, hackers, key recovery attacks, malicious executables, probes, etc.), and any other suitable fault that may cause abnormal application termination or adversely affect the one or more applications.

**[0089]** For example, the system may divide the portions of code based on the size of the application and the number of members in the application community (i.e., size of the application/members in the application community). Alternatively, the system may divide the portions of code based on the amount of available memory in each of the members of the application community. Any suitable approach for determining how to divide up the application's code may also be used. Some suitable approaches are described hereinafter.

**[0090]** For example, the system may examine the total work in the application community,  $W$ , by examining the cost of executing discrete slices of the application's code. Assuming a set of functions,  $F$ , that comprise an application's callgraph, the  $i^{\text{th}}$  member of  $F$  is denoted as  $f_i$ . The cost of executing each  $f_i$  is a function of the amount of computation present in  $f_i$  (i.e.,  $x_i$ ) and the amount of risk in  $f_i$  (i.e.,  $v_i$ ). The calculation of  $x_i$  can be driven by at least two metrics:  $o_i$ , the number of machine instructions executed as part of  $f_i$ , and  $t_i$ , the amount of time spent executing  $f_i$ . Both  $o_i$  and  $t_i$  may vary as a function of time or application workload according to the application's internal logic. For example, an application may perform logging or cleanup duties after the application passes a threshold number of requests.

**[0091]** In some embodiments, a cost function may be provided in two phases. The first phase calculates the cost due to the amount of computation for each  $f_i$ . The second phase normalizes this cost and applies the risk factor  $v_i$  to determine the final cost of each  $f_i$  and the total amount of work in the system. For example, let

$$T = \sum_{i=1}^N x_i$$

If  $C(f_i, x_i) = x_i/T * 100$ , each cost may be normalized by grouping a subset of  $F$  to represent one unit of work.

[0092] In some embodiments, the system may account for the measure of a function's vulnerability. For example, the system treats  $v_i$  as a discrete variable with a value of  $\alpha$ , where  $\alpha$  takes on a range of values according to the amount of risk such that:

$$v_i = \begin{cases} \alpha & (\text{if } f_i \text{ is vulnerable}) \\ 1 & (\text{if } f_i \text{ is not vulnerable}) \end{cases}$$

Given  $v_i$  for each function, the system may determine the total amount of work in the system and the total number of members needed for monitoring:

$$W = N_{\text{vuln}} = \sum_{i=1}^n v_i * r_i$$

After the system (e.g., a controller) or after each application community member has calculated the amount of work in the system, work units can be distributed. In one example, a central controller or one of the workstations may assign each node approximately  $W/N$  work units. In another suitable example, each member of the application community may determine its own work set. Each member may iterate through the list of work units flipping a coin that is weighted with the value  $v_i * r_i$ . Therefore, if the result of the flip is "true," then the member adds that work unit to its work set.

[0093] Alternatively, the system may generate a list having  $n * W$  slots. Each function can be represented by a number of entries on the list (e.g.,  $v_i * r_i$ ). Every member of the application community can iterate through the list, for example, by randomly selecting true or false. If true, the application community member monitors the function of the application for a given time slice. Because heavily weighted functions have more entries in the list, a greater number of users may be assigned to cover the application. The member may stop when its total work reaches  $W/N$ . Such an approach offers statistical coverage of the application.

[0094] In some embodiments, a distributed bidding approach may be used to distribute the workload of monitoring and repairing an application. Each node in the callgraph  $G$  has a weight  $v_i * r_i$ . Some subset of the nodes in  $F$  is assigned to each application community member such that each member does no more work than  $W/N$  work. The threshold can be relaxed to be within some range  $\epsilon$  of  $W/N$ , where  $\epsilon$  is a measure of system fairness. Upon calculating the globally fair amount of work  $W/N$ , each application community member may adjust its workload by bargaining with other members using a distributed bidding approach.

[0095] Two considerations impact the assignment of work units to application community members. First, the system can allocate work units with higher weights, as these

work units likely have a heavier weight due to a high  $v_i$ . Even if the weight is derived solely from the performance cost, assigning more members to the work units with higher weights is beneficial because these members can round-robin the monitoring task so that any one member does not have to assume the full cost. Second, in some situations,  $v_i * r_i$  may be greater than the average amount of work,  $W/N$ . Achieving fairness means that  $v_i * r_i$  defines the quantity of application community members that is assigned to it and the sum of these quantities defines the minimum number of members in the application community.

[0096] In some embodiments, each application community member calculates a table. An example of such a table is shown in FIG. 7. Upon generating the table, application community members may place bids to adjust each of their respective workloads. For example, the system may use tokens for bidding. Tokens may map directly to the number of time quanta that an application community member is responsible for monitoring a work unit or a function of an application. The system ensures that each node does not accumulate more than the total number of tokens allowed by the choice of  $C$ .

[0097] If an application community member monitors more than its share, then the system has increased coverage and can ensure that faults are detected as quickly as possible. As shown in 630 and 640, each application community member may predict that a fault may occur in the assigned portion of code or may detect that a fault has occurred causing the application to abort, where the assigned portion of the code was the source of the fault. As faults are detected, application members may each proactively monitor assigned portions of code containing the fault to prevent the application from further failures. As discussed previously, the application community member may isolate the portion of the code that caused the fault and use the emulator to test vaccines or fixes. At 650, the application community member that detects or predicts the fault may notify the other application community members. Other application members that have succumbed to the fault may be restarted with the protection mechanisms or fixes generated by the application member that detected the fault.

[0098] Assuming a uniform random distribution of new faults across the application community members, the probability of a fault happening at a member,  $k$ , is:  $P(\text{fault}) = 1/N$ . Thus, the probability of  $k$  detecting a new fault is the probability that the fault happens at  $k$  and that  $k$  detects the fault:  $P(\text{fault at } k \wedge \text{detection}) = 1/N * k_i$ , where  $k_i$  is the percentage of coverage at  $k$ . The probability of the application community detecting the fault is:

$$P(\text{AC detect}) = \sum_{i=1}^N \frac{1}{N} * k_i$$

As each  $k_i$  goes to 100%, the above-equation becomes  $\sum_{i=1}^N \frac{1}{N}$ , or  $N/N$ , a probability of 1 that the fault is detected when it first occurs. It will also be understood that various embodiments may be presented in terms of program procedures executed on a computer or network of computers.

**[0099]** A procedure is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. These steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared and otherwise manipulated. It proves convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like. However, all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities.

**[0100]** Further, the manipulations performed are often referred to in terms, such as adding or comparing, which are commonly associated with mental operations performed by a human operator. No such capability of a human operator is necessary, or desirable in many cases, in any of the operations described herein in connection with various embodiments; the operations are machine operations. Useful machines for performing the operation of various embodiments include general purpose digital computers or similar devices.

**[0101]** Some embodiments also provide apparatuses for performing these operations. These apparatuses may be specially constructed for the required purpose or it may comprise a general purpose computer as selectively activated or reconfigured by a computer program stored in the computer. The procedures presented herein are not inherently related to a particular computer or other apparatus. Various general purpose machines may be used with programs written in accordance with the teachings herein, or it may prove more convenient to construct more specialized apparatus to perform the described method. The required structure for a variety of these machines will appear from the description given.

**[0102]** Some embodiments may include a general purpose computer, or a specially programmed special purpose computer. The user may interact with the system via e.g., a personal computer or over PDA, e.g., the Internet an Intranet, etc. Either of these may be implemented as a distributed computer system rather than a single computer. Similarly, the communications link may be a dedicated link, a modem over a POTS line, the Internet and/or any other method of communicating between computers and/or users. Moreover, the

processing could be controlled by a software program on one or more computer systems or processors, or could even be partially or wholly implemented in hardware.

**[0103]** Although a single computer may be used, systems according to one or more embodiments are optionally suitably equipped with a multitude or combination of processors or storage devices. For example, the computer may be replaced by, or combined with, any suitable processing system operative in accordance with the concepts of various embodiments, including sophisticated calculators, hand held, laptop/notebook, mini, mainframe and super computers, as well as processing system network combinations of the same. Further, portions of the system may be provided in any appropriate electronic format, including, for example, provided over a communication line as electronic signals, provided on CD and/or DVD, provided on optical disk memory, etc.

**[0104]** Any presently available or future developed computer software language and/or hardware components can be employed in such embodiments. For example, at least some of the functionality mentioned above could be implemented using Visual Basic, C, C++ or any assembly language appropriate in view of the processor being used. It could also be written in an object oriented and/or interpretive environment such as Java and transported to multiple destinations to various users.

**[0105]** Other embodiments, extensions, and modifications of the ideas presented above are comprehended and within the reach of one skilled in the field upon reviewing the present disclosure. Accordingly, the scope of the present invention in its various aspects is not to be limited by the examples and embodiments presented above. The individual aspects of the present invention, and the entirety of the invention are to be regarded so as to allow for modifications and future developments within the scope of the present disclosure. For example, the set of features, or a subset of the features, described above may be used in any suitable combination. The present invention is limited only by the claims that follow.

What is claimed is:

1. A method for detecting an anomalous sequence of function calls, comprising:  
  
compressing a sequence of function calls made by the execution of a program using a compression model; and  
  
determining the presence of an anomalous sequence of function calls in the sequence of function calls based on the extent to which the sequence of function calls is compressed.
2. The method of claim 1, further comprising  
  
executing at least one known program;  
  
observing at least one sequence of function calls made by the execution of the at least one known program;  
  
assigning each type of function call in the at least one sequence of function calls made by the at least one known program a unique identifier; and  
  
creating at least part of the compression model by recording at least one sequence of unique identifiers based on the unique identifiers assigned to each type of function call and the observed at least one sequence of function calls.
3. The method of claim 2, wherein the type depends on at least one of the value of at least one argument of a function call, the number of arguments of a function call, and/or the identity of the function.
4. The method of claim 1, wherein the determining further comprises providing an indication of an anomalous function call if the sequence of function calls made by the execution of the program using a compression model is not well compressed.
5. The method of claim 1, wherein the determining further comprises providing an indication of an anomalous function call if the sequence of function calls made by the execution of the program using a compression model is well compressed.

6. The method of claim 1, further comprising modifying at least one of the function calls in the sequence of function calls so that the sequence of function calls becomes non-anomalous.

7. The method of claim 1, further comprising randomly selecting the compression model from a plurality of different compression models.

8. The method of claim 1, wherein the compression model comprises at least two compression models created using at least two different digital processing devices.

9. The method of claim 1, wherein the compressing a sequence of function calls made by the execution of the program is performed using Lempel-Ziv-Welch compression.

10. The method of claim 1, wherein the compressing a sequence of function calls made by the execution of the program comprises assigning each function call an identifier based on identifiers assigned to function calls in the compression model.

11. The method of claim 1, wherein the program is executed in an emulator.

12. The method of claim 1, wherein the execution of the program comprises executing only a part of the program.

13. The method of claim 1, further comprising  
in response to determining the presence of the anomalous sequence of function calls in the sequence of function calls, isolating a portion of the program that made the anomalous sequence of function calls;  
constructing a vaccine that makes the anomalous sequence of function calls non-anomalous;  
testing the vaccine in an instrumented version of the isolated portion of the program; and  
in response to verifying that the vaccine made the anomalous sequence of function calls non-anomalous, updating the program with the vaccine.

14. The method of claim 13, wherein the vaccine includes at least one of an emulator-based vaccine, source-code changes, and binary rewriting.

15. A computer-readable medium containing computer-executable instructions that, when executed by a processor, cause the processor to perform a method for detecting an anomalous sequence of function calls, comprising:

compressing a sequence of function calls made by the execution of a program using a compression model; and

determining the presence of an anomalous sequence of function calls in the sequence of function calls based on the extent to which the sequence of function calls is compressed.

16. The computer-readable medium of claim 15, the method further comprising:  
executing at least one known program;

observing at least one sequence of function calls made by the execution of the at least one known program;

assigning each type of function call in the at least one sequence of function calls made by the at least one known program a unique identifier; and

creating at least part of the compression model by recording at least one sequence of unique identifiers based on the unique identifiers assigned to each type of function call and the observed at least one sequence of function calls.

17. The computer-readable medium of claim 15, the method further comprising providing an indication of an anomalous function call if the sequence of function calls made by the execution of the program using a compression model is not well compressed.

18. The computer-readable medium of claim 15, the method further comprising modifying at least one of the function calls in the sequence of function calls so that the sequence of function calls becomes non-anomalous.

19. The computer-readable medium of claim 15, the method further comprising:  
in response to determining the presence of the anomalous sequence of function calls in the sequence of function calls, isolating a portion of the program that made the anomalous sequence of function calls;

constructing a vaccine that makes the anomalous sequence of function calls non-anomalous;

testing the vaccine in an instrumented version of the isolated portion of the program; and

in response to verifying that the vaccine made the anomalous sequence of function calls non-anomalous, updating the program with the vaccine.

20. A system for detecting an anomalous sequence of function calls, comprising:  
a memory; and  
a processor in communication with the memory; wherein the processor:  
compresses a sequence of function calls made by the execution of a program using a compression model; and  
determines the presence of an anomalous sequence of function calls in the sequence of function calls based on the extent to which the sequence of function calls is compressed.

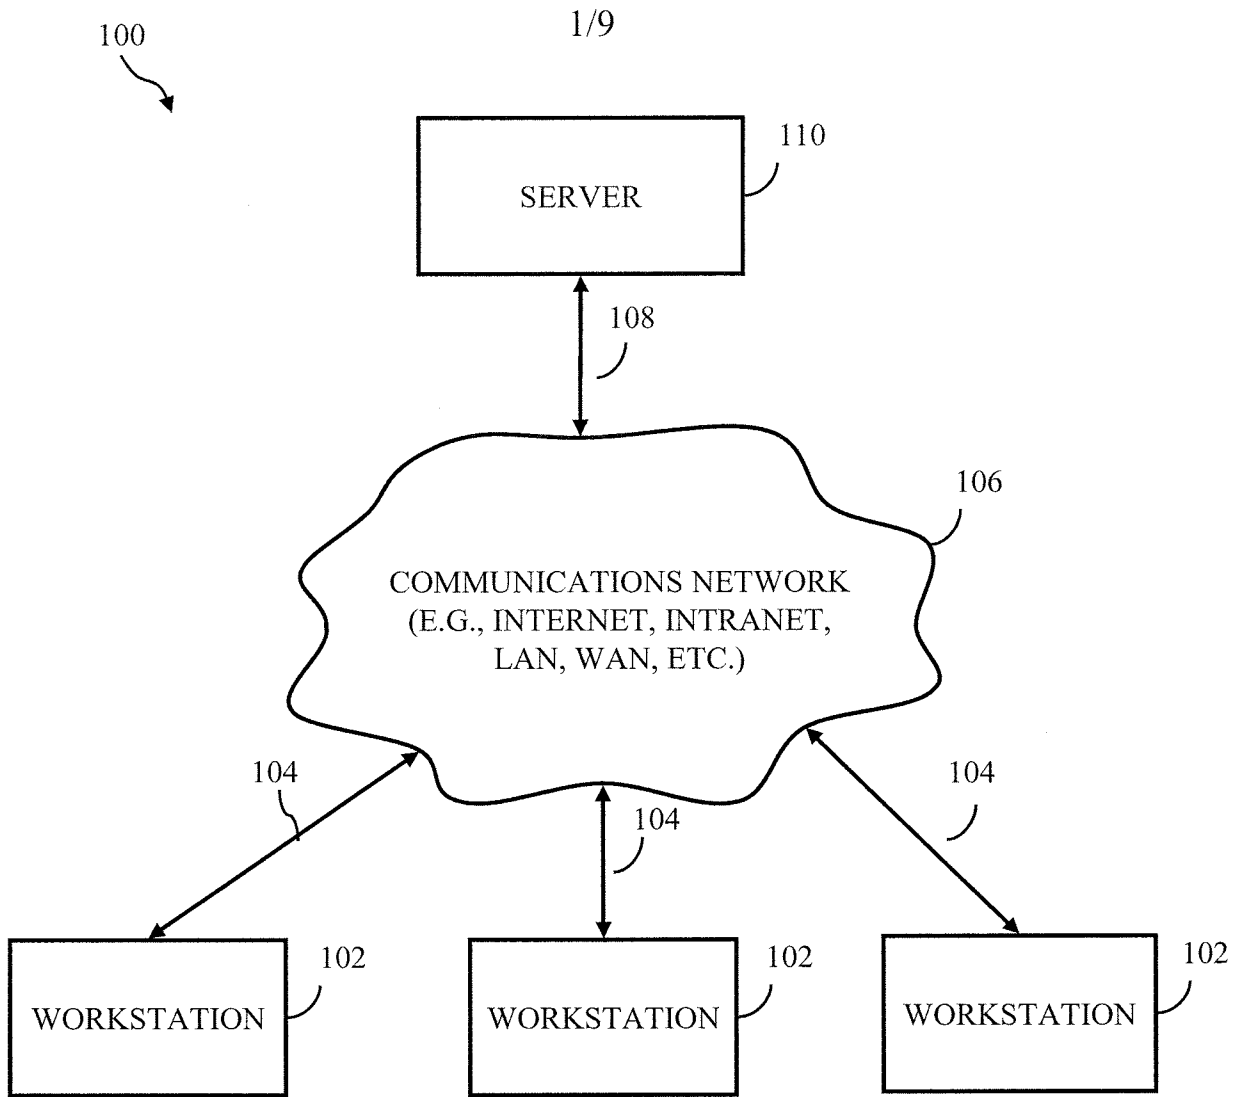
21. The system of claim 20, wherein the processor further:  
executes at least one known program;  
observes at least one sequence of function calls made by the execution of the at least one known program;  
assigns each type of function call in the at least one sequence of function calls made by the at least one known program a unique identifier; and  
creates at least part of the compression model by recording at least one sequence of unique identifiers based on the unique identifiers assigned to each type of function call and the observed at least one sequence of function calls.

22. The system of claim 20, wherein the processor further provides an indication of an anomalous function call if the sequence of function calls made by the execution of the program using a compression model is not well compressed.

23. The system of claim 20, wherein the processor further modifies at least one of the function calls in the sequence of function calls so that the sequence of function calls becomes non-anomalous.

24. The system of claim 20, wherein the processor further:  
in response to determining the presence of the anomalous sequence of function calls in the sequence of function calls, isolates a portion of the program that made the anomalous sequence of function calls;  
constructs a vaccine that makes the anomalous sequence of function calls non-anomalous;  
tests the vaccine in an instrumented version of the isolated portion of the program;  
and  
in response to verifying that the vaccine made the anomalous sequence of function calls non-anomalous, updates the program with the vaccine.

25. The system of claim 24, wherein the vaccine includes at least one of an emulator-based vaccine, source-code changes, and binary rewriting.



**FIG. 1**




FIG. 2

3/9

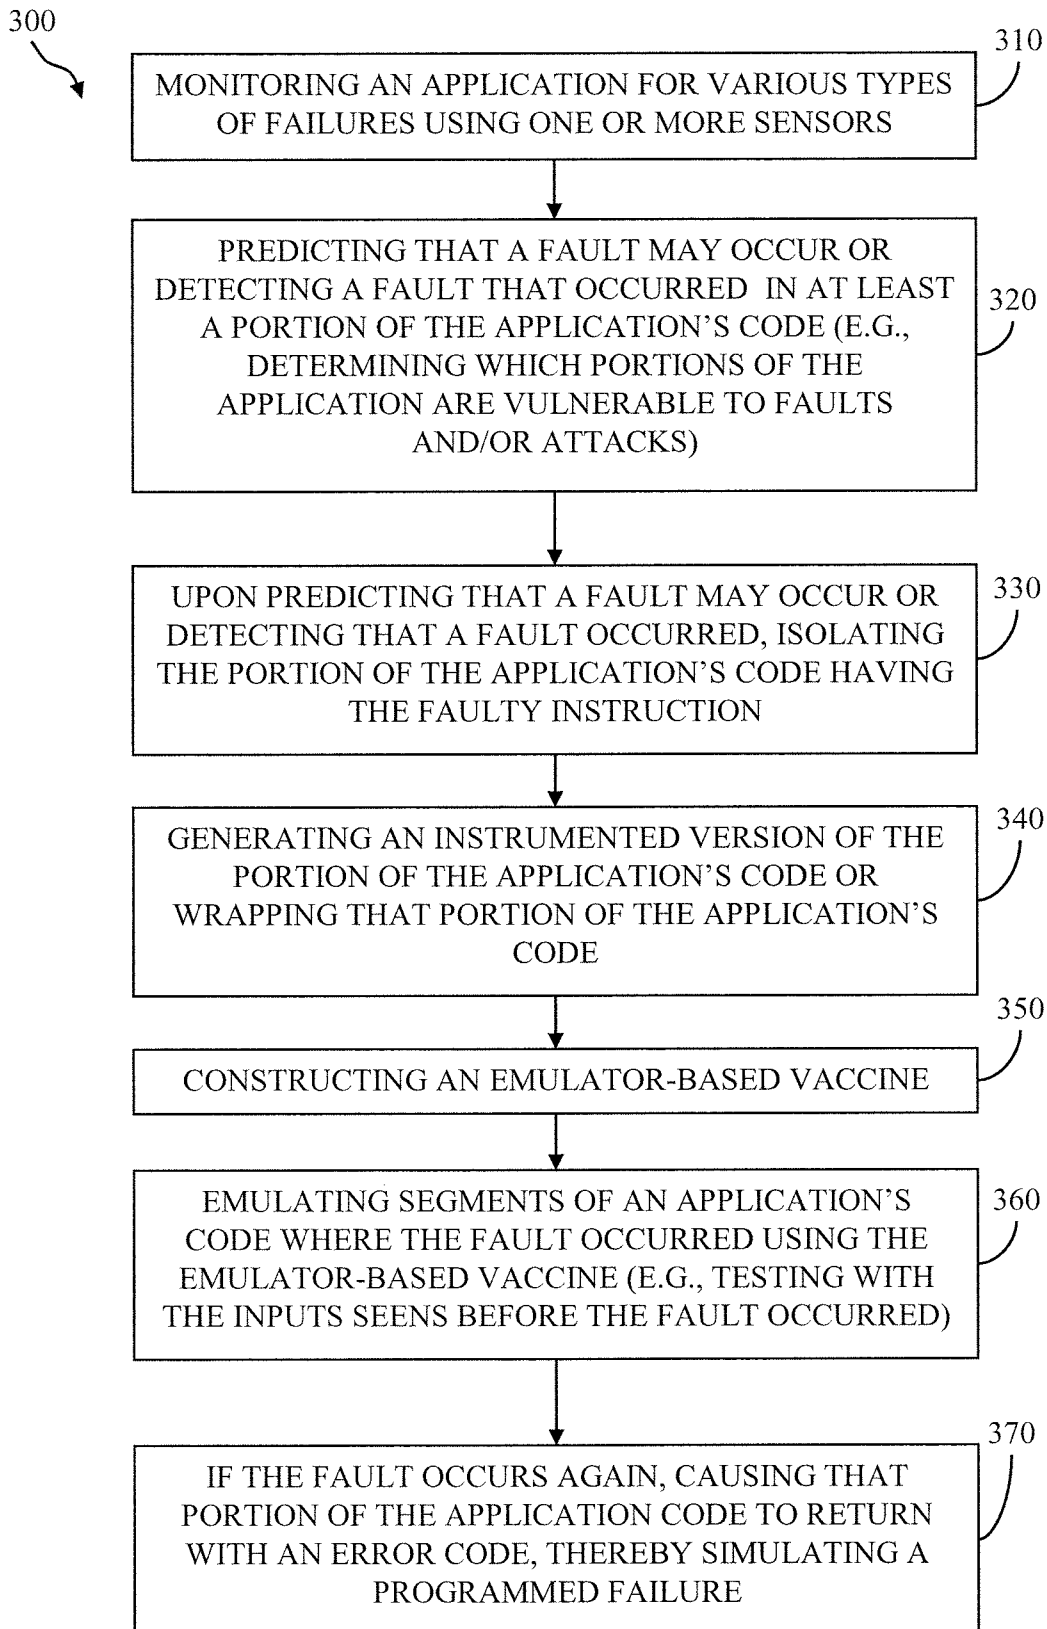


FIG. 3

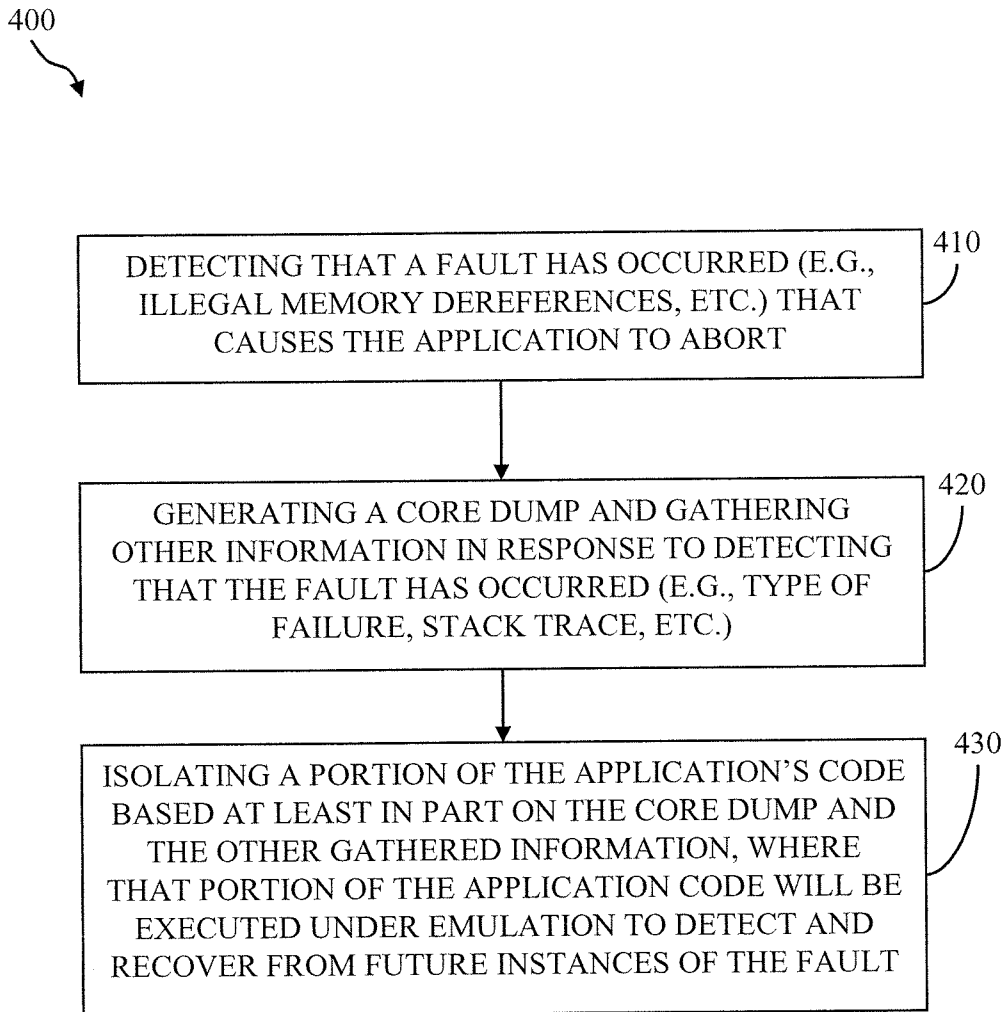



FIG. 4

500



```
void foo() {  
    int a = 1;  
    emulate_init();  
    emulate_begin(p_args);  
    a++;  
    emulate_end();  
    emulate_term();  
    printf("a = %d\n", a);  
}
```

FIG. 5

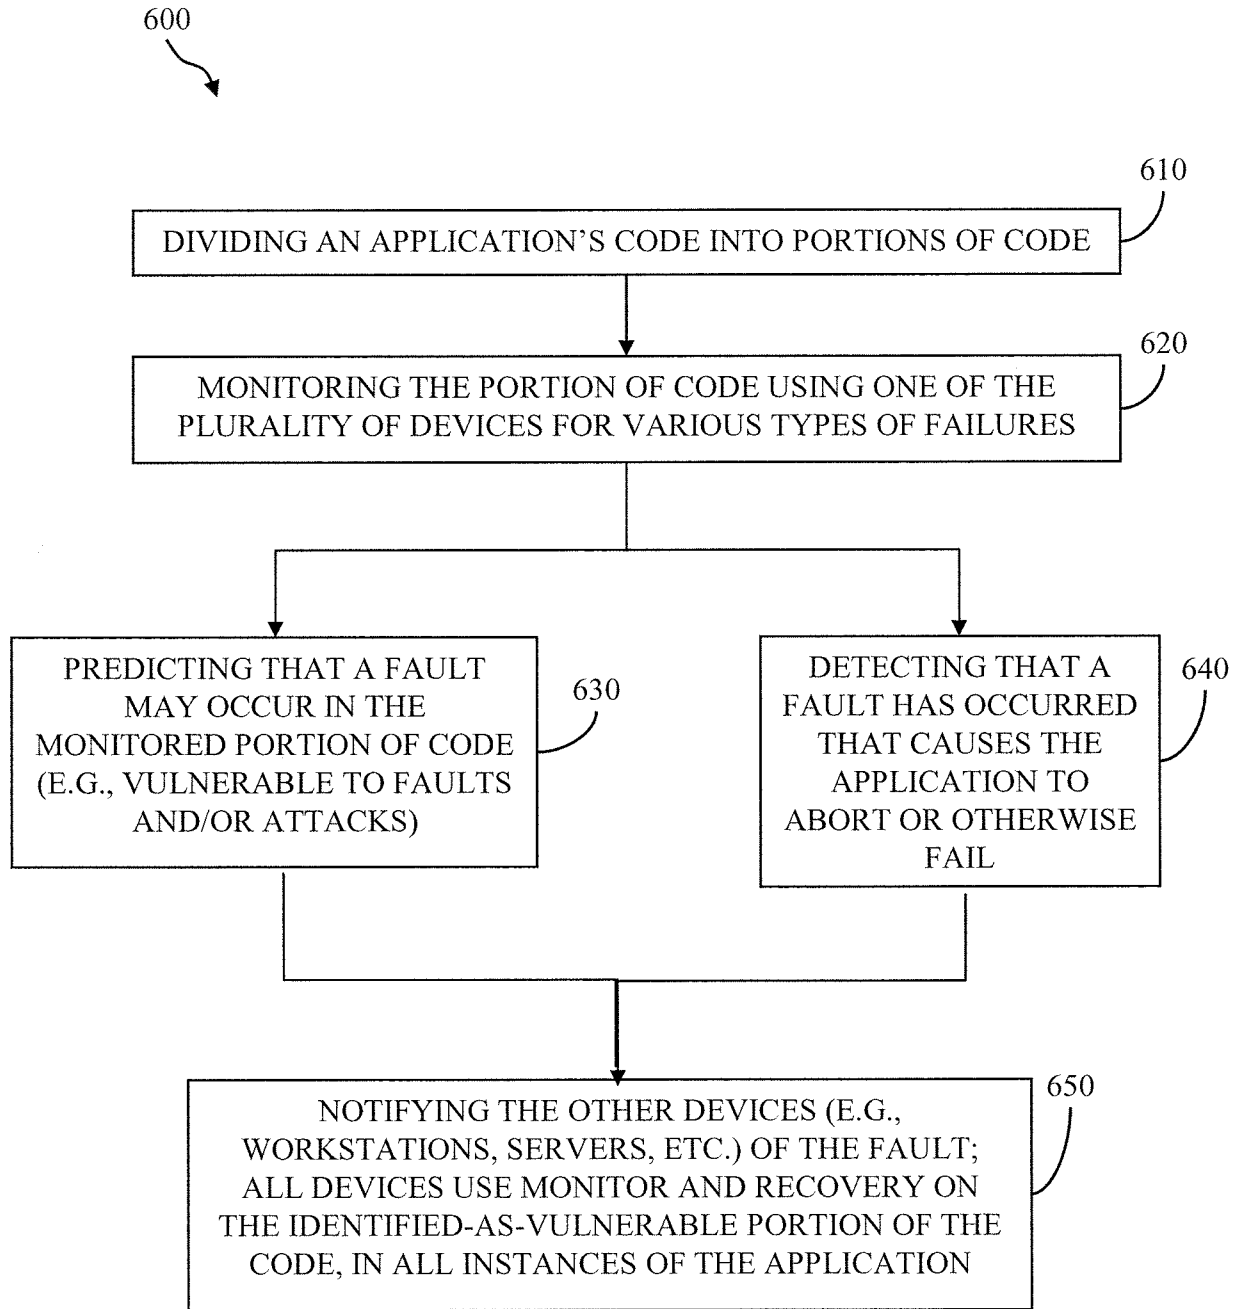


FIG. 6

700  


$f_i$	$x_i$	$r_i$	$v_i$	$T$	$C(f_i, x_i)$	$r_i * v_i$
a0	100	1	$\alpha_1$	600	16	$\alpha_1$
b0	200	2	$\alpha_2$	600	33	$2\alpha_2$
c0	300	3	$\alpha_3$	600	50	$3\alpha_3$

FIG. 7

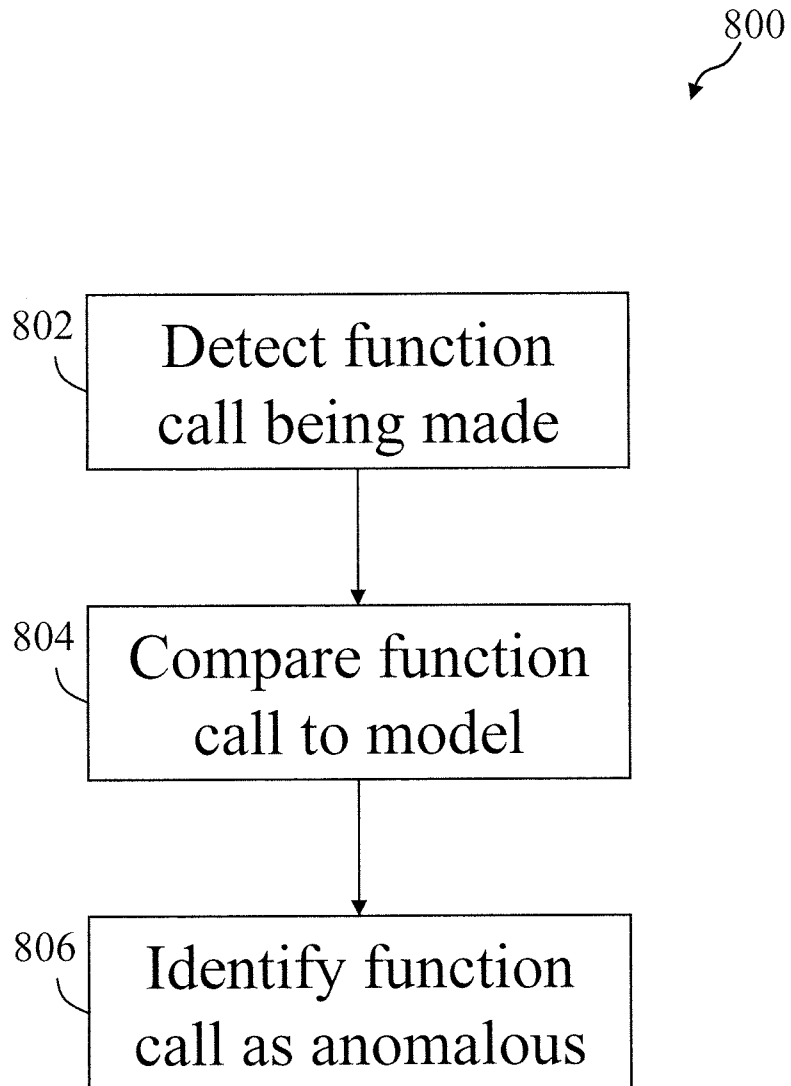


FIG. 8

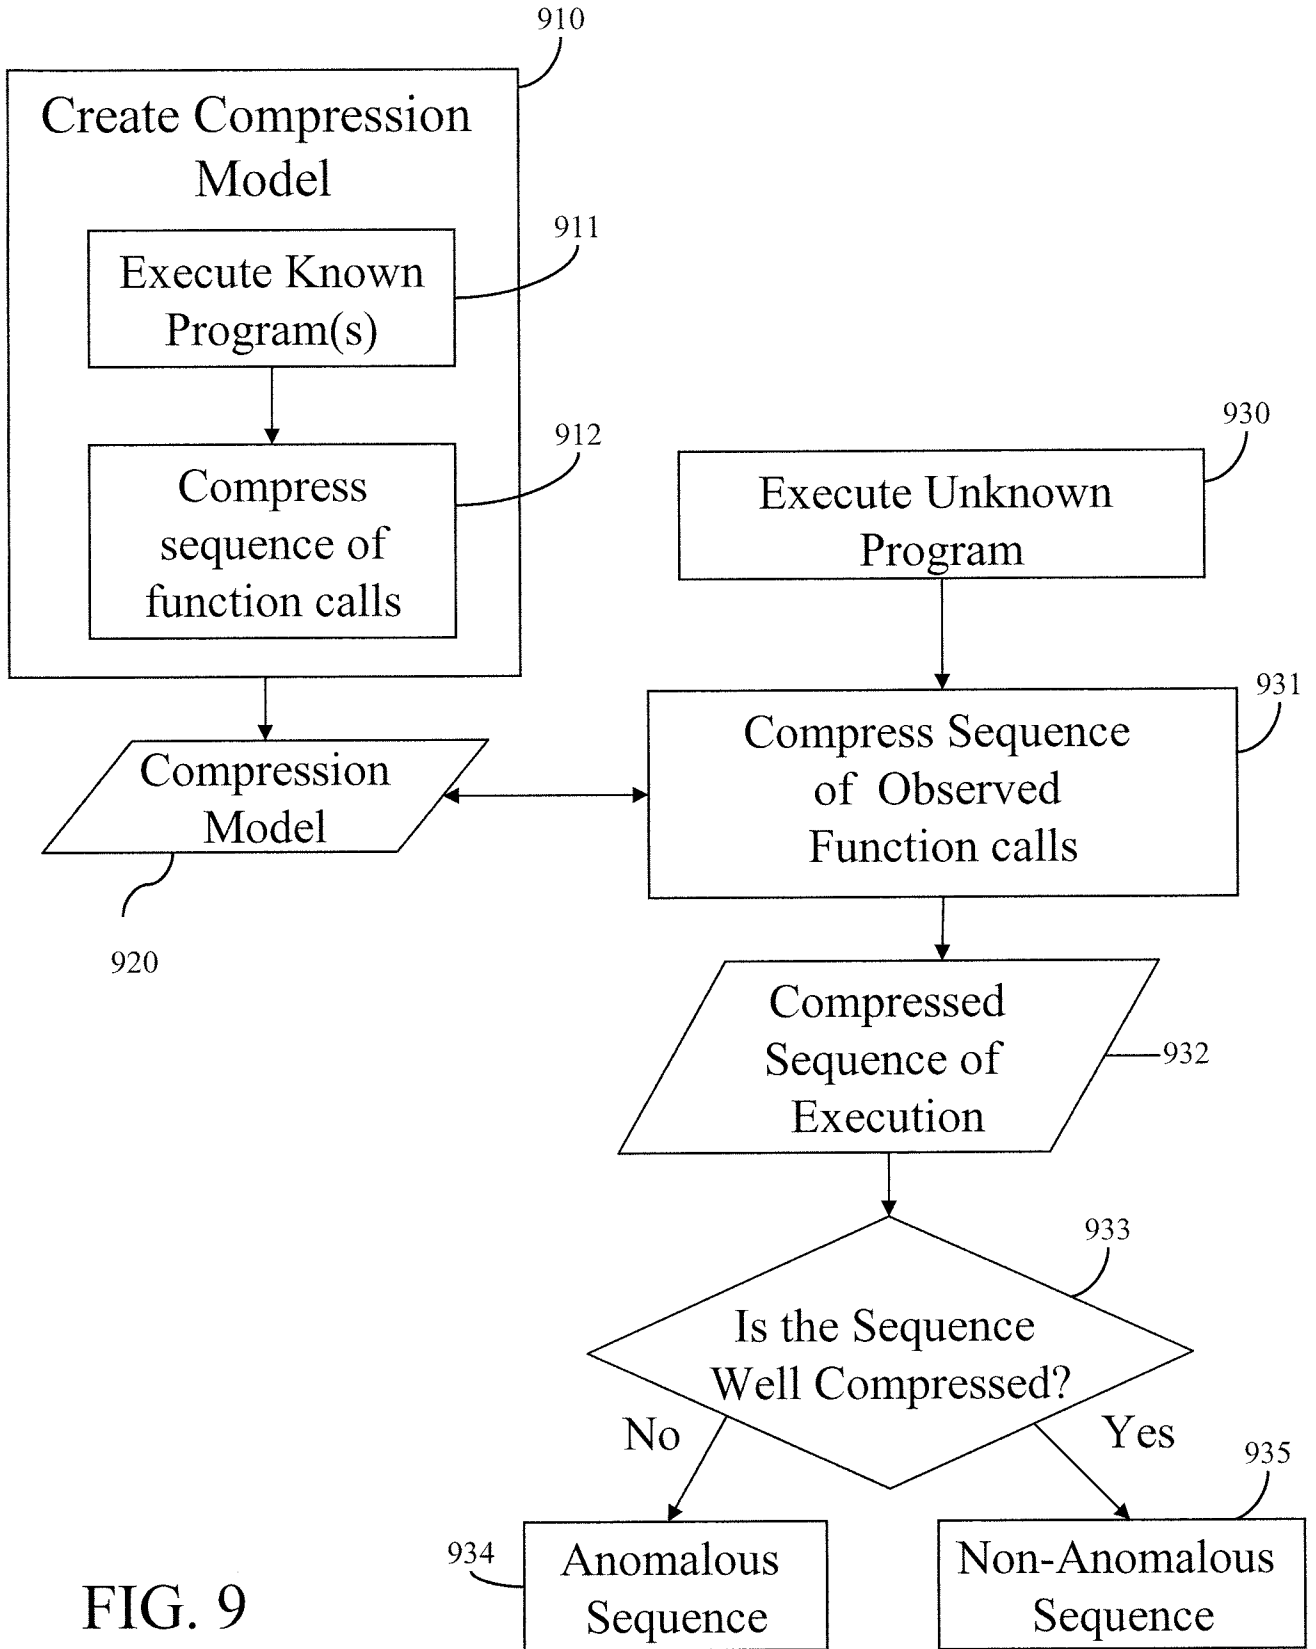


FIG. 9