



(51) International Patent Classification:

H04L 29/06 (2006.01)

(21) International Application Number:

PCT/US2017/016535

(22) International Filing Date:

03 February 2017 (03.02.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/328,224 27 April 2016 (27.04.2016) US

(71) Applicant: ACALVIO TECHNOLOGIES, INC. [US/US]; 2520 Mission College Blvd. #110, Santa Clara, California 95014 (US).

(72) Inventors: SINGH, Satnam; 2520 Mission College Blvd. #110, Santa Clara, California 95014 (US). NEEMA, Nirmesh; 2520 Mission College Blvd. Suite 110, Santa Clara, California 95014 (US). DESAI, Suril; 2520 Mission College Blvd. #110, Santa Clara, California 95014 (US). SAMA, Venkata Babji; 2520 Mission College Blvd. #110, Santa Clara, California 95014 (US). GOPALAKRISHNA, Rajendra; 2520 Mission College Blvd. #110, Santa Clara, California 95014 (US).

(74) Agent: MORSE, Kyle B. et al.; KILPATRICK TOWNSEND & STOCKTON LLP, Mailstop: IP Docketing - 22, 1100 Peachtree Street, Suite 2800, Atlanta, Georgia 30309 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available):

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

Published:

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

(54) Title: CONTEXT-AWARE KNOWLEDGE SYSTEM AND METHODS FOR DEPLOYING DECEPTION MECHANISMS

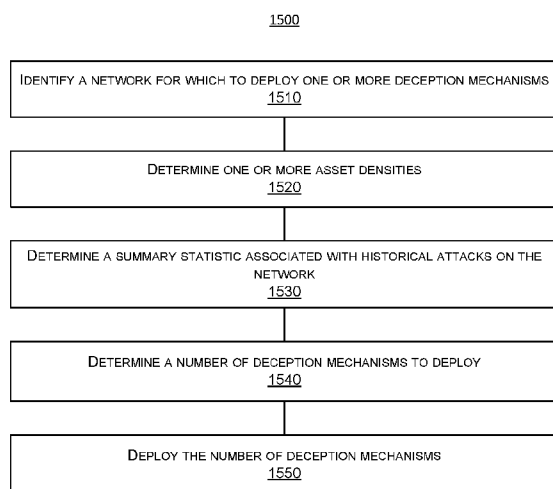


FIG. 15

(57) Abstract: Methods, systems, and computer-readable mediums are described herein to provide context-aware knowledge systems and methods for deploying deception mechanisms. In some examples, a deception profiler can be used to intelligently deploy the deception mechanisms for a network. For example, a method can include identifying a network for which to deploy one or more deception mechanisms (1510). In such an example, a deception mechanism can emulate one or more characteristics of a machine on the network. The method can further include determining one or more asset densities (1520) and a summary statistic (1530). An asset density can be associated with a number of assets connected to the network. The summary statistic can be associated with a number of historical attacks on the network. Using at least one or more of the one or more asset densities, the summary statistic, other information associated with the network, or a combination thereof, the method can further include determining a number of deception mechanisms to deploy (1540), and deploying the number of deception mechanisms (1550).

WO 2017/189071 A1

CONTEXT-AWARE KNOWLEDGE SYSTEM AND METHODS FOR DEPLOYING DECEPTION MECHANISMS

CROSS REFERENCES TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application Number 62/328,224, filed on April 27, 2016, which is incorporated herein by reference in its entirety.

BACKGROUND

[0002] Conventionally, networks were protected from intrusions by attackers using authentication systems. However, as attackers have become more sophisticated, authentication systems have proven ineffective. In response, systems have been created to provide continuous real-time detection of attackers through deception. In particular, deception mechanisms have been created to help lure attackers to access fake systems rather than machines on the network. Unfortunately, the cost of deploying, maintaining, and monitoring deception mechanisms can be onerous. Therefore, there is a need to intelligently create a deception for a network.

BRIEF SUMMARY

[0003] Provided are methods, including computer-implemented methods or methods implemented by a network device, devices including network devices, and computer-program products for deploying one or more deception mechanisms. In some examples, a deception profiler can be used to intelligently deploy the one or more deception mechanisms for a network. For example, a method can include identifying at least a portion of the network for which to deploy the one or more deception mechanisms. In such an example, a deception mechanism can emulate one or more characteristics of a machine on the network. The method can further include determining one or more asset densities and a summary statistic. An asset density can be associated with a number of assets connected to the network and/or the portion of the network. The summary statistic can be associated with a number of historical attacks on the network and/or the portion of the network. Using at least one or more of the one or more asset densities, the summary statistic, other information associated with the network (and/or the portion of the network), or a combination thereof, the method can further include determining a number of deception mechanisms to deploy, and deploying the number of deception mechanisms.

[0004] In some examples, the network can be a subnetwork of a site network. In such examples, the subnetwork can be a logical separation of the site network. In other examples, the network can be a portion of a network. In some examples, a first asset density of the one or more asset densities can be a total number of critical assets in a portion of a network (e.g., a subnetwork, a trunk, one or more machines, or other suitable location in the site network) divided by a total number of critical assets in the network. A second asset density can be a total number of assets in a portion of the network (e.g., a subnetwork, a trunk, one or more machines, or other suitable location in the site network) divided by a total number of assets in the network.

[0005] In some examples, the network can be identified by computing an importance score. The importance score can be computed using the one or more asset densities and machine information associated with the network.

[0006] In some examples, the summary statistic can include a mean, median, or mode of a probability distribution of the historical attacks. The probability distribution of the number of historical attacks can include one or more probability distributions across one or more time periods.

[0007] In some examples, the method can further include updating the summary statistic when a request to access a deception mechanism is received. In some examples, the summary statistic can be updated by determining a new summary statistic using a number of new attacks associated with the network. In such examples, the summary statistic can be updated using the new summary statistic. The new attacks can include attacks that occurred after the historical attacks.

[0008] In some examples, the method can further include computing a population distribution of machine information in the network. The population distribution can identify machine information associated with one or more machines in the network. The method can further include determining one or more configurations for one or more of the deployed deception mechanisms. The one or more configurations can be determined using the population distribution. The machine information can include at least one or more of a MAC address, a host name, an identifier of an operating system, an identifier of a network service, or any combination thereof. A configuration can include at least one or more of a MAC address, a host name, an operating system, a network surface, or any combination thereof.

[0009] In some examples, the method can further include detecting a request to access a deception mechanism, determining a configuration of the accessed deception mechanism, and updating the one or more configurations for the one or more of the deployed deception mechanisms. In such examples, the updating can use the configuration of the accessed deception mechanism.

[0010] In some examples, the method can further include determining a time to deploy a deception mechanism. The determining can use the historical attacks on the network.

[0011] In some examples, historical attacks can be identified using at least one or more of a security operations center, a computer security incident response team, an intrusion detection system, or an intrusion prevention system.

[0012] In another implementation, a method can include identifying a network for which to deploy one or more deception mechanisms. In some examples, a deception mechanism can emulate one or more characteristics of a machine on the network. The method can further include determining a number of deception mechanisms to be deployed for the network. The number of deception mechanisms can be determined using current and historical network information. The method can further include determining one or more configurations for the number of deception mechanisms using a population distribution. In some examples, a configuration can include at least one or more of a MAC address, a host name, an operating system, a network server, or any combination thereof. The population distribution can identify machine information associated with one or more machines in the network. The method can further include determining a schedule to deploy the number of deception mechanisms using the historical network information and deploying the number of deception mechanisms. In some examples, the number of deception mechanisms can be configured, based on the one or more configurations, and deployed using the schedule.

[0013] In some examples, the current network information can include one or more asset densities. An asset density can be associated with a number of assets connected to the network. In some examples, the historical network information can include information on historical attacks on the network.

[0014] In some examples, the schedule can include one or more of a connect time or a disconnect time. The connect time can indicate when to connect a deception mechanism to the network. The disconnect time can indicate when to disconnect a deception mechanism from the network.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] Illustrative examples are described in detail below with reference to the following figures:

[0016] FIG. 1 illustrates an example of a network threat detection and analysis system, in which various implementations of a deception-based security system can be used;

[0017] FIGS. 2A-2D provide examples of different installation configurations that can be used for different customer networks;

[0018] FIG. 3A-3B illustrate examples of customer networks where some of the customer networks' network infrastructure is "in the cloud," that is, is provided by a cloud services provider;

[0019] FIG. 4 illustrates an example of an enterprise network;

[0020] FIG. 5 illustrates a general example of an Internet-of-Things network;

[0021] FIG. 6 illustrates an example of an Internet-of-Things network, here implemented in a private home;

[0022] FIG. 7 illustrates of an Internet-of-Things network, here implemented in a small business;

[0023] FIG. 8 illustrates an example of the basic operation of an industrial control system;

[0024] FIG. 9 illustrates an example of a SCADA system, here used for distributed monitoring and control;

[0025] FIG. 10 illustrates an example of a distributed control;

[0026] FIG. 11 illustrates an example of a PLC implemented in a manufacturing control process;

[0027] FIG. 12 illustrates an example of a network identified by the density engine;

[0028] FIG. 13 illustrates an example of a network interaction graph;

[0029] FIGS. 14A-14C illustrate examples of probability distributions;

[0030] FIG. 15 is a flowchart illustrating an example of a process for deploying a number of deception mechanisms for a network;

[0031] FIG. 16 is a flowchart illustrating an example of a process for updating a probability distribution in order to change a number of deception mechanisms deployed;

[0032] FIG. 17 illustrates examples of host names being determined for deception mechanisms;

[0033] FIG. 18 is a flowchart illustrating an example of a process for updating one or more configurations for one or more deception mechanisms;

[0034] FIG. 19A illustrates an example of a categorical distribution;

[0035] FIG. 19B illustrates an example of a Dirichlet distribution;

[0036] FIG. 20 illustrates an example of a probability distribution of a number of attacks over time; and

[0037] FIG. 21 is a flowchart illustrating another example of a process for intelligently deploying deception mechanisms for a network.

DETAILED DESCRIPTION

[0038] Network deception mechanisms, often referred to as “honeypots,” “honey tokens,” and “honey nets,” among others, defend a network from threats by distracting or diverting the threat. Honeypot-type deception mechanisms can be installed in a network for a particular site, such as a business office, to act as decoys in the site’s network. Honeypot-type deception mechanisms are typically configured to be indistinguishable from active, production systems in the network. Additionally, such deception mechanisms are typically configured to be attractive to a network threat by having seemingly valuable data and/or by appearing vulnerable to infiltration. Though these deception mechanisms can be indistinguishable from legitimate parts of the site network, deception mechanisms are not part of the normal operation of the network, and would not be accessed during normal, legitimate use of the site network. Because normal users of the site network would not normally use or access a deception mechanism, any use or access to the deception mechanism is suspected to be a threat to the network.

[0039] “Normal” operation of a network generally includes network activity that conforms with the intended purpose of a network. For example, normal or legitimate network activity can include the operation of a business, medical facility, government office, education institution, or the ordinary network activity of a private home. Normal network activity can

also include the non-business-related, casual activity of users of a network, such as accessing personal email and visiting websites on personal time, or using network resources for personal use. Normal activity can also include the operations of network security devices, such as firewalls, anti-virus tools, intrusion detection systems, intrusion protection systems, email filters, adware blockers, and so on. Normal operations, however, exclude deception mechanisms, in that deception mechanisms are not intended to take part in business operations or casual use. As such, network users and network systems do not normally access deception mechanisms except perhaps for the most routine network administrative tasks. Access to a deception mechanism, other than entirely routine network administration, may thus indicate a threat to the network.

[0040] Threats to a network can include active attacks, where an attacker interacts or engages with systems in the network to steal information or do harm to the network. An attacker may be a person, or may be an automated system. Examples of active attacks include denial of service (DoS) attacks, distributed denial of service (DDoS) attacks, spoofing attacks, “man-in-the-middle” attacks, attacks involving malformed network requests (e.g. Address Resolution Protocol (ARP) poisoning, “ping of death,” etc.), buffer, heap, or stack overflow attacks, and format string attacks, among others. Threats to a network can also include self-driven, self-replicating, and/or self-triggering malicious software. Malicious software can appear innocuous until activated, upon which the malicious software may attempt to steal information from a network and/or do harm to the network. Malicious software is typically designed to spread itself to other systems in a network. Examples of malicious software include ransomware, viruses, worms, Trojan horses, spyware, keyloggers, rootkits, and rogue security software, among others.

[0041] However, determining a location to deploy a deception mechanism to handle with threats to a network can be difficult. If a deception mechanism is deployed in the wrong location, it is unlikely that an attacker will attempt to access the deception mechanism. On the other hand, if a deception mechanism is deployed in every location, the cost of deploying, maintaining, and monitoring the deception mechanisms can be onerous. Determining a number of deception mechanisms to deploy can also be difficult. If too few deception mechanisms are deployed, it is unlikely that an attacker will attempt to access one of the deception mechanisms. On the other hand, if too many deception mechanisms are deployed, the cost of deploying, maintaining, and monitoring the deception mechanisms can be onerous. Further, determining a configuration for a deception mechanism can be difficult.

Some configurations can cause an attacker to realize that the deception mechanism is fake. On the other hand, some configurations can cause an attacker to not be interested in the deception mechanism. Determining a schedule to deploy a deception mechanism can also be difficult. If the deception mechanism is deployed too late, an attacker might not even see the deception mechanism. On the other hand, if a deception mechanism is deployed at times abnormal to other machines, an attacker might realize that the deception mechanism is fake.

[0042] Provided herein are methods, including computer-implemented methods or methods implemented by a network device, devices including network devices, and computer-program products for deploying one or more deception mechanisms for a network. Examples herein relate to determining a deception mechanism deployment (or deployment profile) for a network. The deployment can include a number of deception mechanisms to deploy for a network, one or more configurations for the number of deception mechanisms, a location to place the number of deception mechanisms, a schedule to deploy the number of deception mechanisms, or any combination thereof.

I. DECEPTION-BASED SECURITY SYSTEMS

[0043] FIG. 1 illustrates an example of a network threat detection and analysis system 100, in which various implementations of a deception-based security system can be used. The network threat detection and analysis system 100, or, more briefly, network security system 100, provides security for a site network 104 using deceptive security mechanisms, a variety of which may be called “honeypots.” The deceptive security mechanisms may be controlled by and inserted into the site network 104 using a deception center 108 and sensors 110, which may also be referred to as deception sensors, installed in the site network 104. In some implementations, the deception center 108 and the sensors 110 interact with a security services provider 106 located outside of the site network 104. The deception center 108 may also obtain or exchange data with sources located on the Internet 150.

[0044] Security mechanisms designed to deceive, sometimes referred to as “honeypots,” may also be used as traps to divert and/or deflect unauthorized use of a network away from the real network assets. A deception-based security mechanism may be a computer attached to the network, a process running on one or more network systems, and/or some other device connected to the network. A security mechanism may be configured to offer services, real or emulated, to serve as bait for an attack on the network. Deception-based security mechanisms that take the form of data, which may be called “honey tokens,” may be mixed in with real

data in devices in the network. Alternatively or additionally, emulated data may also be provided by emulated systems or services.

[0045] Deceptive security mechanisms can also be used to detect an attack on the network. Deceptive security mechanisms are generally configured to appear as if they are legitimate parts of a network. These security mechanisms, however, are not, in fact, part of the normal operation of the network. Consequently, normal activity on the network is not likely to access the security mechanisms. Thus any access over the network to the security mechanism is automatically suspect.

[0046] The network security system 100 may deploy deceptive security mechanisms in a targeted and dynamic fashion. Using the deception center 108 the system 100 can scan the site network 104 and determine the topology of the site network 104. The deception center 108 may then determine devices to emulate with security mechanisms, including the type and behavior of the device. The security mechanisms may be selected and configured specifically to attract the attention of network attackers. The security mechanisms may also be selected and deployed based on suspicious activity in the network. Security mechanisms may be deployed, removed, modified, or replaced in response to activity in the network, to divert and isolate network activity related to an apparent attack, and to confirm that the network activity is, in fact, part of a real attack.

[0047] The site network 104 is a network that may be installed among the buildings of a large business, in the office of a small business, at a school campus, at a hospital, at a government facility, or in a private home. The site network 104 may be described as a local area network (LAN) or a group of LANS. The site network 104 may be one site belonging to an organization that has multiple site networks 104 in one or many geographical locations. In some implementations, the deception center 108 may provide network security to one site network 104, or to multiple site networks 104 belonging to the same entity.

[0048] The site network 104 is where the networking devices and users of the an organizations network may be found. The site network 104 may include network infrastructure devices, such as routers, switches hubs, repeaters, wireless base stations, and/or network controllers, among others. The site network 104 may also include computing systems, such as servers, desktop computers, laptop computers, tablet computers, personal digital assistants, and smart phones, among others. The site network 104 may also include

other analog and digital electronics that have network interfaces, such as televisions, entertainment systems, thermostats, refrigerators, and so on.

[0049] The deception center 108 provides network security for the site network 104 (or multiple site networks for the same organization) by deploying security mechanisms into the site network 104, monitoring the site network 104 through the security mechanisms, detecting and redirecting apparent threats, and analyzing network activity resulting from the apparent threat. To provide security for the site network 104, in various implementations the deception center 108 may communicate with sensors 110 installed in the site network 104, using network tunnels 120. As described further below, the tunnels 120 may allow the deception center 108 to be located in a different sub-network (“subnet”) than the site network 104, on a different network, or remote from the site network 104, with intermediate networks (possibly including the Internet 150) between the deception center 108 and the site network 104.

[0050] In some implementations, the network security system 100 includes a security services provider 106. In these implementations, the security services provider 106 may act as a central hub for providing security to multiple site networks, possibly including site networks controlled by different organizations. For example, the security services provider 106 may communicate with multiple deception centers 108 that each provide security for a different site network 104 for the same organization. In some implementations, the security services provider 106 is located outside the site network 104. In some implementations, the security services provider 106 is controlled by a different entity than the entity that controls the site network. For example, the security services provider 106 may be an outside vendor. In some implementations, the security services provider 106 is controlled by the same entity as that controls the site network 104.

[0051] In some implementations, when the network security system 100 includes a security services provider 106, the sensors 110 and the deception center 108 may communicate with the security services provider 106 in order to be connected to each other. For example, the sensors 110, which may also be referred to as deception sensors, may, upon powering on in the site network 104, send information over a network connection 112 to the security services provider 106, identifying themselves and the site network 104 in which they are located. The security services provider 106 may further identify a corresponding deception center 108 for the site network 104. The security services provider 106 may then provide the network location of the deception center 108 to the sensors 110, and may provide the deception center

108 with the network location of the sensors 110. A network location may take the form of, for example, an Internet Protocol (IP) address. With this information, the deception center 108 and the sensors 110 may be able to configure tunnels 120 to communicate with each other.

[0052] In some implementations, the network security system 100 does not include a security services provider 106. In these implementations, the sensors 110 and the deception center 108 may be configured to locate each other by, for example, sending packets that each can recognize as coming from the other. Using these packets, the sensors 110 and deception center 108 may be able to learn their respective locations on the network. Alternatively or additionally, a network administrator can configure the sensors 110 with the network location of the deception center 108, and vice versa.

[0053] In various implementations, the sensors 110 are a minimal combination of hardware and/or software, sufficient to form a network connection with the site network 104 and a tunnel 120 with the deception center 108. For example, a sensor 110 may be constructed using a low-power processor, a network interface, and a simple operating system. In various implementations, the sensors 110 provide the deception center 108 with visibility into the site network 104, such as for example being able to operate as a node in the site network 104, and/or being able to present or project deceptive security mechanisms into the site network 104, as described further below. Additionally, in various implementations, the sensors 110 may provide a portal through which a suspected attack on the site network 104 can be redirected to the deception center 108, as is also described below.

[0054] In various implementations, the deception center 108 may be configured to profile the site network 104, deploy deceptive security mechanisms for the site network 104, detect suspected threats to the site network 104, analyze the suspected threat, and analyze the site network 104 for exposure and/or vulnerability to the supposed threat.

[0055] To provide the site network 104, the deception center 108 may include a deception profiler 130. In various implementations, the deception profiler may 130 derive information 114 from the site network 104, and determine, for example, the topology of the site network 104, the network devices included in the site network 104, the software and/or hardware configuration of each network device, and/or how the network is used at any given time. Using this information, the deception profiler 130 may determine one or more deceptive security mechanisms to deploy into the site network 104.

[0056] In some examples, the deception profiler 130 may be configured to analyze the site network to determine which deception mechanisms to deploy into the site network, where to deploy them, and/or when to deploy them. The deception profiler 130 may receive network information 114 from the site network. This network information 114 may include information such as subnet addresses, IP addresses in use, an identity and/or configuration of devices in the site network, and/or profiles of usage patterns of the devices in the site network. Using this information, the deception profiler 130 may configure one or more deception mechanisms. For example, the deception profiler 130 may instruct the network emulator 120 to reconfigure the emulated network 116.

[0057] The deception profiler 130 in this example includes a location engine 132, a density engine 134, a configuration engine 136, and a scheduling engine 138. The location engine 132 may determine where in the site network to deploy deception mechanisms. The density engine 134 may determine how many deception mechanisms to deploy. The configuration engine 136 may determine how each deception mechanism is to be configured, and may provide configurations to the network emulator 120. The scheduling engine 138 may determine when a deception mechanism should be deployed and/or activated. The components of the deception profiler 130 are described in further detail below.

[0058] In various implementations, the deception profiler may configure an emulated network 116 to emulate one or more computing systems. Using the tunnels 120 and sensors 110, the emulated computing systems may be projected into the site network 104, where they serve as deceptions. The emulated computing systems may include address deceptions, low-interaction deceptions, and/or high-interaction deceptions. In some implementations, the emulated computing systems may be configured to resemble a portion of the network. In these implementations, this network portion may then be projected into the site network 104.

[0059] In various implementations, a network threat detection engine 140 may monitor activity in the emulated network 116, and look for attacks on the site network 104. For example, the network threat detection engine 140 may look for unexpected access to the emulated computing systems in the emulated network 116. The network threat detection engine 140 may also use information 114 extracted from the site network 104 to adjust the emulated network 116, in order to make the deceptions more attractive to an attack, and/or in response to network activity that appears to be an attack. Should the network threat detection engine 140 determine that an attack may be taking place, the network threat detection engine

140 may cause network activity related to the attack to be redirected to and contained within the emulated network 116.

[0060] In various implementations, the emulated network 116 is a self-contained, isolated, and closely monitored network, in which suspect network activity may be allowed to freely interact with emulated computing systems. In various implementations, questionable emails, files, and/or links may be released into the emulated network 116 to confirm that they are malicious, and/or to see what effect they have. Outside actors can also be allowed to access emulated system, steal data and user credentials, download malware, and conduct any other malicious activity. In this way, the emulated network 116 not only isolated a suspected attack from the site network 104, but can also be used to capture information about an attack. Any activity caused by suspect network activity may be captured in, for example, a history of sent and received network packets, log files, and memory snapshots.

[0061] In various implementations, activity captured in the emulated network 116 may be analyzed using a targeted threat analysis engine 160. The threat analysis engine 160 may examine data collected in the emulated network 116 and reconstruct the course of an attack. For example, the threat analysis engine 160 may correlate various events seen during the course of an apparent attack, including both malicious and innocuous events, and determine how an attacker infiltrated and caused harm in the emulated network 116. In some cases, the threat analysis engine 160 may use threat intelligence 152 from the Internet 150 to identify and/or analyze an attack contained in the emulated network 116. The threat analysis engine 160 may also confirm that suspect network activity was not an attack. The threat analysis engine 160 may produce indicators that describe the suspect network activity, including indicating whether the suspect activity was or was not an actual threat. The threat analysis engine 160 may share these indicators with the security community 180, so that other networks can be defended from the attack. The threat analysis engine 160 may also send the indicators to the security services provider 106, so that the security services provider 106 can use the indicators to defend other site networks.

[0062] In various implementations, the threat analysis engine 160 may also send threat indicators, or similar data, to a behavioral analytics engine 170. The behavioral analytics engine 170 may be configured to use the indicators to probe 118 the site network 104, and see whether the site network 104 has been exposed to the attack, or is vulnerable to the attack. For example, the behavioral analytics engine 170 may search the site network 104 for

computing systems that resemble emulated computing systems in the emulated network 116 that were affected by the attack. In some implementations, the behavioral analytics engine 170 can also repair systems affected by the attack, or identify these systems to a network administrator. In some implementations, the behavioral analytics engine 170 can also reconfigure the site network's 104 security infrastructure to defend against the attack.

[0063] The behavioral analytics engine 170 can work in conjunction with a Security Information and Event Management (SIEM) 172 system. In various implementations, SIEM includes software and/or services that can provide real-time analysis of security alerts generated by network hardware and applications. In various implementations, the deception center 108 can communicate with the SIEM 172 system to obtain information about computing and/or networking systems in the site network 104.

[0064] Using deceptive security mechanisms, the network security system 100 may thus be able to distract and divert attacks on the site network 104. The network security system 100 may also be able to allow, using the emulated network 116, an attack to proceed, so that as much can be learned about the attack as possible. Information about the attack can then be used to find vulnerabilities in the site network 104. Information about the attack can also be provided to the security community 180, so that the attack can be thwarted elsewhere.

II. CUSTOMER INSTALLATIONS

[0065] The network security system, such as the deception-based system described above, may be flexibly implemented to accommodate different customer networks. FIGS. 2A-2D provide examples of different installation configurations 200a-200d that can be used for different customer networks 202. A customer network 202 may generally be described as a network or group of networks that is controlled by a common entity, such as a business, a school, or a person. The customer network 202 may include one or more site networks 204. The customer network's 202 site networks 204 may be located in one geographic location, may be behind a common firewall, and/or may be multiple subnets within one network. Alternatively or additionally, a customer network's 202 site networks 204 may be located in different geographic locations, and be connected to each other over various private and public networks, including the Internet 250.

[0066] Different customer networks 202 may have different requirements regarding network security. For example, some customer networks 202 may have relatively open connections to outside networks such as the Internet 250, while other customer networks 202

have very restricted access to outside networks. The network security system described in FIG. 1 may be configurable to accommodate these variations.

[0067] FIG. 2A illustrates one example of an installation configuration 200a, where a deception center 208 is located within the customer network 202. In this example, being located within the customer network 202 means that the deception center 208 is connected to the customer network 202, and is able to function as a node in the customer network 202. In this example, the deception center 208 may be located in the same building or within the same campus as the site network 204. Alternatively or additionally, the deception center 208 may be located within the customer network 202 but at a different geographic location than the site network 204. The deception center 208 thus may be within the same subnet as the site network 204, or may be connected to a different subnet within the customer network.

[0068] In various implementations, the deception center 208 communicates with sensors 210, which may also be referred to as deception sensors, installed in the site network over network tunnels 220. In this example, the network tunnels 220 may cross one or more intermediate within the customer network 202.

[0069] In this example, the deception center 208 is able to communicate with a security services provider 206 that is located outside the customer network 202, such as on the Internet 250. The security services provider 206 may provide configuration and other information for the deception center 208. In some cases, the security services provider 206 may also assist in coordinating the security for the customer network 202 when the customer network 202 includes multiple site networks 204 located in various geographic areas.

[0070] FIG. 2B illustrates another example of an installation configuration 200b, where the deception center 208 is located outside the customer network 202. In this example, the deception center 208 may be connected to the customer network 202 over the Internet 250. In some implementations, the deception center 208 may be co-located with a security services provider, and/or may be provided by the security services provider.

[0071] In this example, the tunnels 220 connect the deception center 208 to the sensors 210 through a gateway 262. A gateway is a point in a network that connects the network to another network. For example, in this example, the gateway 262 connects the customer network 202 to outside networks, such as the Internet 250. The gateway 262 may provide a firewall, which may provide some security for the customer network 202. The tunnels 220 may be able to pass through the firewall using a secure protocol, such as Secure Socket Shell

(SSH) and similar protocols. Secure protocols typically require credentials, which may be provided by the operator of the customer network 202.

[0072] FIG. 2C illustrates another example of an installation configuration 200c, where the deception center 208 is located inside the customer network 202 but does not have access to outside networks. In some implementations, the customer network 202 may require a high level of network security. In these implementations, the customer network's 202 connections to the other networks may be very restricted. Thus, in this example, the deception center 208 is located within the customer network 202, and does not need to communicate with outside networks. The deception center 208 may use the customer network's 202 internal network to coordinate with and establish tunnels 220 to the sensors 210. Alternatively or additionally, a network administrator may configure the deception center 208 and sensors 210 to enable them to establish the tunnels 220.

[0073] FIG. 2D illustrates another example of an installation configuration 200d. In this example, the deception center 208 is located inside the customer network 202, and further is directly connected to the site network 204. Directly connected, in this example, can mean that the deception center 208 is connected to a router, hub, switch, repeater, or other network infrastructure device that is part of the site network 204. Directly connected can alternatively or additionally mean that the deception center 208 is connected to the site network 204 using a Virtual Local Area Network (VLAN). For example, the deception center 208 can be connected to VLAN trunk port. In these examples, the deception center 208 can project deceptions into the site network 204 with or without the use of sensors, such as are illustrated in FIGS. 2A-2C.

[0074] In the example of FIG. 2D, the deception center 208 can also optionally be connected to an outside security services provider 206. The security services provider 206 can manage the deception center 208, including providing updated security data, sending firmware upgrades, and/or coordinating different deception centers 208 for different site networks 204 belonging to the same customer network 202. In some implementations, the deception center 208 can operate without the assistances of an outside security services provider 206.

III. CUSTOMER NETWORKS

[0075] The network security system, such as the deception-based system discussed above, can be used for variety of customer networks. As noted above, customer networks can come

in wide variety of configurations. For example, a customer network may have some of its network infrastructure “in the cloud.” A customer network can also include a wide variety of devices, including what may be considered “traditional” network equipment, such as servers and routers, and non-traditional, “Internet-of-Things” devices, such as kitchen appliances. Other examples of customer networks include established industrial networks, or a mix of industrial networks and computer networks.

[0076] FIG. 3A-3B illustrate examples of customer networks 302a-302b where some of the customer networks' 302a-302b network infrastructure is “in the cloud,” that is, is provided by a cloud services provider 354. These example customer networks 302a-302b may be defended by a network security system that includes a deception center 308 and sensors 310, which may also be referred to as deception sensors, and may also include an off-site security services provider 306.

[0077] A cloud services provider is a company that offers some component of cloud computer – such as Infrastructure as a Service (IaaS), Software as a Service (SaaS) or Platform as Service (PaaS) – to other businesses and individuals. A cloud services provider may have a configurable pool of computing resources, including, for example, networks, servers, storage, applications, and services. These computing resources can be available on demand, and can be rapidly provisioned. While a cloud services provider's resources may be shared between the cloud service provider's customers, from the perspective of each customer, the individual customer may appear to have a private network within the cloud, including for example having dedicated subnets and IP addresses.

[0078] In the examples illustrated in FIGS. 3A-3B, the customer networks' 302a-302b network is partially in a site network 304, and partially provided by the cloud services provider 354. In some cases, the site network 304 is the part of the customer networks 302a-302b that is located at a physical site owned or controlled by the customer network 302a-302b. For example, the site network 304 may be a network located in the customer network's 302a-302b office or campus. Alternatively or additionally, the site network 304 may include network equipment owned and/or operated by the customer network 302 that may be located anywhere. For example, the customer networks' 302a-302b operations may consist of a few laptops owned by the customer networks 302a-302b, which are used from the private homes of the lap tops' users, from a co-working space, from a coffee shop, or from some other mobile location.

[0079] In various implementations, sensors 310 may be installed in the site network 304. The sensors 310 can be used by the network security system to project deceptions into the site network 304, monitor the site network 304 for attacks, and/or to divert suspect attacks into the deception center 308.

[0080] In some implementations, the sensors 310 may also be able to project deceptions into the part of the customer networks 302a-302b network that is provided by the cloud services provider 354. In most cases, it may not be possible to install sensors 310 inside the network of the cloud services provider 354, but in some implementations, this may not be necessary. For example, as discussed further below, the deception center 308 can acquire the subnet address of the network provided by the cloud services provider 354, and use that subnet address to create deceptions. Though these deceptions are projected from the sensors 310 installed in the site network 304, the deceptions may appear to be within the subnet provided by the cloud services provider 354.

[0081] In illustrated examples, the deception center 308 is installed inside the customer networks 302a-302b. Though not illustrated here, the deception center 308 can also be installed outside the customer networks 302a-302b, such as for example somewhere on the Internet 350. In some implementations, the deception center 308 may reside at the same location as the security service provider 306. When located outside the customer networks 302a-302b, the deception center 308 may connect to the sensors 310 in the site network 304 over various public and/or private networks.

[0082] FIG. 3A illustrates an example of a configuration 300a where the customer network's 302a network infrastructure is located in the cloud and the customer network 302a also has a substantial site network 304. In this example, the customer may have an office where the site network 304 is located, and where the customer's employees access and use the customer network 302a. For example, developers, sales and marketing personnel, human resources and finance employees, may access the customer network 302a from the site network 304. In the illustrated example, the customer may obtain applications and services from the cloud services provider 354. Alternatively or additionally, the cloud services provider 354 may provide data center services for the customer. For example, the cloud services provider 354 may host the customer's repository of data (e.g., music provided by a streaming music service, or video provided by a streaming video provider). In this example,

the customer's own customers may be provided data directly from the cloud services provider 354, rather than from the customer network 302a.

[0083] FIG. 3B illustrates an example of a configuration 300b where the customer network's 302b network is primarily or sometimes entirely in the cloud. In this example, the customer network's 302b site network 304 may include a few laptops, or one or two desktop servers. These computing devices may be used by the customer's employees to conduct the customer's business, while the cloud services provider 354 provides the majority of the network infrastructure needed by the customer. For example, a very small company may have no office space and no dedicated location, and have as computing resources only the laptops used by its employees. This small company may use the cloud services provider 354 to provide its fixed network infrastructure. The small company may access this network infrastructure by connecting a laptop to any available network connection (e.g., in a co-working space, library, or coffee shop). When no laptops are connected to the cloud services provider 354, the customer network 302 may be existing entirely within the cloud.

[0084] In the example provided above, the site network 304 can be found wherever the customer's employees connect to a network and can access the cloud services provider 354. Similarly, the sensors 310 can be co-located with the employees' laptops. For example, whenever an employee connects to a network, she can enable a sensor 310, which can then project deceptions into the network around her. Alternatively or additionally, sensors 310 can be installed in a fixed location (such as the home of an employee of the customer) from which they can access the cloud services provider 354 and project deceptions into the network provided by the cloud services provider 354.

[0085] The network security system, such as the deception-based system discussed above, can provide network security for a variety of customer networks, which may include a diverse array of devices. FIG. 4 illustrates an example of an enterprise network 400, which is one such network that can be defended by a network security system. The example enterprise network 400 illustrates examples of various network devices and network clients that may be included in an enterprise network. The enterprise network 400 may include more or fewer network devices and/or network clients, and/or may include network devices, additional networks including remote sites 452, and/or systems not illustrated here. Enterprise networks may include networks installed at a large site, such as a corporate office, a university campus, a hospital, a government office, or a similar entity. An enterprise network may include

multiple physical sites. Access to an enterprise networks is typically restricted, and may require authorized users to enter a password or otherwise authenticate before using the network. A network such as illustrated by the example enterprise network 400 may also be found at small sites, such as in a small business.

[0086] The enterprise network 400 may be connected to an external network 450. The external network 450 may be a public network, such as the Internet. A public network is a network that has been made accessible to any device that can connect to it. A public network may have unrestricted access, meaning that, for example, no password or other authentication is required to connect to it. The external network 450 may include third-party telecommunication lines, such as phone lines, broadcast coaxial cable, fiber optic cables, satellite communications, cellular communications, and the like. The external network 450 may include any number of intermediate network devices, such as switches, routers, gateways, servers, and/or controllers that are not directly part of the enterprise network 400 but that facilitate communication between the network 400 and other network-connected entities, such as a remote site 452.

[0087] Remote sites 452 are networks and/or individual computers that are generally located outside the enterprise network 400, and which may be connected to the enterprise network 400 through intermediate networks, but that function as if within the enterprise network 400 and connected directly to it. For example, an employee may connect to the enterprise network 400 while at home, using various secure protocols, and/or by connecting to a Virtual Private Network (VPN) provided by the enterprise network 400. While the employee's computer is connected, the employee's home is a remote site 452. Alternatively or additionally, the enterprise network's 400 owner may have a satellite office with a small internal network. This satellite office's network may have a fixed connection to the enterprise network 400 over various intermediate networks. This satellite office can also be considered a remote site.

[0088] The enterprise network 400 may be connected to the external network 450 using a gateway device 404. The gateway device 404 may include a firewall or similar system for preventing unauthorized access while allowing authorized access to the enterprise network 400. Examples of gateway devices include routers, modems (e.g. cable, fiber optic, dial-up, etc.), and the like.

[0089] The gateway device 404 may be connected to a switch 406a. The switch 406a provides connectivity between various devices in the enterprise network 400. In this example, the switch 406a connects together the gateway device 404, various servers 408, 412, 414, 416, 418, and another switch 406b. A switch typically has multiple ports, and functions to direct packets received on one port to another port. In some implementations, the gateway device 404 and the switch 406a may be combined into a single device.

[0090] Various servers may be connected to the switch 406a. For example, a print server 408 may be connected to the switch 406a. The print server 408 may provide network access to a number of printers 410. Client devices connected to the enterprise network 400 may be able to access one of the printers 410 through the printer server 408.

[0091] Other examples of servers connected to the switch 406a include a file server 412, database server 414, and email server 416. The file server 412 may provide storage for and access to data. This data may be accessible to client devices connected to the enterprise network 400. The database server 414 may store one or more databases, and provide services for accessing the databases. The email server 416 may host an email program or service, and may also store email for users on the enterprise network 400.

[0092] As yet another example, a server rack 418 may be connected to the switch 406. The server rack 418 may house one or more rack-mounted servers. The server rack 418 may have one connection to the switch 406a, or may have multiple connections to the switch 406a. The servers in the server rack 418 may have various purposes, including providing computing resources, file storage, database storage and access, and email, among others.

[0093] An additional switch 406b may also be connected to the first switch 406a. The additional switch 406b may be provided to expand the capacity of the network. A switch typically has a limited number of ports (e.g., 8, 16, 32, 64 or more ports). In most cases, however, a switch can direct traffic to and from another switch, so that by connecting the additional switch 406b to the first switch 406a, the number of available ports can be expanded.

[0094] In this example, a server 420 is connected to the additional switch 406b. The server 420 may manage network access for a number of network devices or client devices. For example, the server 420 may provide network authentication, arbitration, prioritization, load balancing, and other management services as needed to manage multiple network devices accessing the enterprise network 400. The server 420 may be connected to a hub 422. The

hub 422 may include multiple ports, each of which may provide a wired connection for a network or client device. A hub is typically a simpler device than a switch, and may be used when connecting a small number of network devices together. In some cases, a switch can be substituted for the hub 422. In this example, the hub 422 connects desktop computers 424 and laptop computers 426 to the enterprise network 400. In this example, each of the desktop computers 424 and laptop computers 426 are connected to the hub 422 using a physical cable.

[0095] In this example, the additional switch 406b is also connected to a wireless access point 428. The wireless access point 428 provides wireless access to the enterprise network 400 for wireless-enabled network or client devices. Examples of wireless-enabled network and client devices include laptops 430, tablet computers 432, and smart phones 434, among others. In some implementations, the wireless access point 428 may also provide switching and/or routing functionality.

[0096] The example enterprise network 400 of FIG. 4 is defended from network threats by a network threat detection and analysis system, which uses deception security mechanisms to attract and divert attacks on the network. The deceptive security mechanisms may be controlled by and inserted into the enterprise network 400 using a deception center 498 and sensors 490, which may also be referred to as deception sensors, installed in various places in the enterprise network 400. In some implementations, the deception center 498 and the sensors 490 interact with a security services provider 496 located outside of the enterprise network 400. The deception center 498 may also obtain or exchange data with sources located on external networks 450, such as the Internet.

[0097] In various implementations, the sensors 490 are a minimal combination of hardware and/or software, sufficient to form a network connection with the enterprise network 400 and a network tunnel 480 with the deception center 498. For example, a sensor 490 may be constructed using a low-power processor, a network interface, and a simple operating system. In some implementations, any of the devices in the enterprise network (e.g., the servers 408, 412, 416, 418 the printers 410, the computing devices 424, 426, 430, 432, 434, or the network infrastructure devices 404, 406a, 406b, 428) can be configured to act as a sensor.

[0098] In various implementations, one or more sensors 490 can be installed anywhere in the enterprise network 400, include being attached switches 406a, hubs 422, wireless access points 428, and so on. The sensors 490 can further be configured to be part of one or more VLANs. The sensors 490 provide the deception center 498 with visibility into the enterprise

network 400, such as for example being able to operate as a node in the enterprise network 400, and/or being able to present or project deceptive security mechanisms into the enterprise network 400. Additionally, in various implementations, the sensors 490 may provide a portal through which a suspected attack on the enterprise network 400 can be redirected to the deception center 498.

[0099] The deception center 498 provides network security for the enterprise network 400 by deploying security mechanisms into the enterprise network 400, monitoring the enterprise network 400 through the security mechanisms, detecting and redirecting apparent threats, and analyzing network activity resulting from the apparent threat. To provide security for the enterprise network 400, in various implementations the deception center 498 may communicate with sensors 490 installed in the enterprise network 400, using, for example, network tunnels 480. The tunnels 480 may allow the deception center 498 to be located in a different sub-network (“subnet”) than the enterprise network 400, on a different network, or remote from the enterprise network 400, with intermediate networks between the deception center 498 and the enterprise network 400. In some implementations, the enterprise network 400 can include more than one deception center 498. In some implementations, the deception center may be located off-site, such as in an external network 450.

[0100] In some implementations, the security services provider 496 may act as a central hub for providing security to multiple site networks, possibly including site networks controlled by different organizations. For example, the security services provider 496 may communicate with multiple deception centers 498 that each provide security for a different enterprise network 400 for the same organization. As another example, the security services provider 496 may coordinate the activities of the deception center 498 and the sensors 490, such as enabling the deception center 498 and the sensors 490 to connect to each other. In some implementations, the security services provider 496 is located outside the enterprise network 400. In some implementations, the security services provider 496 is controlled by a different entity than the entity that controls the site network. For example, the security services provider 496 may be an outside vendor. In some implementations, the security services provider 496 is controlled by the same entity as that controls the enterprise network 400. In some implementations, the network security system does not include a security services provider 496.

[0101] FIG. 4 illustrates one example of what can be considered a “traditional” network, that is, a network that is based on the interconnection of computers. In various implementations, a network security system, such as the deception-based system discussed above, can also be used to defend “non-traditional” networks that include devices other than traditional computers, such as for example mechanical, electrical, or electromechanical devices, sensors, actuators, and control systems. Such “non-traditional” networks may be referred to as the Internet of Things (IoT). The Internet of Things encompasses newly-developed, every-day devices designed to be networked (e.g., drones, self-driving automobiles, etc.) as well as common and long-established machinery that has augmented to be connected to a network (e.g., home appliances, traffic signals, etc.).

[0102] FIG. 5 illustrates a general example of an IoT network 500. The example IoT network 500 can be implemented wherever sensors, actuators, and control systems can be found. For example, the example IoT network 500 can be implemented for buildings, roads and bridges, agriculture, transportation and logistics, utilities, air traffic control, factories, and private homes, among others. In various implementations, the IoT network 500 includes cloud service 554 that collects data from various sensors 510a-510d, 512a-512d, located in various locations. Using the collected data, the cloud service 554 can provide services 520, control of machinery and equipment 514, exchange of data with traditional network devices 516, and/or exchange of data with user devices 518. In some implementations, the cloud service 554 can work with a deception center 528 and/or a security service provider 526 to provide security for the network 500.

[0103] A cloud service, such as the illustrated cloud service 554, is a resource provided over the Internet 550. Sometimes synonymous with “cloud computing,” the resource provided by the cloud services is in the “cloud” in that the resource is provided by hardware and/or software at some location remote from the place where the resource is used. Often, the hardware and software of the cloud service is distributed across multiple physical locations. Generally, the resource provided by the cloud service is not directly associated with specific hardware or software resources, such that use of the resource can continue when the hardware or software is changed. The resource provided by the cloud service can often also be shared between multiple users of the cloud service, without affecting each user’s use. The resource can often also be provided as needed or on-demand. Often, the resource provided by the cloud service 554 is automated, or otherwise capable of operating with little or no assistance from human operators.

[0104] Examples of cloud services include software as a service (SaaS), infrastructure as a service (IaaS), platform as a service (PaaS), desktop as a service (DaaS), managed software as a service (MSaaS), mobile backend as a service (MBaaS), and information technology management as a service (ITMaas). Specific examples of cloud services include data centers, such as those operated by Amazon Web Services and Google Web Services, among others, that provide general networking and software services. Other examples of cloud services include those associated with smartphone applications, or “apps,” such as for example apps that track fitness and health, apps that allow a user to remotely manage her home security system or thermostat, and networked gaming apps, among others. In each of these examples, the company that provides the app may also provide cloud-based storage of application data, cloud-based software and computing resources, and/or networking services. In some cases, the company manages the cloud services provided by the company, including managing physical hardware resources. In other cases, the company leases networking time from a data center provider.

[0105] In some cases, the cloud service 554 is part of one integrated system, run by one entity. For example, the cloud service 554 can be part of a traffic control system. In this example, sensors 510a-510d, 512a-512d can be used to monitor traffic and road conditions. In this example, the cloud service 554 can attempt to optimize the flow of traffic and also provide traffic safety. For example, the sensors 510a-510d, 512a-512d can include a sensor 512a on a bridge that monitors ice formation. When the sensor 512a detects that ice has formed on the bridge, the sensor 512a can alert the cloud service 554. The cloud service 554, can respond by interacting with machinery and equipment 514 that manages traffic in the area of the bridge. For example, the cloud service 554 can turn on warning signs, indicating to drivers that the bridge is icy. Generally, the interaction between the sensor 512, the cloud service 554, and the machinery and equipment 514 is automated, requiring little or no management by human operators.

[0106] In various implementations, the cloud service 554 collects or receives data from sensors 510a-510d, 512a-512d, distributed across one or more networks. The sensors 510a-510d, 512a-512d include devices capable of “sensing” information, such as air or water temperature, air pressure, weight, motion, humidity, fluid levels, noise levels, and so on. The sensors 510a-510d, 512a-512d can alternatively or additionally include devices capable of receiving input, such as cameras, microphones, touch pads, keyboards, key pads, and so on. In some cases, a group of sensors 510a-510d may be common to one customer network 502.

For example, the sensors 510a-510d may be motion sensors, traffic cameras, temperature sensors, and other sensors for monitoring traffic in a city's metro area. In this example, the sensors 510a-510d can be located in one area of the city, or be distributed across the city, and be connected to a common network. In these cases, the sensors 510a-510d can communicate with a gateway device 562, such as a network gateway. The gateway device 562 can further communicate with the cloud service 554.

[0107] In some cases, in addition to receiving data from sensors 510a-510d in one customer network 502, the cloud service 554 can also receive data from sensors 512a-512d in other sites 504a-504c. These other sites 504a-504c can be part of the same customer network 502 or can be unrelated to the customer network 502. For example, the other sites 504a-504c can each be the metro area of a different city, and the sensors 512a-512d can be monitoring traffic for each individual city.

[0108] Generally, communication between the cloud service 554 and the sensors 510a-510d, 512a-512d is bidirectional. For example, the sensors 510a-510d, 512a-512d can send information to the cloud service 554. The cloud service 554 can further provide configuration and control information to the sensors 510a-510d, 512a-512d. For example, the cloud service 554 can enable or disable a sensor 510a-510d, 512a-512d or modify the operation of a sensor 510a-510d, 512a-512d, such as changing the format of the data provided by a sensor 510a-510d, 512a-512d or upgrading the firmware of a sensor 510a-510d, 512a-512d.

[0109] In various implementations, the cloud service 554 can operate on the data received from the sensors 510a-510d, 512a-512d, and use this data to interact with services 520 provided by the cloud service 554, or to interact with machinery and equipment 514, network devices 516, and/or user devices 518 available to the cloud service 554. Services 520 can include software-based services, such as cloud-based applications, website services, or data management services. Services 520 can alternatively or additionally include media, such as streaming video or music or other entertainment services. Services 520 can also include delivery and/or coordination of physical assets, such as for example package delivery, direction of vehicles for passenger pick-up and drop-off, or automate re-ordering and restocking of supplies. In various implementations, services 520 may be delivered to and used by the machinery and equipment 514, the network devices 516, and/or the user devices 518.

[0110] In various implementations, the machinery and equipment 514 can include physical systems that can be controlled by the cloud service 554. Examples of machinery and

equipment 514 include factory equipment, trains, electrical street cars, self-driving cars, traffic lights, gate and door locks, and so on. In various implementations, the cloud service 554 can provide configuration and control of the machinery and equipment 514 in an automated fashion.

[0111] The network devices 516 can include traditional networking equipment, such as server computers, data storage devices, routers, switches, gateways, and so on. In various implementations, the cloud service 554 can provide control and management of the network devices 516, such as for example automated upgrading of software, security monitoring, or asset tracking. Alternatively or additionally, in various implementations the cloud service 554 can exchange data with the network devices 516, such as for example providing websites, providing stock trading data, or providing online shopping resources, among others. Alternatively or additionally, the network devices 516 can include computing systems used by the cloud service provider to manage the cloud service 554.

[0112] The user devices 518 can include individual personal computers, smart phones, tablet devices, smart watches, fitness trackers, medical devices, and so on that can be associated with an individual user. The cloud service 554 can exchange data with the user devices 518, such as for example provide support for applications installed on the user devices 518, providing websites, providing streaming media, providing directional navigation services, and so on. Alternatively or additionally, the cloud service 554 may enable a user to use a user device 518 to access and/or view other devices, such as the sensors 510a-510d, 512a-512d, the machinery and equipment 514, or the network devices 516.

[0113] In various implementations, the services 520, machinery and equipment 514, network devices 516, and user devices 518 may be part of one customer network 506. In some cases, this customer network 506 is the same as the customer network 502 that includes the sensors 510a-510d. In some cases, the services 520, machinery and equipment 514, network devices 516, and user devices 518 are part of the same network, and may instead be part of various other networks 506.

[0114] In various implementations, customer networks can include a deception center 598. The deception center 598 provides network security for the IoT network 500 by deploying security mechanisms into the IoT network 500, monitoring the IoT network 500 through the security mechanisms, detecting and redirecting apparent threats, and analyzing network activity resulting from the apparent threat. To provide security for the IoT network 500, in

various implementations the deception center 598 may communicate with the sensors 510a-5106d, 512a-5012 installed in the IoT network 500, for example through the cloud service 554. In some implementations, the IoT network 500 can include more than one deception center 598. For example, each of customer network 502 and customer networks or other networks 506 can include a deception center 528.

[0115] In some implementations, the deception center 598 and the sensors 510a-510d, 512a-512d interact with a security services provider 596. In some implementations, the security services provider 596 may act as a central hub for providing security to multiple site networks, possibly including site networks controlled by different organizations. For example, the security services provider 596 may communicate with multiple deception centers 598 that each provide security for a different IoT network 500 for the same organization. As another example, the security services provider 596 may coordinate the activities of the deception center 598 and the sensors 510a-510d, 512a-512d, such as enabling the deception center 598 and the sensors 510a-510d, 512a-512d to connect to each other. In some implementations, the security services provider 596 is integrated into the cloud service 554. In some implementations, the security services provider 596 is controlled by a different entity than the entity that controls the site network. For example, the security services provider 596 may be an outside vendor. In some implementations, the security services provider 596 is controlled by the same entity as that controls the IoT network 500. In some implementations, the network security system does not include a security services provider 596.

[0116] IoT networks can also include small networks of non-traditional devices. FIG. 6 illustrates an example of a customer network that is a small network 600, here implemented in a private home. A network for a home is an example of small network that may have both traditional and non-traditional network devices connected to the network 600, in keeping with an Internet of Things approach. Home networks are also an example of networks that are often implemented with minimal security. The average homeowner is not likely to be a sophisticated network security expert, and may rely on his modem or router to provide at least some basic security. The homeowner, however, is likely able to at least set up a basic home network. A deception-based network security device may be as simple to set up as a home router or base station, yet provide sophisticated security for the network 600.

[0117] The example network 600 of FIG. 6 may be a single network, or may include multiple sub-networks. These sub-networks may or may not communicate with each other. For example, the network 600 may include a sub-network that uses the electrical wiring in the house as a communication channel. Devices configured to communicate in this way may connect to the network using electrical outlets, which also provide the devices with power. The sub-network may include a central controller device, which may coordinate the activities of devices connected to the electrical network, including turning devices on and off at particular times. One example of a protocol that uses the electrical wiring as a communication network is X10.

[0118] The network 600 may also include wireless and wired networks, built into the home or added to the home solely for providing a communication medium for devices in the house. Examples of wireless, radio-based networks include networks using protocols such as Z-Wave™, Zigbee™ (also known as Institute of Electrical and Electronics Engineers (IEEE) 802.15.4), Bluetooth™, and Wi-Fi (also known as IEEE 802.11), among others. Wireless networks can be set up by installing a wireless base station in the house. Alternatively or additionally, a wireless network can be established by having at least two devices in the house that are able to communicate with each other using the same protocol.

[0119] Examples of wired networks include Ethernet (also known as IEEE 802.3), token ring (also known as IEEE 802.5), Fiber Distributed Data Interface (FDDI), and Attached Resource Computer Network (ARCNET), among others. A wired network can be added to the house by running cabling through the walls, ceilings, and/or floors, and placing jacks in various rooms that devices can connect to with additional cables. The wired network can be extended using routers, switches, and/or hubs. In many cases, wired networks may be interconnected with wireless networks, with the interconnected networks operating as one seamless network. For example, an Ethernet network may include a wireless base station that provides a Wi-Fi signal for devices in the house.

[0120] As noted above, a small network 600 implemented in a home is one that may include both traditional network devices and non-traditional, everyday electronics and appliances that have also been connected to the network 600. Examples of rooms where one may find non-traditional devices connected to the network are the kitchen and laundry rooms. For example, in the kitchen a refrigerator 604, oven 606, microwave 608, and dishwasher 610 may be connected to the network 600, and in the laundry room a washing machine 612 may

be connected to the network 600. By attaching these appliances to the network 600, the homeowner can monitor the activity of each device (e.g., whether the dishes are clean, the current state of a turkey in the oven, or the washing machine cycle) or change the operation of each device without needing to be in the same room or even be at home. The appliances can also be configured to resupply themselves. For example, the refrigerator 604 may detect that a certain product is running low, and may place an order with a grocery delivery service for the product to be restocked.

[0121] The network 600 may also include environmental appliances, such as a thermostat 602 and a water heater 614. By having these devices connected to the network 600, the homeowner can monitor the current environment of the house (e.g., the air temperature or the hot water temperature), and adjust the settings of these appliances while at home or away. Furthermore, software on the network 600 or on the Internet 650 may track energy usage for the heating and cooling units and the water heater 614. This software may also track energy usage for the other devices, such as the kitchen and laundry room appliances. The energy usage of each appliance may be available to the homeowner over the network 600.

[0122] In the living room, various home electronics may be on the network 600. These electronics may have once been fully analog or may have been standalone devices, but now include a network connection for exchanging data with other devices in the network 600 or with the Internet 650. The home electronics in this example include a television 618, a gaming system 620, and a media device 622 (e.g., a video and/or audio player). Each of these devices may play media hosted, for example, on network attached storage 636 located elsewhere in the network 600, or media hosted on the Internet 650.

[0123] The network 600 may also include home safety and security devices, such as a smoke detector 616, an electronic door lock 624, and a home security system 626. Having these devices on the network may allow the homeowner to track the information monitored and/or sensed by these devices, both when the homeowner is at home and away from the house. For example, the homeowner may be able to view a video feed from a security camera 628. When the safety and security devices detect a problem, they may also inform the homeowner. For example, the smoke detector 616 may send an alert to the homeowner's smartphone when it detects smoke, or the electronic door lock 624 may alert the homeowner when there has been a forced entry. Furthermore, the homeowner may be able to remotely control these devices. For example, the homeowner may be able to remotely open the

electronic door lock 624 for a family member who has been locked out. The safety and security devices may also use their connection to the network to call the fire department or police if necessary.

[0124] Another non-traditional device that may be found in the network 600 is the family car 630. The car 630 is one of many devices, such as laptop computers 638, tablet computers 646, and smartphones 642, that connect to the network 600 when at home, and when not at home, may be able to connect to the network 600 over the Internet 650. Connecting to the network 600 over the Internet 650 may provide the homeowner with remote access to his network. The network 600 may be able to provide information to the car 630 and receive information from the car 630 while the car is away. For example, the network 600 may be able to track the location of the car 630 while the car 630 is away.

[0125] In the home office and elsewhere around the house, this example network 600 includes some traditional devices connected to the network 600. For example, the home office may include a desktop computer 632 and network attached storage 636. Elsewhere around the house, this example includes a laptop computer 638 and handheld devices such as a tablet computer 646 and a smartphone 642. In this example, a person 640 is also connected to the network 600. The person 640 may be connected to the network 600 wirelessly through personal devices worn by the person 640, such as a smart watch, fitness tracker, or heart rate monitor. The person 640 may alternatively or additionally be connected to the network 600 through a network-enabled medical device, such as a pacemaker, heart monitor, or drug delivery system, which may be worn or implanted.

[0126] The desktop computer 632, laptop computer 638, tablet computer 646, and/or smartphone 642 may provide an interface that allows the homeowner to monitor and control the various devices connected to the network. Some of these devices, such as the laptop computer 638, the tablet computer 646, and the smartphone 642 may also leave the house, and provide remote access to the network 600 over the Internet 650. In many cases, however, each device on the network may have its own software for monitoring and controlling only that one device. For example, the thermostat 602 may use one application while the media device 622 uses another, and the wireless network provides yet another. Furthermore, it may be the case that the various sub-networks in the house do not communicate with each other, and/or are viewed and controlled using software that is unique to each sub-network. In many

cases, the homeowner may not have one unified and easily understood view of his entire home network 600.

[0127] The small network 600 in this example may also include network infrastructure devices, such as a router or switch (not shown) and a wireless base station 634. The wireless base station 634 may provide a wireless network for the house. The router or switch may provide a wired network for the house. The wireless base station 634 may be connected to the router or switch to provide a wireless network that is an extension of the wired network. The router or switch may be connected to a gateway device 648 that connects the network 600 to other networks, including the Internet 650. In some cases, a router or switch may be integrated into the gateway device 648. The gateway device 648 is a cable modem, digital subscriber line (DSL) modem, optical modem, analog modem, or some other device that connects the network 600 to an ISP. The ISP may provide access to the Internet 650. Typically, a home network only has one gateway device 648. In some cases, the network 600 may not be connected to any networks outside of the house. In these cases, information about the network 600 and control of devices in the network 600 may not be available when the homeowner is not connected to the network 600; that is, the homeowner may not have access to his network 600 over the Internet 650.

[0128] Typically, the gateway device 648 includes a hardware and/or software firewall. A firewall monitors incoming and outgoing network traffic and, by applying security rules to the network traffic, attempts to keep harmful network traffic out of the network 600. In many cases, a firewall is the only security system protecting the network 600. While a firewall may work for some types of intrusion attempts originating outside the network 600, the firewall may not block all intrusion mechanisms, particularly intrusions mechanisms hidden in legitimate network traffic. Furthermore, while a firewall may block intrusions originating on the Internet 650, the firewall may not detect intrusions originating from within the network 600. For example, an infiltrator may get into the network 600 by connecting to signal from the Wi-Fi base station 634. Alternatively, the infiltrator may connect to the network 600 by physically connecting, for example, to the washing machine 612. The washing machine 612 may have a port that a service technician can connect to service the machine. Alternatively or additionally, the washing machine 612 may have a simple Universal Serial Bus (USB) port. Once an intruder has gained access to the washing machine 612, the intruder may have access to the rest of the network 600.

[0129] To provide more security for the network 600, a deception-based network security device 660 can be added to the network 600. In some implementations, the security device 660 is a standalone device that can be added to the network 600 by connecting it to a router or switch. In some implementations, the security device 660 can alternatively or additionally be connected to the network's 600 wireless sub-network by powering on the security device 660 and providing it with Wi-Fi credentials. The security device 660 may have a touchscreen, or a screen and a keypad, for inputting Wi-Fi credentials. Alternatively or additionally, the homeowner may be able to enter network information into the security device by logging into the security device 660 over a Bluetooth™ or Wi-Fi signal using software on a smartphone, tablet, or laptop, or using a web browser. In some implementations, the security device 660 can be connected to a sub-network running over the home's electrical wiring by connecting the security device 660 to a power outlet. In some implementations, the security device 660 may have ports, interfaces, and/or radio antennas for connecting to the various sub-networks that can be included in the network 600. This may be useful, for example, when the sub-networks do not communicate with each other, or do not communicate with each other seamlessly. Once powered on and connected, the security device 660 may self-configure and monitor the security of each sub-network in the network 600 that it is connected to.

[0130] In some implementations, the security device 660 may be configured to connect between the gateway device 648 and the network's 600 primary router, and/or between the gateway device 648 and the gateway device's 648 connection to the wall. Connected in one or both of these locations, the security device 660 may be able to control the network's 600 connection with outside networks. For example, the security device can disconnect the network 600 from the Internet 650.

[0131] In some implementations, the security device 660, instead of being implemented as a standalone device, may be integrated into one or more of the appliances, home electronics, or computing devices (in this example network 600), or in some other device not illustrated here. For example, the security device 660 – or the functionality of the security device 660 – may be incorporated into the gateway device 648 or a desktop computer 632 or a laptop computer 638. As another example, the security device 660 can be integrated into a kitchen appliance (e.g., the refrigerator 604 or microwave 608), a home media device (e.g., the television 618 or gaming system 620), or the home security system 626. In some implementations, the security device 660 may be a printed circuit board that can be added to another device without requiring significant changes to the other device. In some

implementations, the security device 660 may be implemented using an Application Specific Integrated Circuit (ASIC) or Field Programmable Gate Array (FPGA) that can be added to the electronics of a device. In some implementations, the security device 660 may be implemented as a software module or modules that can run concurrently with the operating system or firmware of a networked device. In some implementations, the security device 660 may have a physical or virtual security barrier that prevents access to it by the device that it is integrated into. In some implementations, the security device's 660 presence in another device may be hidden from the device into which the security device 660 is integrated.

[0132] In various implementations, the security device 660 may scan the network 600 to determine which devices are present in the network 600. Alternatively or additionally, the security device 660 may communicate with a central controller in the network 600 (or multiple central controllers, when there are sub-networks, each with their own central controller) to learn which devices are connected to the network 600. In some implementations, the security device 660 may undergo a learning period, during which the security device 660 learns the normal activity of the network 600, such as what time of day appliances and electronics are used, what they are used for, and/or what data is transferred to and from these devices. During the learning period, the security device 660 may alert the homeowner to any unusual or suspicious activity. The homeowner may indicate that this activity is acceptable, or may indicate that the activity is an intrusion. As described below, the security device 660 may subsequently take preventive action against the intrusion.

[0133] Once the security device 660 has learned the topology and/or activity of the network 600, the security device 660 may be able to provide deception-based security for the network 600. In some implementations, the security device 660 may deploy security mechanisms that are configured to emulate devices that could be found in the network 600. In some implementations, the security device 660 may monitor activity on the network 600, including watching the data sent between the various devices on the network 600, and between the devices and the Internet 650. The security device 660 may be looking for activity that is unusual, unexpected, or readily identifiable as suspect. Upon detecting suspicious activity in the network 600, the security device 660 may deploy deceptive security mechanisms.

[0134] In some implementations, the deceptive security mechanisms are software processes running on the security device 660 that emulate devices that may be found in the network 600. In some implementations, the security device 660 may be assisted in emulating the

security devices by another device on the network 600, such as the desktop computer 632. From the perspective of devices connected to the network 600, the security mechanisms appear just like any other device on the network, including, for example, having an Internet Protocol (IP) address, a Media Access Control (MAC) address, and/or some other identification information, having an identifiable device type, and responding to or transmitting data just as would the device being emulated. The security mechanisms may be emulated by the security device 660 itself; thus, while, from the point of view of the network 600, the network 600 appears to have additional devices, no physical equivalent (other than the security device 660) can be found in the house.

[0135] The devices and data emulated by a security mechanism are selected such that the security mechanism is an attractive target for intrusion attempts. Thus, the security mechanism may emulate valuable data, and/or devices that are easily hacked into, and/or devices that provide easy access to the reset of the network 600. Furthermore, the security mechanisms emulate devices that are likely to be found in the network 600, such as a second television, a second thermostat, or another laptop computer. In some implementations, the security device 660 may contact a service on the Internet 650 for assistance in selecting devices to emulate and/or for how to configure emulated devices. The security devices 660 may select and configure security mechanisms to be attractive to intrusions attempts, and to deflect attention away from more valuable or vulnerable network assets. Additionally, the security mechanisms can assist in confirming that an intrusion into the network 600 has actually taken place.

[0136] In some implementations, the security device 660 may deploy deceptive security mechanisms in advance of detecting any suspicious activity. For example, having scanned the network, the security device 660 may determine that the network 600 includes only one television 618 and one smoke detector 616. The security device 660 may therefore choose to deploy security mechanisms that emulate a second television and a second smoke detector. With security mechanisms preemptively added to the network, when there is an intrusion attempt, the intruder may target the security mechanisms instead of valuable or vulnerable network devices. The security mechanisms thus may serve as decoys and may deflect an intruder away from the network's 600 real devices.

[0137] In some implementations, the security mechanisms deployed by the security device 660 may take into account specific requirements of the network 600 and/or the type of

devices that can be emulated. For example, in some cases, the network 600 (or a sub-network) may assign identifiers to each device connected to the network 600, and/or each device may be required to adopt a unique identifier. In these cases, the security device 660 may assign an identifier to deployed security mechanisms that do not interfere with identifiers used by actual devices in the network 600. As another example, in some cases, devices on the network 600 may register themselves with a central controller and/or with a central service on the Internet 650. For example, the thermostat 602 may register with a service on the Internet 650 that monitors energy use for the home. In these cases, the security mechanisms that emulate these types of devices may also register with the central controller or the central service. Doing so may improve the apparent authenticity of the security mechanism, and may avoid conflicts with the central controller or central service. Alternatively or additionally, the security device 660 may determine to deploy security mechanisms that emulate other devices, and avoid registering with the central controller or central service.

[0138] In some implementations, the security device 660 may dynamically adjust the security mechanisms that it has deployed. For example, when the homeowner adds devices to the network 600, the security device 660 may remove security mechanisms that conflict with the new devices, or change a security mechanism so that the security mechanism's configuration is not incongruous with the new devices (e.g., the security mechanisms should not have the same MAC address as a new device). As another example, when the network owner removes a device from the network 600, the security device 660 may add a security mechanism that mimics the device that was removed. As another example, the security device may change the activity of a security mechanism, for example, to reflect changes in the normal activity of the home, changes in the weather, the time of year, the occurrence of special events, and so on.

[0139] The security device 660 may also dynamically adjust the security mechanisms it has deployed in response to suspicious activity it has detected on the network 600. For example, upon detecting suspicious activity, the security device 660 may change the behavior of a security mechanism or may deploy additional security mechanisms. The changes to the security mechanisms may be directed by the suspicious activity, meaning that if, for example, the suspicious activity appears to be probing for a wireless base station 634, the security device 660 may deploy a decoy wireless base station.

[0140] Changes to the security mechanisms are meant not only to attract a possible intrusion, but also to confirm that an intrusion has, in fact occurred. Since the security mechanisms are not part of the normal operation of the network 600, normal occupants of the home are not expected to access the security mechanisms. Thus, in most cases, any access of a security mechanism is suspect. Once the security device 660 has detected an access to a security mechanism, the security device 660 may next attempt to confirm that an intrusion into the network 600 has taken place. An intrusion can be confirmed, for example, by monitoring activity at the security mechanism. For example, login attempts, probing of data emulated by the security mechanism, copying of data from the security mechanism, and attempts to log into another part of the network 600 from the security mechanism indicate a high likelihood that an intrusion has occurred.

[0141] Once the security device 660 is able to confirm an intrusion into the network 600, the security device 660 may alert the homeowner. For example, the security device 660 may sound an audible alarm, send an email or text message to the homeowner or some other designated persons, and/or send an alert to an application running on a smartphone or tablet. As another example, the security device 660 may access other network devices and, for example, flash lights, trigger the security system's 626 alarm, and/or display messages on devices that include display screens, such as the television 618 or refrigerator 604. In some implementations, depending on the nature of the intrusion, the security device 660 may alert authorities such as the police or fire department.

[0142] In some implementations, the security device 660 may also take preventive actions. For example, when an intrusion appears to have originated outside the network 600, the security device 660 may block the network's 600 access to the Internet 650, thus possibly cutting off the intrusion.. As another example, when the intrusion appears to have originated from within the network 600, the security device 660 may isolate any apparently compromised devices, for example by disconnecting them from the network 600. When only its own security mechanisms are compromised, the security device 660 may isolate itself from the rest of the network 600. As another example, when the security device 660 is able to determine that the intrusion very likely included physical intrusion into the house, the security device 660 may alert the authorities. The security device 660 may further lock down the house by, for example, locking any electronic door locks 624.

[0143] In some implementations, the security device 660 may be able to enable a homeowner to monitor the network 600 when a suspicious activity has been detected, or at any other time. For example, the homeowner may be provided with a software application that can be installed on a smartphone, tablet, desktop, and/or laptop computer. The software application may receive information from the security device 660 over a wired or wireless connection. Alternatively or additionally, the homeowner may be able to access information about his network through a web browser, where the security device 660 formats webpages for displaying the information. Alternatively or additionally, the security device 660 may itself have a touchscreen or a screen and key pad that provide information about the network 600 to the homeowner.

[0144] The information provided to the homeowner may include, for example, a list and/or graphic display of the devices connected to the network 600. The information may further provide a real-time status of each device, such as whether the device is on or off, the current activity of the device, data being transferred to or from the device, and/or the current user of the device, among other things. The list or graphic display may update as devices connect and disconnect from the network 600, such as for example laptops and smartphones connecting to or disconnecting from a wireless sub-network in the network 600. The security device 660 may further alert the homeowner when a device has unexpectedly been disconnected from the network 600. The security device 660 may further alert the homeowner when an unknown device connects to the network 600, such as for example when a device that is not known to the homeowner connects to the Wi-Fi signal.

[0145] The security device 660 may also maintain historic information. For example, the security device 660 may provide snapshots of the network 600 taken once a day, once a week, or once a month. The security device 660 may further provide a list of devices that have, for example, connected to the wireless signal in the last hour or day, at what times, and for how long. The security device 660 may also be able to provide identification information for these devices, such as MAC addresses or usernames. As another example, the security device 660 may also maintain usage statistics for each device in the network 600, such as for example the times at which each device was in use, what the device was used for, how much energy the device used, and so on.

[0146] The software application or web browser or display interface that provides the homeowner with information about his network 600 may also enable the homeowner to make

changes to the network 600 or to devices in the network 600. For example, through the security device 660, the homeowner may be able to turn devices on or off, change the configuration of a device, change a password for a device or for the network, and so on.

[0147] In some implementations, the security device 660 may also display currently deployed security mechanisms and their configuration. In some implementations, the security device 660 may also display activity seen at the security mechanisms, such as for example a suspicious access to a security mechanism. In some implementations, the security device 660 may also allow the homeowner to customize the security mechanisms. For example, the homeowner may be able to add or remove security mechanisms, modify data emulated by the security mechanisms, modify the configuration of security mechanism, and/or modify the activity of a security mechanism.

[0148] A deception-based network security device 660 thus can provide sophisticated security for a small network. The security device 660 may be simple to add to a network, yet provide comprehensive protection against both external and internal intrusions. Moreover, the security device 660 may be able to monitor multiple sub-networks that are each using different protocols. The security device 660, using deceptive security mechanisms, may be able to detect and confirm intrusions into the network 600. The security device 660 may be able to take preventive actions when an intrusion occurs. The security device 660 may also be able to provide the homeowner with information about his network, and possibly also control over devices in the network.

[0149] FIG. 7 illustrates another example of a small network 700, here implemented in a small business. A network in a small business may have both traditional and non-traditional devices connected to the network 700. Small business networks are also examples of networks that are often implemented with minimal security. A small business owner may not have the financial or technical resources, time, or expertise to configure a sophisticated security infrastructure for her network 700. The business owner, however, is likely able to at least set up a network 700 for the operation of the business. A deception-based network security device that is at least as simple to set up as the network 700 itself may provide inexpensive and simple yet sophisticated security for the network 700.

[0150] The example network 700 may be one, single network, or may include multiple sub-networks. For example, the network 700 may include a wired sub-network, such as an Ethernet network, and a wireless sub-network, such as an 802.11 Wi-Fi network. The wired

sub-network may be implemented using cables that have been run through the walls and/or ceilings to the various rooms in the business. The cables may be connected to jacks in the walls that devices can connect to in order to connect to the network 700. The wireless network may be implemented using a wireless base station 720, or several wireless base stations, which provide a wireless signal throughout the business. The network 700 may include other wireless sub-networks, such as a short-distance Bluetooth™ network. In some cases, the sub-networks communicate with one another. For example, the Wi-Fi sub-network may be connected to the wired Ethernet sub-network. In some cases, the various sub-networks in the network 700 may not be configured to or able to communicate with each other.

[0151] As noted above, the small business network 700 may include both computers, network infrastructure devices, and other devices not traditionally found in a network. The network 700 may also include electronics, machinery, and systems that have been connected to the network 700 according to an Internet-of-Things approach. Workshop machinery that was once purely analog may now have computer controls. Digital workshop equipment may be network-enabled. By connecting shop equipment and machinery to the network 700, automation and efficiency of the business can be improved and orders, materials, and inventory can be tracked. Having more devices on the network 700, however, may increase the number of vulnerabilities in the network 700. Devices that have only recently become network-enabled may be particularly vulnerable because their security systems have not yet been hardened through use and attack. A deception-based network security device may provide simple-to-install and sophisticated security for a network that may otherwise have only minimal security.

[0152] The example small business of FIG. 7 includes a front office. In the front office, the network may include devices for administrative tasks. These devices may include, for example, a laptop computer 722 and a telephone 708. These devices may be attached to the network 700 in order to, for example, access records related to the business, which may be stored on a server 732 located elsewhere in the building. In the front office, security devices for the building may also be found, including, for example, security system controls 724 and an electronic door lock 726. Having the security devices on the network 700 may enable the business owner to remotely control access to the building. The business owner may also be able to remotely monitor the security of building, such as for example being able to view video streams from security cameras 742. The front office may also be where environmental

controls, such as a thermostat 702, are located. Having the thermostat 702 on the network 700 may allow the business owner to remotely control the temperature settings. A network-enabled thermostat 702 may also track energy usage for the heating and cooling systems. The front office may also include safety devices, such as a network-connected smoke alarm 728. A network-connected smoke alarm may be able to inform the business owner that there is a problem in the building by connecting to the business owner's smartphone or computer.

[0153] Another workspace in this example small business is a workshop. In the workshop, the network 700 may include production equipment for producing the goods sold by the business. The production equipment may include, for example, manufacturing machines 704 (e.g. a milling machine, a Computer Numerical Control (CNC) machine, a 3D printer, or some other machine tool) and a plotter 706. The production equipment may be controlled by a computer on the network 700, and/or may receive product designs over the network 700 and independently execute the designs. In the workshop, one may also find other devices related to the manufacturing of products, such as radiofrequency identification (RFID) scanners, barcode or Quick Response (QR) code generators, and other devices for tracking inventory, as well as electronic tools, hand tools, and so on.

[0154] In the workshop and elsewhere in the building, mobile computing devices and people 738 may also be connected to the network 700. Mobile computing devices include, for example, tablet computers 734 and smartphones 736. These devices may be used to control production equipment, track supplies and inventory, receive and track orders, and/or for other operations of the business. People 738 may be connected to the network through network-connected devices worn or implanted in the people 738, such as for example smart watches, fitness trackers, heart rate monitors, drug delivery systems, pacemakers, and so on.

[0155] At a loading dock, the example small business may have a delivery van 748 and a company car 746. When these vehicles are away from the business, they may be connected to the network 700 remotely, for example over the Internet 750. By being able to communicate with the network 700, the vehicles may be able to receive information such as product delivery information (e.g., orders, addresses, and/or delivery times), supply pickup instructions, and so on. The business owner may also be able to track the location of these vehicles from the business location, or over the Internet 750 when away from the business, and/or track who is using the vehicles.

[0156] The business may also have a back office. In the back office, the network 700 may include traditional network devices, such as computers 730, a multi-function printer 716, a scanner 718, and a server 732. In this example, the computers 730 may be used to design products for manufacturing in the workshop, as well as for management of the business, including tracking orders, supplies, inventory, and/or human resources records. The multi-function printer 716 and scanner 718 may support the design work and the running of the business. The server 732 may store product designs, orders, supply records, and inventory records, as well as administrative data, such as accounting and human resources data.

[0157] The back office may also be where a gateway device 770 is located. The gateway device 770 connects the small business to other networks, including the Internet 750. Typically, the gateway device 770 connects to an ISP, and the ISP provides access to the Internet 750. In some cases, a router may be integrated into the gateway device 770. In some cases, gateway device 770 may be connected to an external router, switch, or hub, not illustrated here. In some cases, the network 700 is not connected to any networks outside of the business's own network 700. In these cases, the network 700 may not have a gateway device 770.

[0158] The back office is also where the network 700 may have a deception-based network security device 760. The security device 760 may be a standalone device that may be enabled as soon as it is connected to the network 700. Alternatively or additionally, the security device 760 may be integrated into another device connected to the network 700, such as the gateway device 770, a router, a desktop computer 730, a laptop computer 722, the multi-function printer 716, or the thermostat 702, among others. When integrated into another device, the security device 760 may use the network connection of the other device, or may have its own network connection for connecting to the network 700. The security device 760 may connect to the network 700 using a wired connection or a wireless connection.

[0159] Once connected to the network 700, the security device 760 may begin monitoring the network 700 for suspect activity. In some implementations, the security device 760 may scan the network 700 to learn which devices are connected to the network 700. In some cases, the security device 760 may learn the normal activity of the network 700, such as what time the various devices are used, for how long, by whom, for what purpose, and what data is transferred to and from each device, among other things.

[0160] In some implementations, having learned the configuration and/or activity of the network 700, the security device 760 may deploy deceptive security mechanisms. These security mechanisms may emulate devices that may be found on the network 700, including having an identifiable device type and/or network identifiers (such as a MAC address and/or IP address), and being able to send and receive network traffic that a device of a certain time would send and receive. For example, for the example small business, the security device 760 may configure a security mechanism to emulate a 3D printer, a wide-body scanner, or an additional security camera. The security device 760 may further avoid configuring a security mechanism to emulate a device that is not likely to be found in the small business, such as a washing machine. The security device 760 may use the deployed security mechanisms to monitor activity on the network 700.

[0161] In various implementations, when the security device 760 detects suspect activity, the security device 760 may deploy additional security mechanisms. These additional security mechanisms may be selected based on the nature of suspect activity. For example, when the suspect activity appears to be attempting to break into the shop equipment, the security device 760 may deploy a security mechanism that looks like shop equipment that is easy to hack. In some implementations, the security device 760 may deploy security mechanisms only after detecting suspect activity on the network 700.

[0162] The security device 760 selects devices to emulate that are particularly attractive for an infiltration, either because the emulated device appears to have valuable data or because the emulated device appears to be easy to infiltrate, or for some other reason. In some implementations, the security device 760 connects to a service on the Internet 750 for assistance in determining which devices to emulate and/or how to configure the emulated device. Once deployed, the security mechanisms serve as decoys to attract the attention of a possible infiltrator away from valuable network assets. In some implementations, the security device 760 emulates the security mechanisms using software processes. In some implementations, the security device 760 may be assisted in emulating security mechanisms by a computer 730 on the network.

[0163] In some implementations, the security device 760 may deploy security mechanisms prior to detecting suspicious activity on the network 700. In these implementations, the security mechanisms may present more attractive targets for a possible, future infiltration, so

that if an infiltration occurs, the infiltrator will go after the security mechanisms instead of the actual devices on the network 700.

[0164] In various implementations, the security device 760 may also change the security mechanisms that it has deployed. For example, the security device 760 may add or remove security mechanisms as the operation of the business changes, as the activity on the network 700 changes, as devices are added or removed from the network 700, as the time of year changes, and so on.

[0165] Besides deflecting a possible network infiltration away from valuable or vulnerable network devices, the security device 760 may use the security mechanisms to confirm that the network 700 has been infiltrated. Because the security mechanisms are not part of actual devices in use by the business, any access to them over the network is suspect. Thus, once the security device 760 detects an access to one of its security mechanisms, the security device 760 may attempt to confirm that this access is, in fact, an unauthorized infiltration of the network 700.

[0166] To confirm that a security mechanism has been infiltrated, the security device 760 may monitor activity seen at the security mechanism. The security device 760 may further deploy additional security mechanisms, to see if, for example, it can present an even more attractive target to the possible infiltrator. The security device 760 may further look for certain activity, such as log in attempts to other devices in the network, attempts to examine data on the security mechanism, attempts to move data from the security mechanism to the Internet 750, scanning of the network 700, password breaking attempts, and so on.

[0167] Once the security device 760 has confirmed that the network 700 has been infiltrated, the security device 760 may alert the business owner. For example, the security device 760 may sound an audible alarm, email or send text messages to the computers 730 and/or handheld devices 734, 736, send a message to the business's cars 746, 748, flash lights, or trigger the security system's 724 alarm. In some implementations, the security device 760 may also take preventive measures. For example, the security device 760 may disconnect the network 700 from the Internet 750, may disconnect specific devices from the network 700 (e.g., the server 732 or the manufacturing machines 704), may turn some network-connected devices off, and/or may lock the building.

[0168] In various implementations, the security device 760 may allow the business owner to monitor her network 700, either when an infiltration is taking place or at any other time.

For example, the security device 760 may provide a display of the devices currently connected to the network 700, including flagging any devices connected to the wireless network that do not appear to be part of the business. The security device 760 may further display what each device is currently doing, who is using them, how much energy each device is presently using, and/or how much network bandwidth each device is using. The security device 760 may also be able to store this information and provide historic configuration and/or usage of the network 700.

[0169] The security device 760 may have a display it can use to show information to the business owner. Alternatively or additionally, the security device 760 may provide this information to a software application that can run on a desktop or laptop computer, a tablet, or a smartphone. Alternatively or additionally, the security device 760 may format this information for display through a web browser. The business owner may further be able to control devices on the network 700 through an interface provided by the security device 760, including, for example, turning devices on or off, adjusting settings on devices, configuring user accounts, and so on. The business owner may also be able to view any security mechanisms presently deployed, and may be able to re-configure the security mechanisms, turn them off, or turn them on.

[0170] IoT networks can also include industrial control systems. Industrial control system is a general term that encompasses several types of control systems, including supervisory control and data acquisition (SCADA) systems, distributed control systems (DCS) and other control system configurations, such as Programmable Logic Controllers (PLCs), often found in the industrial sectors and infrastructures. Industrial control systems are often found in industries such as electrical, water and wastewater, oil and natural gas, chemical, transportation, pharmaceutical, pulp and paper, food and beverage, and discrete manufacturing (e.g., automotive, aerospace, and durable goods). While a large percentage of industrial control systems may be privately owned and operated, federal agencies also operate many industrial processes, such as air traffic control systems and materials handling (e.g., Postal Service mail handling).

[0171] FIG. 8 illustrates an example of the basic operation of an industrial control system 800. Generally, an industrial control system 800 may include a control loop 802, a human-machine interface 806, and remote diagnostics and maintenance 808. In some implementations, the example industrial control system can be defended by a network threat

detection and analysis system, which can include a deception center 898 and a security services provider 896.

[0172] A control loop 802 may consist of sensors 812, controller 804 hardware such as PLCs, actuators 810, and the communication of variables 832, 834. The sensors 812 may be used for measuring variables in the system, while the actuators 810 may include, for example, control valves breakers, switches, and motors. Some of the sensors 812 may be deception sensors. Controlled variables 834 may be transmitted to the controller 804 from the sensors 812. The controller 804 may interpret the controlled variables 834 and generates corresponding manipulated variables 832, based on set points provided by controller interaction 830. The controller 804 may then transmit the manipulated variables 832 to the actuators 810. The actuators 810 may drive a controlled process 814 (e.g., a machine on an assembly line). The controlled process 814 may accept process inputs 822 (e.g., raw materials) and produce process outputs 824 (e.g., finished products). New information 820 provided to the controlled process 814 may result in new sensor 812 signals, which identify the state of the controlled process 814 and which may also transmitted to the controller 804.

[0173] In some implementations, at least some of the sensors 812 can also provide the deception center 898 with visibility into the industrial control system 800, such as for example being able to present or project deceptive security mechanisms into the industrial control system. Additionally, in various implementations, the sensors 812 may provide a portal through which a suspected attack on the industrial control system can be redirected to the deception center 898. The deception center 898 and the sensors 812 may be able to communicate using network tunnels 880.

[0174] The deception center 898 provides network security for the industrial control system 800 by deploying security mechanisms into the industrial control system 800, monitoring the industrial control system through the security mechanisms, detecting and redirecting apparent threats, and analyzing network activity resulting from the apparent threat. In some implementations, the industrial control system 800 can include more than one deception center 898. In some implementations, the deception center may be located off-site, such as on the Internet.

[0175] In some implementations, the deception center 898 may interact with a security services provider 896 located outside the industrial control system 800. The security services provider 896 may act as a central hub for providing security to multiple sites that are part of

the industrial control system 800, and/or for multiple separate, possibly unrelated, industrial control systems. For example, the security services provider 896 may communicate with multiple deception centers 898 that each provide security for a different industrial control system 800 for the same organization. As another example, the security services provider 896 may coordinate the activities of the deception center 898 and the sensors 812, such as enabling the deception center 898 and the sensors 812 to connect to each other. In some implementations, the security services provider 896 is located outside the industrial control system 800. In some implementations, the security services provider 896 is controlled by a different entity than the entity that controls the site network. For example, the security services provider 896 may be an outside vendor. In some implementations, the security services provider 896 is controlled by the same entity as that controls the industrial control system. In some implementations, the network security system does not include a security services provider 896.

[0176] The human-machine interface 806 provides operators and engineers with an interface for controller interaction 830. Controller interaction 830 may include monitoring and configuring set points and control algorithms, and adjusting and establishing parameters in the controller 804. The human-machine interface 806 typically also receives information from the controller 804 that allows the human-machine interface 806 to display process status information and historical information about the operation of the control loop 802.

[0177] The remote diagnostics and maintenance 808 utilities are typically used to prevent, identify, and recover from abnormal operation or failures. For diagnostics, the remote diagnostics and maintenance utilities 808 may monitor the operation of each of the controller 804, sensors 812, and actuators 810. To recover after a problem, the remote diagnostics and maintenance 808 utilities may provide recovery information and instructions to one or more of the controller 804, sensors 812, and/or actuators 810.

[0178] A typical industrial control system contains many control loops, human-machine interfaces, and remote diagnostics and maintenance tools, built using an array of network protocols on layered network architectures. In some cases, multiple control loops are nested and/or cascading, with the set point for one control loop being based on process variables determined by another control loop. Supervisory-level control loops and lower-level control loops typically operate continuously over the duration of a process, with cycle times ranging from milliseconds to minutes.

[0179] One type of industrial control system that may include many control loops, human-machine interfaces, and remote diagnostics and maintenance tools is a supervisory control and data acquisition (SCADA) system. SCADA systems are used to control dispersed assets, where centralized data acquisition is typically as important as control of the system. SCADA systems are used in distribution systems such as, for example, water distribution and wastewater collection systems, oil and natural gas pipelines, electrical utility transmission and distribution systems, and rail and other public transportation systems, among others. SCADA systems typically integrate data acquisition systems with data transmission systems and human-machine interface software to provide a centralized monitoring and control system for numerous process inputs and outputs. SCADA systems are typically designed to collect field information, transfer this information to a central computer facility, and to display the information to an operator in a graphic and/or textual manner. Using this displayed information, the operator may, in real time, monitor and control an entire system from a central location. In various implementations, control of any individual sub-system, operation, or task can be automatic, or can be performed by manual commands.

[0180] FIG. 9 illustrates an example of a SCADA system 900, here used for distributed monitoring and control. This example SCADA system 900 includes a primary control center 902 and three field sites 930a-930c. A backup control center 904 provides redundancy in case of there is a malfunction at the primary control center 902. The primary control center 902 in this example includes a control server 906 – which may also be called a SCADA server or a Master Terminal Unit (MTU) – and a local area network (LAN) 908. The primary control center 902 may also include a human-machine interface station 908, a data historian 910, engineering workstations 912, and various network equipment such as printers 914, each connected to the LAN 918.

[0181] The control server 906 typically acts as the master of the SCADA system 900. The control server 906 typically includes supervisory control software that controls lower-level control devices, such as Remote Terminal Units (RTUs) and PLCs, located at the field sites 930a-930c. The software may tell the system 900 what and when to monitor, what parameter ranges are acceptable, and/or what response to initiate when parameters are outside of acceptable values.

[0182] The control server 906 of this example may access Remote Terminal Units and/or PLCs at the field sites 930a-930c using a communications infrastructure, which may include

radio-based communication devices, telephone lines, cables, and/or satellites. In the illustrated example, the control server 906 is connected to a modem 916, which provides communication with serial-based radio communication 920, such as a radio antenna. Using the radio communication 920, the control server 906 can communicate with field sites 930a-930b using radiofrequency signals 922. Some field sites 930a-930b may have radio transceivers for communicating back to the control server 906.

[0183] A human-machine interface station 908 is typically a combination of hardware and software that allows human operators to monitor the state of processes in the SCADA system 900. The human-machine interface station 908 may further allow operators to modify control settings to change a control objective, and/or manually override automatic control operations, such as in the event of an emergency. The human-machine interface station 908 may also allow a control engineer or operator to configure set points or control algorithms and parameters in a controller, such as a Remote Terminal Unit or a PLC. The human-machine interface station 908 may also display process status information, historical information, reports, and other information to operators, administrators, managers, business partners, and other authorized users. The location, platform, and interface of a human-machine interface station 908 may vary. For example, the human-machine interface station 908 may be a custom, dedicated platform in the primary control center 902, a laptop on a wireless LAN, or a browser on a system connected to the Internet.

[0184] The data historian 910 in this example is a database for logging all process information within the SCADA system 900. Information stored in this database can be accessed to support analysis of the system 900, for example for statistical process control or enterprise level planning.

[0185] The backup control center 904 may include all or most of the same components that are found in the primary control center 902. In some cases, the backup control center 904 may temporarily take over for components at the primary control center 902 that have failed or have been taken offline for maintenance. In some cases, the backup control center 904 is configured to take over all operations of the primary control center 902, such as when the primary control center 902 experiences a complete failure (e.g., is destroyed in a natural disaster).

[0186] The primary control center 902 may collect and log information gathered by the field sites 930a-930c and display this information using the human-machine interface station

908. The primary control center 902 may also generate actions based on detected events. The primary control center 902 may, for example, poll field devices at the field sites 930a-930c for data at defined intervals (e.g., 5 or 60 seconds), and can send new set points to a field device as required. In addition to polling and issuing high-level commands, the primary control center 902 may also watch for priority interrupts coming from the alarm systems at the field sites 930a-930c.

[0187] In this example, the primary control center 902 uses point-to-point connections to communication with three field sites 930a-930c, using radio telemetry for two communications with two of the field sites 930a-930b. In this example, the primary control center 902 uses a wide area network (WAN) 960 to communicate with the third field site 930c. In other implementations, the primary control center 902 may use other communication topologies to communicate with field sites. Other communication topologies include rings, stars, meshes, trees, lines or series, and busses or multi-drops, among others. Standard and proprietary communication protocols may be used to transport information between the primary control center 902 and field sites 930a-930c. These protocols may use telemetry techniques such as provided by telephone lines, cables, fiber optics, and/or radiofrequency transmissions such as broadcast, microwave, and/or satellite communications.

[0188] The field sites 930a-930c in this example perform local control of actuators and monitor local sensors. For example, a first field site 930a may include a PLC 932. A PLC is a small industrial computer originally designed to perform the logic functions formerly executed by electrical hardware (such as relays, switches, and/or mechanical timers and counters). PLCs have evolved into controllers capable of controlling complex processes, and are used extensively in both SCADA systems and distributed control systems. Other controllers used at the field level include process controllers and Remote Terminal Units, which may provide the same level of control as a PLC but may be designed for specific control applications. In SCADA environments, PLCs are often used as field devices because they are more economical, versatile, flexible, and configurable than special-purpose controllers.

[0189] The PLC 932 at a field site, such as the first field site 930a, may control local actuators 934, 936 and monitor local sensors 938, 940, 942. Examples of actuators include valves 934 and pumps 936, among others. Examples of sensors include level sensors 938, pressure sensors 940, and flow sensors 942, among others. Any of the actuators 934, 936 or

sensors 938, 940, 942 may be “smart” actuators or sensors, more commonly called intelligent electronic devices (IEDs). Intelligent electronic devices may include intelligence for acquiring data, communicating with other devices, and performing local processing and control. An intelligent electronic device could combine an analog input sensor, analog output, low-level control capabilities, a communication system, and/or program memory in one device. The use of intelligent electronic devices in SCADA systems and distributed control systems may allow for automatic control at the local level. Intelligent electronic devices, such as protective relays, may communicate directly with the control server 906. Alternatively or additionally, a local Remote Terminal Unit may poll intelligent electronic devices to collect data, which it may then pass to the control server 906.

[0190] Field sites 930a-930c are often equipped with remote access capability that allows field operators to perform remote diagnostics and repairs. For example, the first remote 930a may include a modem 916 connected to the PLC 932. A remote access 950 site may be able to, using a dial up connection, connect to the modem 916. The remote access 950 site may include its own modem 916 for dialing into to the field site 930a over a telephone line. At the remote access 950 site, an operator may use a computer 952 connected to the modem 916 to perform diagnostics and repairs on the first field site 930a.

[0191] The example SCADA system 900 includes a second field site 930b, which may be provisioned in substantially the same way as the first field site 930a, having at least a modem and a PLC or Remote Terminal that controls and monitors some number of actuators and sensors.

[0192] The example SCADA system 900 also includes a third field site 930c that includes a network interface card (NIC) 944 for communicating with the system’s 900 WAN 960. In this example, the third field site 930c includes a Remote Terminal Unit 946 that is responsible for controlling local actuators 934, 936 and monitoring local sensors 938, 940, 942. A Remote Terminal Unit, also called a remote telemetry unit, is a special-purpose data acquisition and control unit typically designed to support SCADA remote stations. Remote Terminal Units may be field devices equipped with wireless radio interfaces to support remote situations where wire-based communications are unavailable. In some cases, PLCs are implemented as Remote Terminal Units.

[0193] The SCADA system 900 of this example also includes a regional control center 970 and a corporate enterprise network 980. The regional control center 970 may provide a higher

level of supervisory control. The regional control center 970 may include at least a human-machine interface station 908 and a control server 906 that may have supervisory control over the control server 906 at the primary control center 902. The corporate enterprise network 980 typically has access, through the system's 900 WAN 960, to all the control centers 902, 904 and to the field sites 930a-930c. The corporate enterprise network 980 may include a human-machine interface station 908 so that operators can remotely maintain and troubleshoot operations.

[0194] Another type of industrial control system is the distributed control system (DCS). Distributed control systems are typically used to control production systems within the same geographic location for industries such as oil refineries, water and wastewater management, electric power generation plants, chemical manufacturing plants, and pharmaceutical processing facilities, among others. These systems are usually process control or discrete part control systems. Process control systems may be processes that run continuously, such as manufacturing processes for fuel or steam flow in a power plant, for petroleum production in a refinery, or for distillation in a chemical plant. Discrete part control systems have processes that have distinct processing steps, typically with a distinct start and end to each step, such as found in food manufacturing, electrical and mechanical parts assembly, and parts machining. Discrete-based manufacturing industries typically conduct a series of steps on a single item to create an end product.

[0195] A distributed control system typically uses a centralized supervisory control loop to mediate a group of localized controllers that share the overall tasks of carrying out an entire production process. By modularizing the production system, a distributed control system may reduce the impact of a single fault on the overall system. A distributed control system is typically interfaced with a corporate network to give business operations a view of the production process.

[0196] FIG. 10 illustrates an example of a distributed control system 1000. This example distributed control system 1000 encompasses a production facility, including bottom-level production processes at a field level 1004, supervisory control systems at a supervisory level 1002, and a corporate or enterprise layer.

[0197] At the supervisory level 1002, a control server 1006, operating as a supervisory controller, may communicate with subordinate systems via a control network 1018. The control server 1006 may send set points to distributed field controllers, and may request data

from the distributed field controllers. The supervisory level 1002 may include multiple control servers 1006, with one acting as the primary control server and the rest acting as redundant, back-up control servers. The supervisory level 1002 may also include a main human-machine interface 1008 for use by operators and engineers, a data historian 1010 for logging process information from the system 1000, and engineering workstations 1012.

[0198] At the field level 1004, the system 1000 may include various distributed field controllers. In the illustrated example, the distributed control system 1000 includes a machine controller 1020, a PLC 1032, a process controller 1040, and a single loop controller 1044. The distributed field controllers may each control local process actuators, based on control server 1006 commands and sensor feedback from local process sensors.

[0199] In this example, the machine controller 1020 drives a motion control network 1026. Using the motion control network 1026, the machine controller 1020 may control a number of servo drives 1022, which may each drive a motor. The machine controller 1020 may also drive a logic control bus 1028 to communicate with various devices 1024. For example, the machine controller 1020 may use the logic control bus 1028 to communicate with pressure sensors, pressure regulators, and/or solenoid valves, among other devices. One or more of the devices 1024 may be an intelligent electronic device. A human-machine interface 1008 may be attached to the machine controller 1020 to provide an operator with local status information about the processes under control of the machine controller 1020, and/or local control of the machine controller 1020. A modem 1016 may also be attached to the machine controller 1020 to provide remote access to the machine controller 1020.

[0200] The PLC 1032 in this example system 1000 uses a fieldbus 1030 to communicate with actuators 1034 and sensors 1036 under its control. These actuators 1034 and sensors 1036 may include, for example, direct current (DC) servo drives, alternating current (AC) servo drives, light towers, photo eyes, and/or proximity sensors, among others. A human-machine interface 1008 may also be attached to the fieldbus 1030 to provide operators with local status and control for the PLC 1032. A modem 1016 may also be attached to the PLC 1032 to provide remote access to the PLC 1032.

[0201] The process controller 1040 in this example system 1000 also uses a fieldbus 1030 to communicate with actuators and sensors under its control, one or more of which may be intelligent electronic devices. The process controller 1040 may communicate with its fieldbus 1030 through an input/output (I/O) server 1042. An I/O server is a control component

typically responsible for collecting, buffering, and/or providing access to process information from control sub-components. An I/O server may be used for interfacing with third-party control components. Actuators and sensors under control of the process controller 1040 may include, for example, pressure regulators, pressure sensors, temperature sensors, servo valves, and/or solenoid valves, among others. The process controller 1040 may be connected to a modem 1016 so that a remote access 1050 site may access the process controller 1040. The remote access 1050 site may include a computer 1052 for use by an operator to monitor and control the process controller 1040. The computer 1052 may be connected to a local modem 1016 for dialing in to the modem 1016 connected to the process controller 1040.

[0202] The illustrated example system 1000 also includes a single loop controller 1044. In this example, the single loop controller 1044 interfaces with actuators 1034 and sensors 1036 with point-to-point connections, instead of a fieldbus. Point-to-point connections require a dedicated connection for each actuator 1034 and each sensor 1036. Fieldbus networks, in contrast, do not need point-to-point connections between a controller and individual field sensors and actuators. In some implementations, a fieldbus allows greater functionality beyond control, including field device diagnostics. A fieldbus can accomplish control algorithms within the fieldbus, thereby avoiding signal routing back to a PLC for every control operation. Standard industrial communication protocols are often used on control networks and fieldbus networks.

[0203] The single loop controller 1044 in this example is also connected to a modem 1016, for remote access to the single loop controller.

[0204] In addition to the supervisory level 1002 and field level 1004 control loops, the distributed control system 1000 may also include intermediate levels of control. For example, in the case of a distributed control system controlling a discrete part manufacturing facility, there could be an intermediate level supervisor for each cell within the plant. This intermediate level supervisor could encompass a manufacturing cell containing a machine controller that processes a part, and a robot controller that handles raw stock and final products. Additionally, the distributed control system could include several of these cells that manage field-level controllers under the main distributed control system supervisory control loop.

[0205] In various implementations, the distributed control system may include a corporate or enterprise layer, where an enterprise network 1080 may connect to the example production

facility. The enterprise network 1080 may be, for example, located at a corporate office co-located with the facility, and connected to the control network 1018 in the supervisory level 1002. The enterprise network 1080 may provide engineers and managers with control and visibility into the facility. The enterprise network 1080 may further include Manufacturing Execution Systems (MES) 1092, control systems for managing and monitoring work-in-process on a factory floor. An MES can track manufacturing information in real time, receiving up-to-the-minute data from robots, machine monitors and employees. The enterprise network 1080 may also include Management Information Systems (MIS) 1094, software and hardware applications that implement, for example, decision support systems, resource and people management applications, project management, and database retrieval applications, as well as basic business functions such as order entry and accounting. The enterprise network 1080 may further include Enterprise Resource Planning (ERP) systems 1096, business process management software that allows an organization to use a system of integrated applications to manage the business and automate many back office functions related to technology, services, and human resources.

[0206] The enterprise network 1080 may further be connected to a WAN 1060. Through the WAN 1060, the enterprise network 1080 may connect to a distributed plant 1098, which may include control loops and supervisory functions similar to the illustrated facility, but which may be at a different geographic location. The WAN 1060 may also connect the enterprise network to the outside world 1090, that is, to the Internet and/or various private and public networks. In some cases, the WAN 1060 may itself include the Internet, so that the enterprise network 1080 accesses the distributed plant 1098 over the Internet.

[0207] As described above, SCADA systems and distributed control systems use Programmable Logic Controllers (PLCs) as the control components of an overall hierarchical system. PLCs can provide local management of processes through feedback control, as described above. In a SCADA implementation, a PLC can provide the same functionality as a Remote Terminal Unit. When used in a distributed control system, PLCs can be implemented as local controllers within a supervisory scheme. PLCs can have user-programmable memory for storing instructions, where the instructions implement specific functions such as I/O control, logic, timing, counting, proportional-integral-derivative (PID) control, communication, arithmetic, and data and file processing.

[0208] FIG. 11 illustrates an example of a PLC 1132 implemented in a manufacturing control process. The PLC 1132 in this example monitors and controls various devices over fieldbus network 1130. The PLC 1132 may be connected to a LAN 1118. An engineering workstation 1112 may also be connected to the LAN 1118, and may include a programming interface that provides access to the PLC 1132. A data historian 1110 on the LAN 1118 may store data produced by the PLC 1132.

[0209] The PLC 1132 in this example may control a number of devices attached to its fieldbus network 1130. These devices may include actuators, such as a DC servo drive 1122, an AC drive 1124, a variable frequency drive 1134, and/or a light tower 1138. The PLC 1132 may also monitor sensors connected to the fieldbus network 1130, such as proximity sensors 1136, and/or a photo eye 1142. A human-machine interface 1108 may also be connected to the fieldbus network 1130, and may provide local monitoring and control of the PLC 1132.

[0210] Most industrial control systems were developed years ago, long before public and private networks, desktop computing, or the Internet were a common part of business operations. These well-established industrial control systems were designed to meet performance, reliability, safety, and flexibility requirements. In most cases, they were physically isolated from outside networks and based on proprietary hardware, software, and communication protocols that included basic error detection and correction capabilities, but lacked secure communication capabilities. While there was concern for reliability, maintainability, and availability when addressing statistical performance and failure, the need for cyber security measures within these systems was not anticipated. At the time, security for industrial control systems mean physically securing access to the network and the consoles that controlled the systems.

[0211] Internet-based technologies have since become part of modern industrial control systems. Widely available, low-cost IP devices have replaced proprietary solutions, which increases the possibility of cyber security vulnerabilities and incidents. Industrial control systems have adopted Internet-based solutions to promote corporate connectivity and remote access capabilities, and are being designed and implemented using industry standard computers, operating systems (OS) and network protocols. As a result, these systems may resemble computer networks. This integration supports new networking capabilities, but provides less isolation for industrial control systems from the outside world than predecessor systems. Networked industrial control systems may be exposed to similar threats as are seen

in computer networks, and an increased likelihood that an industrial control system can be compromised.

[0212] Industrial control system vendors have begun to open up their proprietary protocols and publish their protocol specifications to enable third-party manufacturers to build compatible accessories. Organizations are also transitioning from proprietary systems to less expensive, standardized technologies such as Microsoft Windows and Unix-like operating systems as well as common networking protocols such as TCP/IP to reduce costs and improve performance. Another standard contributing to this evolution of open systems is Open Platform Communications (OPC), a protocol that enables interaction between control systems and PC-based application programs. The transition to using these open protocol standards provides economic and technical benefits, but also increases the susceptibility of industrial control systems to cyber incidents. These standardized protocols and technologies have commonly known vulnerabilities, which are susceptible to sophisticated and effective exploitation tools that are widely available and relatively easy to use.

[0213] Industrial control systems and corporate networking systems are often interconnected as a result of several changes in information management practices, operational, and business needs. The demand for remote access has encouraged many organizations to establish connections to the industrial control system that enable of industrial control systems engineers and support personnel to monitor and control the system from points outside the control network. Many organizations have also added connections between corporate networks and industrial control systems networks to allow the organization's decision makers to obtain access to critical data about the status of their operational systems and to send instructions for the manufacture or distribution of product.

[0214] In early implementations this might have been done with custom applications software or via an OPC server/gateway, but, in the past ten years this has been accomplished with TCP/IP networking and standardized IP applications like File Transfer Protocol (FTP) or Extensible Markup Language (XML) data exchanges. Often, these connections were implemented without a full understanding of the corresponding security risks. In addition, corporate networks are often connected to strategic partner networks and to the Internet. Control systems also make more use of WANs and the Internet to transmit data to their remote or local stations and individual devices. This integration of control system networks with public and corporate networks increases the accessibility of control system

vulnerabilities. These vulnerabilities can expose all levels of the industrial control system network architecture to complexity-induced error, adversaries and a variety of cyber threats, including worms and other malware.

[0215] Many industrial control system vendors have delivered systems with dial-up modems that provide remote access to ease the burdens of maintenance for the technical field support personnel. Remote access can be accomplished, for example, using a telephone number, and sometimes an access control credential (e.g., valid ID, and/or a password). Remote access may provide support staff with administrative-level access to a system. Adversaries with war dialers – simple personal computer programs that dial consecutive phone numbers looking for modems – and password cracking software could gain access to systems through these remote access capabilities. Passwords used for remote access are often common to all implementations of a particular vendor's systems and may have not been changed by the end user. These types of connections can leave a system highly vulnerable because people entering systems through vendor-installed modems are may be granted high levels of system access.

[0216] Organizations often inadvertently leave access links such as dial-up modems open for remote diagnostics, maintenance, and monitoring. Also, control systems increasingly utilize wireless communications systems, which can be vulnerable. Access links not protected with authentication and/or encryption have the increased risk of adversaries using these unsecured connections to access remotely controlled systems. This could lead to an adversary compromising the integrity of the data in transit as well as the availability of the system, both of which can result in an impact to public and plant safety. Data encryption may be a solution, but may not be the appropriate solution in all cases.

[0217] Many of the interconnections between corporate networks and industrial control systems require the integration of systems with different communications standards. The result is often an infrastructure that is engineered to move data successfully between two unique systems. Because of the complexity of integrating disparate systems, control engineers often fail to address the added burden of accounting for security risks. Control engineers may have little training in security and often network security personnel are not involved in security design. As a result, access controls designed to protect control systems from unauthorized access through corporate networks may be minimal. Protocols, such as

TCP/IP and others have characteristics that often go unchecked, and this may counter any security that can be done at the network or the application levels.

[0218] Public information regarding industrial control system design, maintenance, interconnection, and communication may be readily available over the Internet to support competition in product choices as well as to enable the use of open standards. Industrial control system vendors also sell toolkits to help develop software that implements the various standards used in industrial control system environments. There are also many former employees, vendors, contractors, and other end users of the same industrial control system equipment worldwide who have inside knowledge about the operation of control systems and processes.

[0219] Information and resources are available to potential adversaries and intruders of all calibers around the world. With the available information, it is quite possible for an individual with very little knowledge of control systems to gain unauthorized access to a control system with the use of automated attack and data mining tools and a factory-set default password. Many times, these default passwords are never changed.

IV. DECEPTION PROFILER

[0220] In various implementations, the systems and methods discussed above can be used to implement a deception profiler (e.g., the deception profiler 1230) to manage a deployment of one or more deception mechanisms. As described above, the deception profiler can include at least one or more of a location engine (e.g., the location engine 1232), a density engine (e.g., the density engine 1234), a configuration engine (e.g., the configuration engine 1236), a scheduling engine (e.g., the scheduling engine 1238), or any combination thereof.

[0221] The location engine can determine where to deploy one or more deception mechanisms (e.g., a network, a subnetwork, a trunk, one or more machines, a portion of a network, or other suitable location in a network). In some examples, a trunk can include a plurality of subnetworks. In some examples, the one or more deception mechanisms can be deployed in the network. In other examples, the one or more deception mechanisms can be deployed in another network such that the one or more deception mechanisms can be projected in the network (as described above).

[0222] The density engine can determine a number of deception mechanisms to be deployed for the location. The configuration engine can determine one or more configurations for the number of deception mechanisms. A configuration for a deception

mechanism can include a MAC address, an Internet Protocol (IP) address, a type of operating system, a version of the operating system, one or more types of network services, any combination thereof, and any other suitable configurations. The scheduling engine can determine when to deploy the number of deception mechanisms. For example, the scheduling engine can determine at least one or more of a connect time, a disconnect time, or any combination thereof. The connect time can indicate when to connect a deception mechanism of the number of deception mechanisms to the network. The disconnect time can indicate when to disconnect a deception mechanism of the number of deception mechanisms from the network.

[0223] In some examples, the deception profiler can receive information associated with the network to use with the engines described above. For example, the deception profiler can receive a network topology. In some examples, the network topology can be determined using an active directory. In other examples, the network topology can be determined using a network discovery tool. In some examples, a user can change the network topology, including adding, modifying, or removing data in the network topology.

[0224] The network topology can include network information associated with one or more machines of the network. For example, the network information can include a number of subnetworks that are in the network and the machines that are in each subnetwork. The network information can also include a type of a subnetwork. Examples of types of subnetworks include human resource, finance, privileged users, source code, user's data, and data-backup systems.

[0225] The network information can also include information associated with a difficulty level of deploying a deception mechanism for a subnetwork. The difficulty level can be based on the number of deception mechanisms in the network. For example, a number of deception mechanisms in a network can cause a higher difficulty level. In some examples, the number of deception mechanisms can be relevant because the deception mechanisms must be maintained. For example, a list of deception mechanisms with their configurations and locations can be maintained. In addition, a need to refresh, alter, restart, or in some way remove a complication from a deception mechanism can arise when the deception mechanism is compromised.

[0226] The network information can also include a number and distribution of assets in a subnetwork in relation to the network. The number and distribution of assets can be separated

by category. Examples of categories can include type of server (e.g., email, DHCP server, database server, or other suitable type of server), type of machine (e.g., privilege user machine, end-user machine, security operations center machine, an active directory, or other type of machines), and type of asset (e.g., asset, critical asset, or other type of asset).

[0227] The deception profiler can also receive machine information of one or more machines on the network. The machine information for a machine can include one or more of a MAC address, an IP address, a type of operating system, a version of the operating system, one or more types of network services, or any combination thereof. In some examples, a user can change the machine information, including adding, modifying, or removing machines in the network topology.

[0228] The deception profiler can also receive historical attack information. The source of the historical attack information can depend on the type of system implemented in the network. The historical attack information can be received from at least one or more of a security operations center (SOC), a computer security incident response team (CSIRT), an intrusion detection system (IDS), an intrusion prevention system (IPS), or any combination thereof. The SOC can be a centralized unit that monitors, assesses, and defends a network. The SOC can perform real-time monitoring and management of a network, including aggregate logs, aggregate data, and coordinate response and remediation. The SOC can also report attacks and perform post-attack analysis. Post-attack analysis can include forensics and investigation to determine a source of an attacker. The CSIRT is a system that receives reports of security breaches, conducts an analysis of the reports, and reacts to the attack. The IDS is a system that monitors network and system activities for malicious activities. The IPS also monitors network and system activities for malicious activity, but also actively prevent/block intrusions that are detected.

[0229] Other data sources for the historical attack information can include existing deception mechanism attack information, threat feeds, vulnerabilities, and privilege user management data. In some examples, the existing deception mechanism attack information can be associated with attacks detected on one or more machines in the network. In other examples, the existing deception mechanism attack information can be associated with one or more networks other than the network. In some examples, the historical attack information can include a distribution of attacks on a type of mechanism (e.g., a honeypot) using threat intelligence feeds of historical attack data. In other examples, the historical attack information

can include a distribution of threats intelligence for an industry. In some examples, the deception profiler can receive a distribution of historical attacks for a data source. In other examples, the deception profiler can determine a distribution of historical attacks for a data source. In some examples, a user can change the historical information, including adding, modifying, or removing data used for the historical information.

[0230] FIG. 12 illustrates an example of a subnetwork 1210 identified by a location engine (e.g., the location engine 132). In some examples, the subnetwork 1210 can be in a network (e.g., the network 400 described above). However, it should be recognized that the network can be in any form, including as described in FIGS. 5-11. In FIG. 12, the subnetwork 1210 includes the hub 422, the laptop computers 426, and the desktop computers 424. However, it should be recognized that a subnetwork can be identified that includes other devices (e.g., all devices connected to the addition switch 406b, the laptop computers 426, the site network 400, or some other division. In some examples, an indication of a location can be received by the location engine from a user.

[0231] The location engine can identify the location by computing an importance score. The importance score can use a network topology of a network to compare portions of the network (a subnetwork can be an example of a portion of the network). In such examples, the location engine can compare the one or more asset densities, as described below, that are associated with a plurality of portions of a network to identify the location with the highest score. For example, the location engine can identify a subnetwork that includes the most critical assets. In some examples, the importance score can further use machine information associated with the network. For example, the importance score can use the types of assets in a portion of the network.

[0232] In other examples, the location engine can identify a location using a distribution of historical attacks on the network. For example, the location engine can identify a subnetwork that includes the most historical attacks on the network as a location for deploying a number of deception mechanisms.

[0233] In some examples, the location engine can update the location of one or more deception mechanisms. For example, the location can be updated when an attack occurs on the network. In such an example, when an attack occurs on the network, the deception profiler can determine the location where the attack occurred. In such examples, the deception profiler can detect a request to access a deception mechanism. In other examples,

the deception mechanism can send a notification to the deception profiler that a request to access a deception mechanism has been received. In response to the request, the deception profiler can determine a location of the accessed deception mechanism in order to update the location of one or more of the deployed deception mechanisms using the location of the accessed deception mechanism.

[0234] In some examples, the location engine can update the location of one or more deception mechanisms when a certain number of attacks occur on the network. For example, the location engine can determine a probability distribution of the attacks on the network. The probability distribution can statistically represent the number of attacks on the network over a time period. In some examples, the probability distribution can include one or more types of attacks on the network. The location engine can use the probability distribution of the attacks on the network to determine a location that includes more attacks. For example, the location engine can determine that more attacks have occurred on a particular portion of the network than another. In such an example, the location engine can determine to analyze the portion of the network to determine a number of deception mechanisms to deploy.

[0235] FIG. 13 illustrates an example of a network interaction graph 1300. In some examples, the network interaction graph 1300 can be based on one or more devices and/or systems in a network. In such examples, the network can be identified for deploying one or more deception mechanisms. For example, the network interaction graph 1300 can include one or more routers (e.g., a first router 1310, a second router 1320, and a third router 1330) and one or more machines (e.g., M₁ 1340, M₂ 1342, M₃ 1344, M₄ 1346, M₅ 1348, M₆ 1350, M₇ 1352, M₈ 1354, M₉ 1356, and M₁₀ 1358). The lines between the machines and the routers in the network 1300 can indicate a communication. However, it should be recognized that the network 1300 can resemble any of the networks described herein or the like, including the networks described in FIGS. 1-11.

[0236] In some examples, network interactions illustrated by the network interaction graph 1300 can be captured using network interaction events (e.g., Bro logs and network topology). In such examples, the network interactions can occur among endpoints and network elements (e.g., routers, machines, and the like). Using graph analytics, central nodes can be determined in the network interaction graph 1300. For example, a number of incoming connections to a network element can be compared between network elements to identify a network element

with the most incoming connections. It should be recognized that other methods of graph analytics can be used to identify central nodes.

[0237] In some examples, central nodes can be limited to a type of network element (e.g., routers or machines). For example, if a central node is limited to routers, the first router 1310 can be identified as the central node when the central node is based on a number of incoming connections. However, if a central node is limited to machines, the M₉ 1356 can be identified as the central node when the central node is based on a number of incoming connections.

[0238] In some examples, once a central node is determined, one or more deception mechanisms can be deployed on or near the central node. By placing the one or more deception mechanisms on or near the central node, there can be a higher probability of catching an attacker because the attacker is likely to at least go through the central node.

[0239] As described above, the deception profiler can include the density engine. The density engine can determine a number of deception mechanisms to deploy for the network using at least one or more of a network topology, machine information, historical attack information, or any combination thereof (as described herein). In some examples, the density engine can analyze each portion of the network. In other examples, the density engine can analyze a portion of the network identified by the location engine. In some examples, the density engine can use the network topology, the machine information, and/or the historical attack information to determine one or more asset densities, a summary statistic, or any combination thereof. In some examples, a user can modify the one or more asset densities and/or a summary statistic after they are computed.

[0240] In some examples, the density engine can determine one or more asset densities. An asset density can be associated with a number of assets connected to the network. In some examples, an asset can be a critical asset. For example, the asset density can be a total number of critical assets in a portion of the network (e.g., a subnetwork, a trunk, one or more machines, or other suitable location in the site network) divided by a total number of critical assets in the network. An asset that is critical can be in terms of information security. For example, a critical asset can include a machine that stores network data or a privileged user account that has broad access to the network. In some examples, a critical asset can be user-defined. In other examples, a critical asset can be industry specific. In some examples, an asset density can be a total number of assets (whether critical or not) in a portion of the

network (e.g., a subnetwork, a trunk, one or more machines, or other suitable location in the site network) divided by a total number of assets in the network.

[0241] The density engine can also determine one or more summary statistics. A summary statistic can be associated with a number of historical attacks on the network. In some examples, the summary statistic can include a mean, medium, or mode of a probability distribution of the number of historical attacks on the network. In some examples, the probability distribution can be received by the deception profiler. In other examples, the deception profiler can determine the probability distribution. Because a summary statistic of a probability distribution is used, the probability distribution can be in a parametric form (e.g., normal distribution), a nonparametric form, or any other form that can be summarized using a mean, medium, or mode.

[0242] Using at least one or more of the one or more asset densities, the summary statistic, or any combination thereof, the density engine can compute a mixture density model that is used to determine a number of deception mechanisms to deploy. In particular the number of deception mechanisms to deploy in the network can be determined by:

[0243] $N_i =$

$$w_1 * p_i^c(s) * N_s + w_2 * p_i(s) * N_s + \left(\frac{w_4 * p_t(h) * N_{h_t} + w_5 * p_t(ids) * N_{ids_t} + w_6 * p_t(ips) * N_{ips_t}}{N_s} \right), \text{ where:}$$

[0244] (1): N_i is the number of deception mechanisms to deploy in a subnetwork i ;

[0245] (2): N_s is the total number of subnetworks;

[0246] (3): N_{h_t} is the total number of historical attacks over time t using SOC/CSIRT;

[0247] (4): N_{ids_t} is the total number of historical attacks over time t using IDS;

[0248] (5): N_{ips_t} is the total number of historical attacks over time t using IPS;

[0249] (6): $w = \{w_1, w_2, w_3, \dots, w_n\} =$ a set of weights;

[0250] (7): $p_i^c(s) = \frac{a_{c_i}}{\sum_{i=1}^{N_s} a_{c_i}} =$ the probability of placing a deception mechanism

in a subnetwork (i) based on critical assets, where a_{c_i} is the number of critical assets;

[0251] (8): $p_i(s) = \frac{a_i}{\sum_{i=1}^{N_s} a_i} =$ *the probability of placing a deception mechanism*

in a subnetwork i based on assets;

[0252] (9): $p_t(h) = \frac{f(\mu_h)}{N_{h_t}},$ *the probability of an attack over a time t using SOC/CSIRT;*

[0253]

$$(10): p_t(ids) = \frac{f(\mu_{ids})}{N_{ids_t}} =$$

the probability of an attack over a time t using IDS;

[0254]

$$(11): p_t(ips) = \frac{f(\mu_{ips})}{N_{ips_t}} =$$

the probability of an attack over a time t using IPS;

[0255] (12): $\mu =$ *a summary statistic of a probability distribution of historical attacks over a time t , where the average can be a mean, meadium, or mode; and*

[0256] (13): $f(x) =$ *a math function, such as logarithmic or square root.*

[0257] In the above equations, it is assumed that all of the above data sources are available. If a data source is unavailable, a term associated with the data source can be dropped from the equation N_i .

[0258] As described above, the equation N_i illustrates that the number of deception mechanisms to deploy in a portion of a network (e.g., a subnetwork, a trunk, one or more machines, or other suitable location in the network) can depend on information associated with the portion of the network. For example, when more assets are on the portion of the network, the number of deception mechanisms can increase. In another example, an increased number of attacks on a portion of the network can increase the number of deception mechanisms in all portions of the networks a similar amount in some examples. In other examples, the probability distributions can be associated with a probability of an attack in a portion of the network. In such examples, an increased number of attacks in the portion of the network can increase the number of deception mechanisms to deploy in the portion of the network. In some examples, a probability distribution can be modified by a user.

[0259] FIG. 14A illustrates an example of a probability distribution. In particular, FIG. 14A illustrates a histogram 1400 of IDS attacks. The x-axis can indicate types of attacks. For example, a type of attack can include an IP Mobility activity. The IP Mobility activity represents a denial-of-service vulnerability on Cisco IOS in an IPv4 network environment. Another type of attack can include a Bind NXT Overflow. The Bind NXT Overflow exploits a vulnerability in Bind, a DNS daemon. Other attacks can include, but are not limited to, Oracle TNS Command Execution, Oracle Database DBMS TNS Listener Denial of Service, Ascend Router, Microsoft Windows SMTP Service DNS Response Spoofing, LinkSys Cable/DSL Router Gozila DomainName Parameter denial of service, SCO Openserver named Bugger Overflow (1), ISC Bind CNAME RRSIG Query with RPZ denial of service, and other examples listed in FIG. 14A. One of ordinary skill in the art will appreciate that the attacks shown in FIG. 14A are illustrative examples, and that any other suitable attacks can be represented in the histogram 1400.

[0260] The y-axis of the histogram 1400 can indicate a number of the types of attacks. For example, the IP Mobility activity attack can include a bar that indicates 32 attacks. In such an example, the 32 IP Mobility activity attacks might have occurred on the network. In other examples, the y-axis can indicate probabilities of types of attacks. A probability of a type of attack can be a measure of how likely the type of attack is to occur out of a number of possible attacks.

[0261] FIG. 14B illustrates another example of a probability distribution. In particular, FIG. 14B illustrates a MAC address histogram 1410. For the MAC address histogram 1410, the x-axis can indicate types of MAC addresses. In some examples, types of MAC addresses can include, but are not limited to, Cisco Systems, Inc., Apple, Samsung Electronics, Arris Group, Inc., Intel Corporation, and other examples listed in FIG. 14B. One of ordinary skill in the art will appreciate that the MAC address types shown in FIG. 14B are illustrative examples, and that any other suitable MAC address types can be represented in the histogram 1410.

[0262] The y-axis of the MAC address histogram 1410 can indicate a number of machines with the types of MAC addresses. For example, the Cisco Systems, Inc. MAC address can include a bar that indicates 280 machines with a Cisco Systems, Inc. MAC address. A probability of a type of MAC address can be computed by dividing a number of a type of MAC address by the total number of machines on a network.

[0263] FIG. 14C illustrates another example of a probability distribution. In particular, FIG. 14C illustrates a network service histogram 1420. For the network service histogram 1420, the x-axis can indicate types of network services running on one or more ports in a network. In some examples, types of network services can include, but are not limited to, File Transfer Protocol (FTP), Secure Shell (SSH), and Telnet. One of ordinary skill in the art will appreciate that the network services shown in FIG. 14C are illustrative examples, and that any other suitable network services can be represented in the histogram 1420.

[0264] The y-axis of the network service histogram 1420 can indicate a number of ports on the network with the types of network services. For example, the FTP service can include a bar that indicates 8 ports with a FTP service. A probability of a number of ports with a type of network service can be computed by dividing a number of ports on a network with a type of network service by the total number of ports on the network.

[0265] FIG. 15 is a flowchart illustrating an example of a process for deploying a number of deception mechanisms for a network. In some aspects, the process 1500 can be performed by a deception profiler. While specific examples can be given of a deception mechanism, one of ordinary skill in the art will appreciate that other deception mechanisms can be included in the process 1500.

[0266] The process 1500 is illustrated as a logical flow diagram, the operation of which represent a sequence of operations that can be implemented in hardware, computer instructions, or a combination thereof. In the context of computer instructions, the operations represent computer-executable instructions stored on one or more computer-readable storage media that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular data types. The order in which the operations are described is not intended to be construed as a limitation, and any number of the described operations can be combined in any order and/or in parallel to implement the processes.

[0267] Additionally, the process 1500 can be performed under the control of one or more computer systems configured with executable instructions and can be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations thereof. As noted above, the code can be stored on a machine-readable storage medium, for example, in

the form of a computer program comprising a plurality of instructions executable by one or more processors. The machine-readable storage medium can be non-transitory.

[0268] At step 1510, the process 1500 includes identifying a network for which to deploy one or more deception mechanisms. In some examples, the network can be a portion of a network, a subnetwork, a trunk, one or more machines, or other suitable location in a network. In some examples, a deception mechanism can emulate one or more characteristics of a machine on the network. Identifying can include receiving an identification of the network. In such an example, the network can be identified as the network to determine a number of deception mechanisms to deploy. In other examples, identifying can include randomly selecting the network from a plurality of networks. In other examples, one or more asset densities, as described above, can be determined for each network of the plurality of networks. In such examples, the network with the highest asset density can be identified. A person of ordinary skill in the art will recognize that a network can be identified in other ways. For example, the network can be identified using an importance score. In some examples, the importance score can be computed using one or more asset densities and machine information associated with the network. The machine information can include types of assets in the network.

[0269] At step 1520, the process 1500 further includes determining one or more asset densities. An asset density can be associated with a number of assets connected to the network. In some examples, the network can be a subnetwork of a site network. In such examples, the asset density can be a total number of assets in a portion of the site network (e.g., a subnetwork, a trunk, one or more machines, or other suitable location in the site network) divided by a total number of assets in the site network. In other examples, the asset density can be a total number of critical assets in a portion of the site network divided by a total number of critical assets in the site network. In some examples, a first asset density can be associated with the total number of assets and a second asset density can be associated with the total number of critical assets.

[0270] At step 1530, the process 1500 further includes determining a summary statistic associated with a number of historical attacks on the network. In some examples, the summary statistic can include a mean, medium, or mode of a probability distribution of the number of historical attacks. In some examples, the probability distribution of the number of historical attacks can include one or more probability distributions for one or more time

periods. In other examples, the probability distribution of the number of historical attacks can include one or more probability distributions for one or more types of attacks. In some examples, the summary statistic can include one or more data source summary statistics. A data source summary statistic can be associated with a data source described above (e.g., SOC, CSIRT, IDS, or IPS).

[0271] At steps 1540 and 1550, the process 1500 further includes determining a number of deception mechanisms to deploy and deploying the number of deception mechanisms. In some examples, the number of deception mechanisms can be determined using at least one or more of the one or more asset densities, the summary statistic, or a combination thereof.

[0272] In some examples, the density engine can update the number of deception mechanisms. In particular, the deception profiler can monitor requests to access deception mechanisms in order to determine whether to deploy more deception mechanisms. In some examples, the density engine can update the summary statistic. In particular, the density engine can compute an updated probability distribution of historical attacks by combining the probability distribution of historical attacks and a probability distribution of new attacks. The probability distribution of new attacks can be the attacks that have occurred since the historical attacks were received.

[0273] In some examples, to compute the updated probability distribution of historical attacks, the probability distribution of historical attacks and the probability distribution of new attacks can each be associated with a weight such that the updated probability distribution favors one of the probability distributions over the other. In some examples, the density engine will favor the probability distribution of historical attacks when the recent probability distribution has a small sample size. For example, a small sample size can occur when only a few attacks have occurred on the network. A person of ordinary skill will recognize how to weigh the two probability distributions based on the circumstances. In addition, if the one or more machines on the network have changed, the machine information component of the mixture density model can be modified.

[0274] FIG. 16 is a flowchart illustrating an example of a process for updating a probability distribution in order to change the number of deception mechanisms deployed. In some aspects, the process 1600 can be performed by a deception profiler. While specific examples can be given of a deception mechanism, one of ordinary skill in the art will appreciate that other deception mechanisms can be included in the process 1600.

[0275] The process 1600 is illustrated as a logical flow diagram, the operation of which represent a sequence of operations that can be implemented in hardware, computer instructions, or a combination thereof. In the context of computer instructions, the operations represent computer-executable instructions stored on one or more computer-readable storage media that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular data types. The order in which the operations are described is not intended to be construed as a limitation, and any number of the described operations can be combined in any order and/or in parallel to implement the processes.

[0276] Additionally, the process 1600 can be performed under the control of one or more computer systems configured with executable instructions and can be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations thereof. As noted above, the code can be stored on a machine-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors. The machine-readable storage medium can be non-transitory.

[0277] At step 1610, the process 1600 includes receiving a probability distribution of historical attacks on the network. In some examples, the probability distribution of historical attacks on the network can indicate the number of attacks on the network over a time. In some examples, the probability distribution of historical attacks can be the same probability distribution of historical attacks used to determine the number of deception mechanisms to deploy.

[0278] At step 1620, the process 1600 further includes receiving a probability distribution of new attacks on the network. In some examples, the probability distribution of new attacks can be received after a particular number of attacks have occurred on the network. In other examples, the probability distribution of new attacks can be received after a set amount of time has lapsed. For example, the probability distribution of new attacks can be received once a week, every two weeks, every month, or any other suitable period of time. In other examples, the probability distribution of new attacks can be determined by the deception profiler.

[0279] At step 1630, the process 1600 further includes computing a posteriori distribution in order to update the probability distribution of historical attacks used for determining the number of deception mechanisms. In some examples, the summary statistic of the posteriori distribution can be a combination of the summary statistics of the probability distribution of historical attacks and the probability distribution of new attacks. In some examples, weights are associated with each probability distribution. The standard deviation for the posteriori distribution can similarly be computed.

[0280] In some examples, the process 1600 can return to step 1620 after a set number of attacks occur on the network. For example, the process can return to step 1620 after another attack occurs on the network. In another example, the process can return to step 1620 after N attacks have occurred on the network, similar to described above in reference to step 1620. In other examples, the process can return to step 1620 after an amount of time has passed since the last probability distribution of new attacks was received. A person of ordinary skill in the art will recognize that other measures can be used to determine when to update the number of deception mechanisms deployed.

[0281] To illustrate the updating process of the number of deception mechanisms described above, the following equations can be used:

[0282] *the probability distribution of historical attacks* $\sim N(\mu, \sigma^2)$;

[0283] *The probability distribution of new attacks* $\sim N(\alpha, \beta^2)$;

[0284] *The posteriori distribution* $\sim N(\mu_p, \sigma_p^2)$;

[0285] *The mean of the posteriori distribution* $\mu_p = \frac{\beta^2\mu + \sigma^2\alpha}{\beta^2 + \sigma^2}$; and

[0286] *The standard deviation of the posteriori distribution* $\sigma_p^2 = \frac{\beta^2\sigma^2}{\beta^2 + \sigma^2}$.

[0287] The procedure above can repeat when either a particular number of attacks have occurred or after a set amount of time has lapsed. In some examples, any of the distributions described herein (including the probability distribution and the posteriori distribution, can be modified by a user.

[0288] As described above, the deception profiler can include the configuration engine. The configuration engine can determine one or more configurations of the number of deception mechanisms. In some examples, a configuration can be associated with machine

information of a deception mechanism. For example, the configuration can include at least one or more of a MAC address, an IP address, a host name, a type of operating system, a version of the operating system, one or more network services, or any combination thereof.

[0289] In some examples, the configuration engine can compute a population distribution for one or more machines of the network. The population distribution can identify machine information associated with the one or more machines of the network. For example, the population distribution can identify the number of the one or more machines that are running a type of operating system, a type of network service, a MAC address, or any other machine information associated with a machine.

[0290] Using the population distribution, a configuration for a deception mechanism can be determined. The configuration engine can analyze the population distribution to determine a configuration that is included in the population distribution. In some examples, the configuration can be determined based on one or more other configurations that have been determined. For example, the configuration engine can assign a percentage to each type of machine information that corresponds to the percent of the one or more machines on the network with that type of machine information. The configuration engine can then configure deception mechanisms in the network to correspond to the percentage of each type of machine information of the one or more machines. For example, the configuration engine can determine to configure a first number of deception mechanisms with Microsoft Windows, a second number of deception mechanisms with Mac OS, and a third number of deception mechanisms with Linux.

[0291] To illustrate determining a configuration, one or more machines of the network can be identified. A type of machine information can be identified for the one or more machines of the network. In some examples, the type of machine information can be identified based on undetermined machine information associated with a deception mechanism. For example, a configuration associated with a deception mechanism might not have an operating system determined yet. In such an example, a population distribution associated with types of operating systems can be used to determine an operating system for the configuration associated with the deception mechanism. In some examples, a configuration can be determined to blend with the one or more machines, or a portion of the one or more machines. Blending can include configuring the deception mechanism to match one or more characteristics of a machine on the network.

[0292] In some examples, the machine information used for the population distribution can include data that is unique to a machine. For example, a MAC address and a host name can be unique to a machine. For a MAC address, the population distribution can indicate an organizational unique identifier (OUI) to be used for the MAC address. For a host name, the population distribution can indicate a portion of the host name for a deception mechanism. The rest of the host name can be determined by randomly including at least one or more of a character or a number to the portion of the host name. In some examples, at least one or more of the character, the number, or a combination thereof can be randomly chosen.

[0293] To determine one or more network services, the configuration engine can identify a distribution of network services on one or more machines in the network. The distribution of network services can be a probability distribution of the network services. For example, the probability distribution of the network services can indicate a number of times that a file transfer protocol service is present and a number of times that a secure shell protocol is present on a machine in the network. The configuration engine can determine to configure one or more deception mechanisms using the probability distribution of the network services (e.g., file transfer protocol service, secure shell protocol, or other suitable network service). In particular, using the information from the population distribution to identify configurations of machines on the network, intelligent configurations can be determined for deception mechanisms.

[0294] FIG. 17 illustrates examples of host names being determined for deception mechanisms. The examples can be associated with a network (e.g., 10.75.36.0/23). In some examples, the network can include two different patterns (e.g., a first pattern 1710 and a second pattern 1740). However, it should be recognized that networks can include more or less patterns. In some examples, a system to identify a pattern can utilize various data mining techniques (e.g., clustering, similarity distance, or the like).

[0295] In some examples, a pattern can be associated with one or more devices. Host names (or some other type of configuration of a device) of the one or more devices can be according the pattern. For example, the first pattern 1710 can be “CTRL-##-SERVER,” where each host name of each device associated with the first pattern 1710 is in this form with “##” being a number that is counting up from 01. For example, a first device can include a host name of “CTRL-01-SERVER” and a second device can include a host name of “CTRL-02-SERVER.” Based on the first pattern 1710, it can be determined that a deception

mechanism that is to be deployed according to the first pattern 1710. For example, a host name of the deception mechanism may be “CTRL-03-SERVER,” where the CTRL and SERVER is the same as the other host names and the number (i.e., “03”) is incremented 1 from the other host names.

[0296] While the first pattern 1710 illustrates a simple example, the second pattern 1740 illustrates an example with numbers that are not incremented and numbers that are skipped. For example, the second pattern can be “METRO-W01-#.” As can be seen, the W01 does not change between the devices on the network (e.g., METRO-W01-5, METRO-W01-6, and METRO-W01-8). In addition, it can be seen that METRO-W01-7 was skipped. Therefore, in some examples, a host name of a deception mechanism deployed according to the second pattern 1740 may be “METRO-W01-9.” In other examples, a host name of a deception mechanism deployed according to the second pattern 1740 may be “METRO-W01-7.” It should be recognized that other patterns can be recognized and other methods of changing the patterns for deception mechanisms can be used.

[0297] In some examples, the configuration engine can update the one or more configurations after an event occurs. The event can include an attack associated with the network. In other examples, the event can include a number of attacks associated with the network or an amount of time.

[0298] FIG. 18 is a flowchart illustrating an example of a process for updating one or more configurations for one or more deception mechanisms. In some aspects, the process 1800 can be performed by a deception profiler. While specific examples can be given of a deception mechanism, one of ordinary skill in the art will appreciate that other deception mechanisms can be included in the process 1800.

[0299] The process 1800 is illustrated as a logical flow diagram, the operation of which represent a sequence of operations that can be implemented in hardware, computer instructions, or a combination thereof. In the context of computer instructions, the operations represent computer-executable instructions stored on one or more computer-readable storage media that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular data types. The order in which the operations are described is not intended to be construed as a

limitation, and any number of the described operations can be combined in any order and/or in parallel to implement the processes.

[0300] Additionally, the process 1800 can be performed under the control of one or more computer systems configured with executable instructions and can be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations thereof. As noted above, the code can be stored on a machine-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors. The machine-readable storage medium can be non-transitory.

[0301] At step 1810, the process 1800 includes computing a categorical distribution for one or more deception mechanisms deployed for the network. The categorical distribution can be associated with a type of machine information. For example, the categorical distribution can be associated with a type of operating system, a MAC address, a type of service, or any other type of machine information.

[0302] FIG. 19A illustrates an example of a categorical distribution 1910. In the categorical distribution 1910, the x-axis can indicate each type of a machine information and the y-axis can indicate a number of occurrences. The categorical distribution 1910 can include a bar associated with α_1 , a bar associated with α_2 , and a bar associated with α_3 . The categorical distribution 1910 can include more or less bars, depending on the number of types of machine information. The categorical distribution 1910 can describe a frequency of each type of operating system for the one or more deception mechanisms. For example, the categorical distribution 1910 can describe that there are seven deception mechanisms with α_1 , which can be associated with Microsoft Windows, two deception mechanisms with α_2 , which can be associated with Linux, and one deception mechanism with α_3 , which can be associated with Mac OS. In some examples, the configuration engine can receive the categorical distribution 1910. In some examples, the categorical distribution 1910 can be modified by a user.

[0303] At step 1820, the process 1800 further includes computing a distribution (e.g., a Dirichlet distribution) for one or more deception mechanisms that were associated with an attack during a first event. As discussed above, the first event can be one or more attacks (e.g., 100 attacks) associated with the network. In such an example, the Dirichlet distribution can include each deception mechanism that was attacked during the 100 attacks. In other examples, the first event can be an amount of time. In such examples, the Dirichlet

distribution can include each deception mechanism that was attacked during the amount of time prior to a current time. The Dirichlet distribution can be associated with a type of machine information. In some examples, the Dirichlet distribution can be modified by a user.

[0304] FIG. 19B illustrates an example of a Dirichlet distribution 1920. In the Dirichlet distribution 1920, the x-axis can indicate each type of category and the y-axis can indicate a number of occurrences. The Dirichlet distribution 1910 can include a bar associated with c_1 , a bar associated with c_2 , and a bar associated with c_3 . The Dirichlet distribution 1920 can include more or less bars, depending on the number of types of categories. The Dirichlet distribution 1920 can describe a frequency of each type of operating system for one or more deception mechanisms that are associated with an attack during the first event. For example, the Dirichlet distribution 1920 can describe that there are twenty deception mechanisms attacked with operating system c_1 , which can be associated with Microsoft Windows, seventy deception mechanisms attacked with operating system c_2 , which can be associated with Linux, and ten deception mechanisms attacked with operating system c_3 , which can be associated with Mac OS. In some examples, the configuration engine can receive the Dirichlet distribution 1920.

[0305] At step 1830, the process 1800 further includes performing a Bayesian update for the one or more deception mechanisms deployed for the network using a posteriori distribution. The posteriori distribution can be a Dirichlet distribution that is derived using the categorical distribution as a prior distribution and the Dirichlet distribution as a likelihood distribution. The posteriori distribution can be associated with a type of machine information. To continue the example from the steps 1810 and 1820, the posteriori distribution can be associated with a type of operating system. One of ordinary skill in the art will appreciate that the posteriori distribution can be associated with other machine information, such as a MAC address, a type of service on a port, or the like. In such an example, the posteriori distribution can describe an expected value of each type of operating system for deception mechanisms.

The expected value can equal a percentage of deception mechanisms that should be configured to be a particular type of operating system. For example, the expected value for a first operating system can be 0.2, which indicates that 20% of the deception mechanisms should be configured with the first operating system. The expected value equation can be

$E(P_i | X, \alpha) = \frac{c_i + \alpha_i}{N + \sum_k \alpha_k}$. Using the expected value equation with the distributions above, the

expected value for Microsoft Windows can equal $\frac{20+7}{100+(7+2+1)} = \frac{27}{110} = 0.245$. Using the

expected value equation with the distributions above, the expected value for Linux can equal $\frac{70+2}{100+(7+2+1)} = \frac{72}{110} = 0.654$. Using the expected value equation with the distributions above, the expected value for Mac OS can equal $\frac{10+1}{100+(7+2+1)} = \frac{11}{110} = 0.1$. In such examples, the deception mechanisms deployed for the network should be updated to include 24.5% Microsoft Windows, 65.4% Linux, and 10% Mac OS.

[0306] When the Bayesian update is complete, the method can return to step 1820 after a second set of attacks. In other examples, the method can return to step 1820 after a second event. In some examples, the second event can be similar to the first event. For example, if the first event is 100 attacks, the second event can be an additional 100 attacks. The updating process can be performed periodically based on number of events, time, or other suitable event as described herein.

[0307] The scheduling engine can determine a time to deploy a deception mechanism. In some examples, the scheduling engine can use the historical attack information to determine the time. In other examples, the scheduling engine can use the machine information. In particular, the scheduling engine can analyze at least one or more of a connect time, a disconnect time, or a combination thereof of one or more machines on the network. The scheduling engine can determine a connect time and a disconnect time for one or more deception mechanisms so as to blend a visible, or active, time of the one or more deception mechanisms with the one or more machines on the network. The visible, or active, time of a deception mechanism can be the time that the deception mechanism is connected to the network. In some examples, the visible time can include time that there is a threshold of activity on the network. For example, the scheduling engine can determine to connect a deception mechanism to the network before a network is busy, and disconnect the deception mechanism after the network is busy.

[0308] While the connect time and disconnect time for the one or more machines on the network can be associated with an actual connect and disconnect from the network, the connect time and disconnect time for the one or more deception mechanisms can indicate when to have the one or more deception mechanisms appear to connect to and disconnect from the network. In some examples, the one or more deception mechanisms can appear to connect and disconnect by becoming visible and invisible to a machine on the network. In other examples, a deception mechanism can be connected to another network such that the deception mechanism is visible on the network. In such an example, the deception

mechanism can remain connected to the other network when the deception mechanism appears to disconnect from the network.

[0309] In some examples, the scheduling engine can use a distribution to determine a schedule. FIG. 20 illustrates an example of a probability distribution of a number of attacks over time that can be used to determine a schedule. In some examples, the x-axis can indicate a time period (e.g., 8 AM to 5 PM Monday through Friday, a previous week, a previous month, or any other time period). In other examples, the x-axis can indicate groups of time periods. For example a group of time periods can include morning, afternoon, and evening. In another example, a group of time periods can include weekdays and weekends.

[0310] The y-axis can indicate a number of attacks on the network. A probability of a number of attacks can be computed by dividing a number of attacks on a network during a subset of a time period by a total number of attacks on the network during the time period. While FIG. 20 illustrates a probability distribution using percentages, in some examples, a number of attacks on the network can be a number of the actual attacks rather than a probability. Using a probability distribution, the scheduling engine can determine to have deception mechanisms connected to a network during times that there are more or equal to a certain number/probability of attacks.

[0311] In some examples, the schedule of the one or more deception mechanisms can be updated when a certain amount of time has passed since the last schedule was determined. For example, the schedule can update weekly. In other examples, the schedule can update when a certain number of attacks have occurred on the network. In such examples, the deception profiler can detect a request to access a deception mechanism. In other examples, the deception mechanism can send a notification to the deception profiler that a request has been received by the deception mechanism. In response to the access, the deception profiler can determine a time that the deception mechanism was accessed. The deception profiler can then update the schedule using the time of the accessed deception mechanism.

[0312] FIG. 21 is a flowchart illustrating an example of a process for deploying a number of deception mechanisms for a network. In some aspects, the process 2100 can be performed by a deception profiler. While specific examples can be given of a deception mechanism, one of ordinary skill in the art will appreciate that other deception mechanisms can be included in the process 2100.

[0313] The process 2100 is illustrated as a logical flow diagram, the operation of which represent a sequence of operations that can be implemented in hardware, computer instructions, or a combination thereof. In the context of computer instructions, the operations represent computer-executable instructions stored on one or more computer-readable storage media that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular data types. The order in which the operations are described is not intended to be construed as a limitation, and any number of the described operations can be combined in any order and/or in parallel to implement the processes.

[0314] Additionally, the process 2100 can be performed under the control of one or more computer systems configured with executable instructions and can be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations thereof. As noted above, the code can be stored on a machine-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors. The machine-readable storage medium can be non-transitory.

[0315] At step 2110, the process 2100 includes identifying a network for which to deploy one or more deception mechanisms. In some examples, a deception mechanism can emulate one or more characteristics of a machine on the network.

[0316] At step 2120, the process 2100 further includes determining a number of deception mechanisms to be deployed for the network. In some examples, the number of deception mechanisms can be determined using current and historical network information. In some examples, the current network information can include one or more asset densities. An asset density can be associated with a number of assets connected to the network. For example, a first asset density can be associated with total assets and a second asset density can be associated with critical assets. In some examples, the historical network information can include information on historical attacks on the network.

[0317] At step 2130, the process 2100 further includes determining one or more configurations of the number of deception mechanisms using a population distribution. A configuration can include at least one or more of a MAC address, a host name, an operating

system, or a network service. In some examples, the population distribution can identify machine information associated with one or more machines in the network.

[0318] At step 2140, the process 2100 further includes determining a schedule to deploy the number of deception mechanisms using the historical network information. In some examples, the schedule can include at least one or more of a connect time, a disconnect time, or a combination thereof. The connect time can indicate when to connect a deception mechanism to the network. The disconnect time can indicate when to disconnect a deception mechanism from the network.

[0319] At step 2150, the process 2100 further includes deploying the number of deception mechanisms. In some examples, the number of deception mechanisms can be configured based on the one or more configurations. In some examples, the number of deception mechanisms can be deployed using the schedule. In some examples, the number of deception mechanisms can be deployed by a deception center.

[0320] Specific details were given in the preceding description to provide a thorough understanding of various implementations of systems and components for network threat detection and analysis. It will be understood by one of ordinary skill in the art, however, that the implementations described above may be practiced without these specific details. For example, circuits, systems, networks, processes, and other components may be shown as components in block diagram form in order not to obscure the examples in unnecessary detail. In other instances, well-known circuits, processes, algorithms, structures, and techniques may be shown without unnecessary detail in order to avoid obscuring the examples.

[0321] It is also noted that individual implementations may be described as a process which is depicted as a flowchart, a flow diagram, a data flow diagram, a structure diagram, or a block diagram. Although a flowchart may describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged. A process is terminated when its operations are completed, but could have additional steps not included in a figure. A process may correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination can correspond to a return of the function to the calling function or the main function.

[0322] The term “computer-readable medium” includes, but is not limited to, portable or non-portable storage devices, optical storage devices, and various other mediums capable of storing, containing, or carrying instruction(s) and/or data. A computer-readable medium may include a non-transitory medium in which data can be stored and that does not include carrier waves and/or transitory electronic signals propagating wirelessly or over wired connections. Examples of a non-transitory medium may include, but are not limited to, a magnetic disk or tape, optical storage media such as compact disk (CD) or digital versatile disk (DVD), flash memory, memory or memory devices. A computer-readable medium may have stored thereon code and/or machine-executable instructions that may represent a procedure, a function, a subprogram, a program, a routine, a subroutine, a module, a software package, a class, or any combination of instructions, data structures, or program statements. A code segment may be coupled to another code segment or a hardware circuit by passing and/or receiving information, data, arguments, parameters, or memory contents. Information, arguments, parameters, data, etc. may be passed, forwarded, or transmitted via any suitable means including memory sharing, message passing, token passing, network transmission, or the like.

[0323] The various examples discussed above may further be implemented by hardware, software, firmware, middleware, microcode, hardware description languages, or any combination thereof. When implemented in software, firmware, middleware or microcode, the program code or code segments to perform the necessary tasks (e.g., a computer-program product) may be stored in a computer-readable or machine-readable medium. A processor(s), implemented in an integrated circuit, may perform the necessary tasks.

[0324] Where components are described as being “configured to” perform certain operations, such configuration can be accomplished, for example, by designing electronic circuits or other hardware to perform the operation, by programming programmable electronic circuits (e.g., microprocessors, or other suitable electronic circuits) to perform the operation, or any combination thereof.

[0325] The various illustrative logical blocks, modules, circuits, and algorithm steps described in connection with the implementations disclosed herein may be implemented as electronic hardware, computer software, firmware, or combinations thereof. To clearly illustrate this interchangeability of hardware and software, various illustrative components, blocks, modules, circuits, and steps have been described above generally in terms of their

functionality. Whether such functionality is implemented as hardware or software depends upon the particular application and design constraints imposed on the overall system. Skilled artisans may implement the described functionality in varying ways for each particular application, but such implementation decisions should not be interpreted as causing a departure from the scope of the present disclosure.

[0326] The techniques described herein may also be implemented in electronic hardware, computer software, firmware, or any combination thereof. Such techniques may be implemented in any of a variety of devices such as general purposes computers, wireless communication device handsets, or integrated circuit devices having multiple uses including application in wireless communication device handsets and other devices. Any features described as modules or components may be implemented together in an integrated logic device or separately as discrete but interoperable logic devices. If implemented in software, the techniques may be realized at least in part by a computer-readable data storage medium comprising program code including instructions that, when executed, performs one or more of the methods described above. The computer-readable data storage medium may form part of a computer program product, which may include packaging materials. The computer-readable medium may comprise memory or data storage media, such as random access memory (RAM) such as synchronous dynamic random access memory (SDRAM), read-only memory (ROM), non-volatile random access memory (NVRAM), electrically erasable programmable read-only memory (EEPROM), FLASH memory, magnetic or optical data storage media, and the like. The techniques additionally, or alternatively, may be realized at least in part by a computer-readable communication medium that carries or communicates program code in the form of instructions or data structures and that can be accessed, read, and/or executed by a computer, such as propagated signals or waves.

[0327] The program code may be executed by a processor, which may include one or more processors, such as one or more digital signal processors (DSPs), general purpose microprocessors, an application specific integrated circuits (ASICs), field programmable logic arrays (FPGAs), or other equivalent integrated or discrete logic circuitry. Such a processor may be configured to perform any of the techniques described in this disclosure. A general purpose processor may be a microprocessor; but in the alternative, the processor may be any conventional processor, controller, microcontroller, or state machine. A processor may also be implemented as a combination of computing devices, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in

conjunction with a DSP core, or any other such configuration. Accordingly, the term “processor,” as used herein may refer to any of the foregoing structure, any combination of the foregoing structure, or any other structure or apparatus suitable for implementation of the techniques described herein. In addition, in some aspects, the functionality described herein may be provided within dedicated software modules or hardware modules configured for network threat detection and analysis.

[0328] As used below, any reference to a series of examples is to be understood as a reference to each of those examples disjunctively (e.g., “Examples 1-4” is to be understood as “Examples 1, 2, 3, or 4”).

[0329] Example 1 is a method comprising: identifying a network for which to deploy one or more deception mechanisms, wherein a deception mechanism emulates one or more characteristics of a machine on the network; determining one or more asset densities, wherein an asset density is associated with a number of assets connected to the network; determining a summary statistic associated with historical attacks on the network; determining a number of deception mechanisms to deploy, wherein the number of deception mechanisms is determined using at least one or more of the one or more asset densities or the summary statistic; and deploying the number of deception mechanisms.

[0330] Example 2 is the method of example 1, wherein the network is a portion of another network.

[0331] Example 3 is the method of example 2, wherein an asset density of the one or more asset densities is a total number of critical assets in the network divided by a total number of critical assets in the other network.

[0332] Example 4 is the method of examples 2-3, wherein an asset density of the one or more asset densities is a total number of assets in the network divided by a total number of assets in the other network.

[0333] Example 5 is the method of examples 1-4, wherein the network is identified by computing an importance score, and wherein the importance score is computed using the one or more asset densities and machine information associated with the network.

[0334] Example 6 is the method of examples 1-5, further comprising identifying a central node of the network, wherein a deception of the number of deception mechanisms is deployed on or near the central node.

[0335] Example 7 is the method of examples 1-6, wherein the summary statistic includes a mean, median, or mode of a probability distribution of the historical attacks.

[0336] Example 8 is the method of example 7, wherein the probability distribution of the number of historical attacks includes one or more probability distributions across one or more time periods.

[0337] Example 9 is the method of examples 1-8, wherein the summary statistic is updated when a request to access a deception mechanism is received, wherein updating the summary statistic includes: determining a new summary statistic using a number of new attacks associated with the network, wherein the new attacks include attacks that occurred after the historical attacks; and updating the summary statistic using the new summary statistic.

[0338] Example 10 is the method of examples 1-9, further comprising determining one or more configurations for one or more of the deployed deception mechanisms.

[0339] Example 11 is the method of example 10, wherein the one or more configurations are determined using a population distribution of machine information in the network, and wherein the population distribution identifies machine information associated with one or more machines in the network.

[0340] Example 12 is the method of examples 10-11, wherein the machine information includes at least one or more of a MAC address, a host name, an identifier of an operating system, or an identifier of a network service.

[0341] Example 13 is the method of examples 10-12, wherein a configuration includes at least one or more of a MAC address, a host name, an operating system, or a network service, and wherein the host name is determined by: determining a pattern for host names based on one or more machines in the network; and determining the host name according to the pattern.

[0342] Example 14 is the method of examples 10-13, further comprising: detecting a request to access a deception mechanism; determining a configuration of the accessed deception mechanism; and updating the one or more configurations for the one or more of the deployed deception mechanisms using the configuration of the accessed deception mechanism.

[0343] Example 15 is the method of examples 1-14, further comprising determining a schedule to deploy a deception mechanism.

[0344] Example 16 is the method of example 15, wherein the schedule is determined using the historical attacks on the network.

[0345] Example 17 is the method of examples 15-16, wherein the schedule includes one or more of a connect time or a disconnect time, wherein the connect time indicates when to connect a deception mechanism to the network, and wherein the disconnect time indicates when to disconnect a deception mechanism from the network.

[0346] Example 18 is the method of examples 1-17, wherein historical attacks are identified using at least one or more of a security operations center, a computer security incident response team, an intrusion detection system, or an intrusion prevention system.

[0347] Example 19 is a system comprising: one or more processors; and a non-transitory computer-readable medium containing instructions that, when executed by the one or more processors, cause the one or more processors to perform operations including: identify a network for which to deploy one or more deception mechanisms, wherein a deception mechanism emulates one or more characteristics of a machine on the network; determine one or more asset densities, wherein an asset density is associated with a number of assets connected to the network; determine a summary statistic associated with historical attacks on the network; determine a number of deception mechanisms to deploy, wherein the number of deception mechanisms is determined using at least one or more of the one or more asset densities or the summary statistic; and deploy the number of deception mechanisms.

[0348] Example 20 is the system of example 19, wherein the network is a portion of another network.

[0349] Example 21 is the system of example 20, wherein an asset density of the one or more asset densities is a total number of critical assets in the network divided by a total number of critical assets in the other network.

[0350] Example 22 is the system of examples 20-11, wherein an asset density of the one or more asset densities is a total number of assets in the network divided by a total number of assets in the other network.

[0351] Example 23 is the system of examples 19-22, wherein the network is identified by computing an importance score, and wherein the importance score is computed using the one or more asset densities and machine information associated with the network.

[0352] Example 24 is the system of examples 19-23, wherein the non-transitory computer-readable medium further includes instructions that, when executed by the one or more processors, cause the one or more processors to perform operations including to identify a central node of the network, wherein a deception of the number of deception mechanisms is deployed on or near the central node.

[0353] Example 25 is the system of examples 19-24, wherein the summary statistic includes a mean, median, or mode of a probability distribution of the historical attacks.

[0354] Example 26 is the system of example 25, wherein the probability distribution of the number of historical attacks includes one or more probability distributions across one or more time periods.

[0355] Example 27 is the system of examples 19-26, wherein the summary statistic is updated when a request to access a deception mechanism is received, wherein updating the summary statistic includes: determining a new summary statistic using a number of new attacks associated with the network, wherein the new attacks include attacks that occurred after the historical attacks; and updating the summary statistic using the new summary statistic.

[0356] Example 28 is the system of examples 19-27, wherein the non-transitory computer-readable medium further includes instructions that, when executed by the one or more processors, cause the one or more processors to perform operations including to determine one or more configurations for one or more of the deployed deception mechanisms.

[0357] Example 29 is the system of example 28, wherein the one or more configurations are determined using a population distribution of machine information in the network, and wherein the population distribution identifies machine information associated with one or more machines in the network.

[0358] Example 30 is the system of examples 28-29, wherein the machine information includes at least one or more of a MAC address, a host name, an identifier of an operating system, or an identifier of a network service.

[0359] Example 31 is the system of examples 28-30, wherein a configuration includes at least one or more of a MAC address, a host name, an operating system, or a network service, and wherein the host name is determined by: determining a pattern for host names based on one or more machines in the network; and determining the host name according to the pattern.

[0360] Example 32 is the system of examples 28-31, wherein the non-transitory computer-readable medium further includes instructions that, when executed by the one or more processors, cause the one or more processors to perform operations including to detect a request to access a deception mechanism; determining a configuration of the accessed deception mechanism; and updating the one or more configurations for the one or more of the deployed deception mechanisms using the configuration of the accessed deception mechanism.

[0361] Example 33 is the system of examples 19-32, wherein the non-transitory computer-readable medium further includes instructions that, when executed by the one or more processors, cause the one or more processors to perform operations including to determine a schedule to deploy a deception mechanism.

[0362] Example 34 is the system of example 33, wherein the schedule is determined using the historical attacks on the network.

[0363] Example 35 is the system of examples 33-34, wherein the schedule includes one or more of a connect time or a disconnect time, wherein the connect time indicates when to connect a deception mechanism to the network, and wherein the disconnect time indicates when to disconnect a deception mechanism from the network.

[0364] Example 36 is the system of examples 19-35, wherein historical attacks are identified using at least one or more of a security operations center, a computer security incident response team, an intrusion detection system, or an intrusion prevention system.

[0365] Example 37 is a computer-program product tangibly embodied in a non-transitory machine-readable storage medium of a network device, including instructions that, when executed by the one or more processors, cause the one or more processors to: identify a network for which to deploy one or more deception mechanisms, wherein a deception mechanism emulates one or more characteristics of a machine on the network; determine one or more asset densities, wherein an asset density is associated with a number of assets

connected to the network; determine a summary statistic associated with historical attacks on the network; determine a number of deception mechanisms to deploy, wherein the number of deception mechanisms is determined using at least one or more of the one or more asset densities or the summary statistic; and deploy the number of deception mechanisms.

[0366] Example 38 is the computer-program product of example 37, wherein the network is a portion of another network.

[0367] Example 39 is the computer-program product of example 38, wherein an asset density of the one or more asset densities is a total number of critical assets in the network divided by a total number of critical assets in the other network.

[0368] Example 40 is the computer-program product of examples 38-39, wherein an asset density of the one or more asset densities is a total number of assets in the network divided by a total number of assets in the other network.

[0369] Example 41 is the computer-program product of examples 37-40, wherein the network is identified by computing an importance score, and wherein the importance score is computed using the one or more asset densities and machine information associated with the network.

[0370] Example 42 is the computer-program product of examples 37-41, further including instructions that, when executed by the one or more processors, cause the one or more processors to identify a central node of the network, wherein a deception of the number of deception mechanisms is deployed on or near the central node.

[0371] Example 43 is the computer-program product of examples 37-42, wherein the summary statistic includes a mean, median, or mode of a probability distribution of the historical attacks.

[0372] Example 44 is the computer-program product of example 43, wherein the probability distribution of the number of historical attacks includes one or more probability distributions across one or more time periods.

[0373] Example 45 is the computer-program product of examples 37-44, wherein the summary statistic is updated when a request to access a deception mechanism is received, wherein updating the summary statistic includes: determining a new summary statistic using a number of new attacks associated with the network, wherein the new attacks include attacks

that occurred after the historical attacks; and updating the summary statistic using the new summary statistic.

[0374] Example 46 is the computer-program product of examples 37-45, further including instructions that, when executed by the one or more processors, cause the one or more processors to determine one or more configurations for one or more of the deployed deception mechanisms.

[0375] Example 47 is the computer-program product of example 46, wherein the one or more configurations are determined using a population distribution of machine information in the network, and wherein the population distribution identifies machine information associated with one or more machines in the network.

[0376] Example 48 is the computer-program product of examples 46-47, wherein the machine information includes at least one or more of a MAC address, a host name, an identifier of an operating system, or an identifier of a network service.

[0377] Example 49 is the computer-program product of examples 46-48, wherein a configuration includes at least one or more of a MAC address, a host name, an operating system, or a network service, and wherein the host name is determined by: determining a pattern for host names based on one or more machines in the network; and determining the host name according to the pattern.

[0378] Example 50 is the computer-program product of examples 46-49, further including instructions that, when executed by the one or more processors, cause the one or more processors to: detect a request to access a deception mechanism; determine a configuration of the accessed deception mechanism; and update the one or more configurations for the one or more of the deployed deception mechanisms using the configuration of the accessed deception mechanism.

[0379] Example 51 is the computer-program product of examples 37-50, further including instructions that, when executed by the one or more processors, cause the one or more processors to determine a schedule to deploy a deception mechanism.

[0380] Example 52 is the computer-program product of example 51, wherein the schedule is determined using the historical attacks on the network.

[0381] Example 53 is the computer-program product of examples 51-52, wherein the schedule includes one or more of a connect time or a disconnect time, wherein the connect time indicates when to connect a deception mechanism to the network, and wherein the disconnect time indicates when to disconnect a deception mechanism from the network.

[0382] Example 54 is the computer-program product of examples 37-53, wherein historical attacks are identified using at least one or more of a security operations center, a computer security incident response team, an intrusion detection system, or an intrusion prevention system.

WHAT IS CLAIMED IS:

1. A method comprising:
 - identifying a network for which to deploy one or more deception mechanisms, wherein a deception mechanism emulates one or more characteristics of a machine on the network;
 - determining one or more asset densities, wherein an asset density is associated with a number of assets connected to the network;
 - determining a summary statistic associated with historical attacks on the network;
 - determining a number of deception mechanisms to deploy, wherein the number of deception mechanisms is determined using at least one or more of the one or more asset densities or the summary statistic; and
 - deploying the number of deception mechanisms.
2. The method of claim 1, wherein the network is a portion of another network.
3. The method of claim 2, wherein an asset density of the one or more asset densities is a total number of critical assets in the network divided by a total number of critical assets in the other network.
4. The method of claim 2, wherein an asset density of the one or more asset densities is a total number of assets in the network divided by a total number of assets in the other network.
5. The method of claim 1, wherein the network is identified by computing an importance score, and wherein the importance score is computed using the one or more asset densities and machine information associated with the network.
6. The method of claim 1, further comprising:
 - identifying a central node of the network, wherein a deception of the number of deception mechanisms is deployed on or near the central node.
7. The method of claim 1, wherein the summary statistic includes a mean, median, or mode of a probability distribution of the historical attacks.

8. The method of claim 7, wherein the probability distribution of the number of historical attacks includes one or more probability distributions across one or more time periods.

9. The method of claim 1, wherein the summary statistic is updated when a request to access a deception mechanism is received, wherein updating the summary statistic includes:

determining a new summary statistic using a number of new attacks associated with the network, wherein the new attacks include attacks that occurred after the historical attacks; and

updating the summary statistic using the new summary statistic.

10. The method of claim 1, further comprising:

determining one or more configurations for one or more of the deployed deception mechanisms.

11. The method of claim 10, wherein the one or more configurations are determined using a population distribution of machine information in the network, and wherein the population distribution identifies machine information associated with one or more machines in the network.

12. The method of claim 10, wherein the machine information includes at least one or more of a MAC address, a host name, an identifier of an operating system, or an identifier of a network service.

13. The method of claim 10, wherein a configuration includes at least one or more of a MAC address, a host name, an operating system, or a network service, and wherein the host name is determined by:

determining a pattern for host names based on one or more machines in the network; and

determining the host name according to the pattern.

14. The method of claim 10, further comprising:

detecting a request to access a deception mechanism;

determining a configuration of the accessed deception mechanism; and

updating the one or more configurations for the one or more of the deployed deception mechanisms using the configuration of the accessed deception mechanism.

15. The method of claim 1, further comprising:
determining a schedule to deploy a deception mechanism.

16. The method of claim 15, wherein the schedule is determined using the historical attacks on the network.

17. The method of claim 15, wherein the schedule includes one or more of a connect time or a disconnect time, wherein the connect time indicates when to connect a deception mechanism to the network, and wherein the disconnect time indicates when to disconnect a deception mechanism from the network.

18. The method of claim 1, wherein historical attacks are identified using at least one or more of a security operations center, a computer security incident response team, an intrusion detection system, or an intrusion prevention system.

19. A system comprising:
one or more processors; and
a non-transitory computer-readable medium containing instructions that, when executed by the one or more processors, cause the one or more processors to perform operations including:

identify a network for which to deploy one or more deception mechanisms, wherein a deception mechanism emulates one or more characteristics of a machine on the network;

determine one or more asset densities, wherein an asset density is associated with a number of assets connected to the network;

determine a summary statistic associated with historical attacks on the network;

determine a number of deception mechanisms to deploy, wherein the number of deception mechanisms is determined using at least one or more of the one or more asset densities or the summary statistic; and

deploy the number of deception mechanisms.

20. A computer-program product tangibly embodied in a non-transitory machine-readable storage medium of a network device, including instructions that, when executed by the one or more processors, cause the one or more processors to:

identify a network for which to deploy one or more deception mechanisms, wherein a deception mechanism emulates one or more characteristics of a machine on the network;

determine one or more asset densities, wherein an asset density is associated with a number of assets connected to the network;

determine a summary statistic associated with historical attacks on the network;

determine a number of deception mechanisms to deploy, wherein the number of deception mechanisms is determined using at least one or more of the one or more asset densities or the summary statistic; and

deploy the number of deception mechanisms.




FIG. 1




FIG. 2A




FIG. 2B

200c ↗




FIG. 2C




FIG. 2D




FIG. 3A




FIG. 3B




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 9

900 ↗

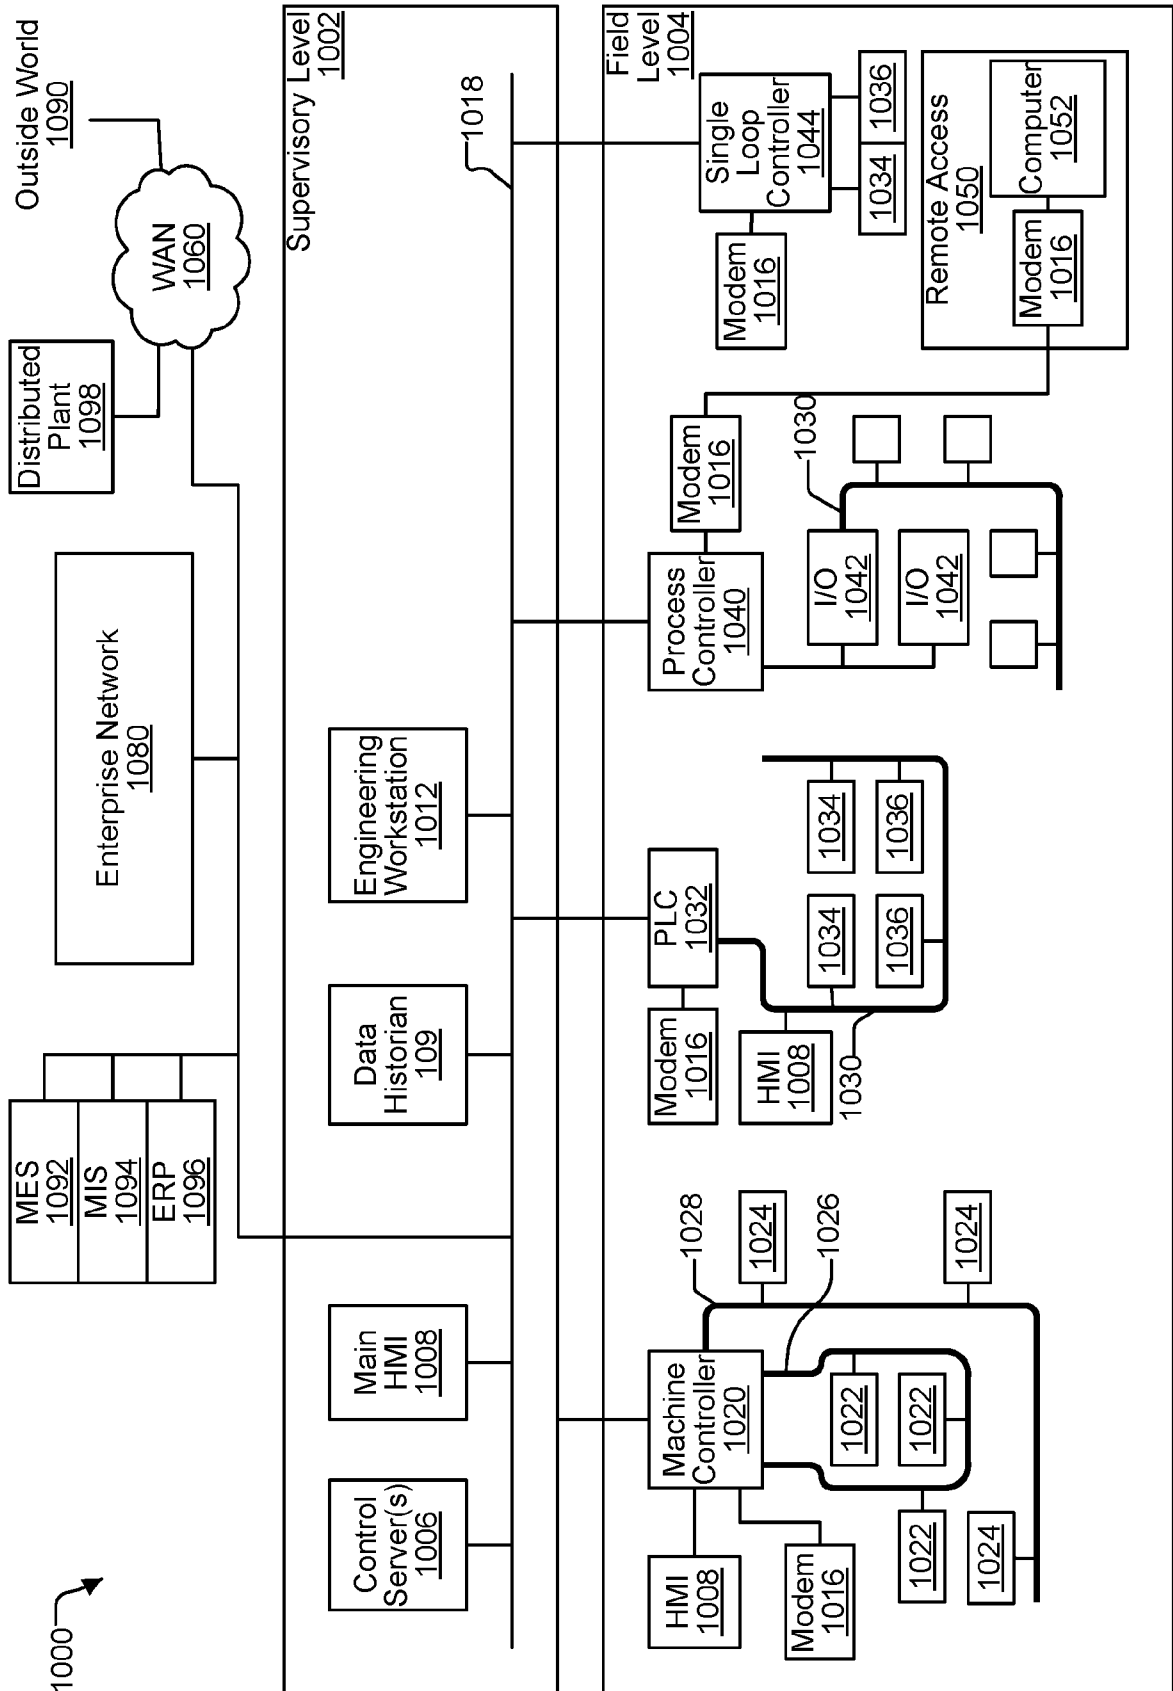


FIG. 10




FIG. 11




FIG. 12

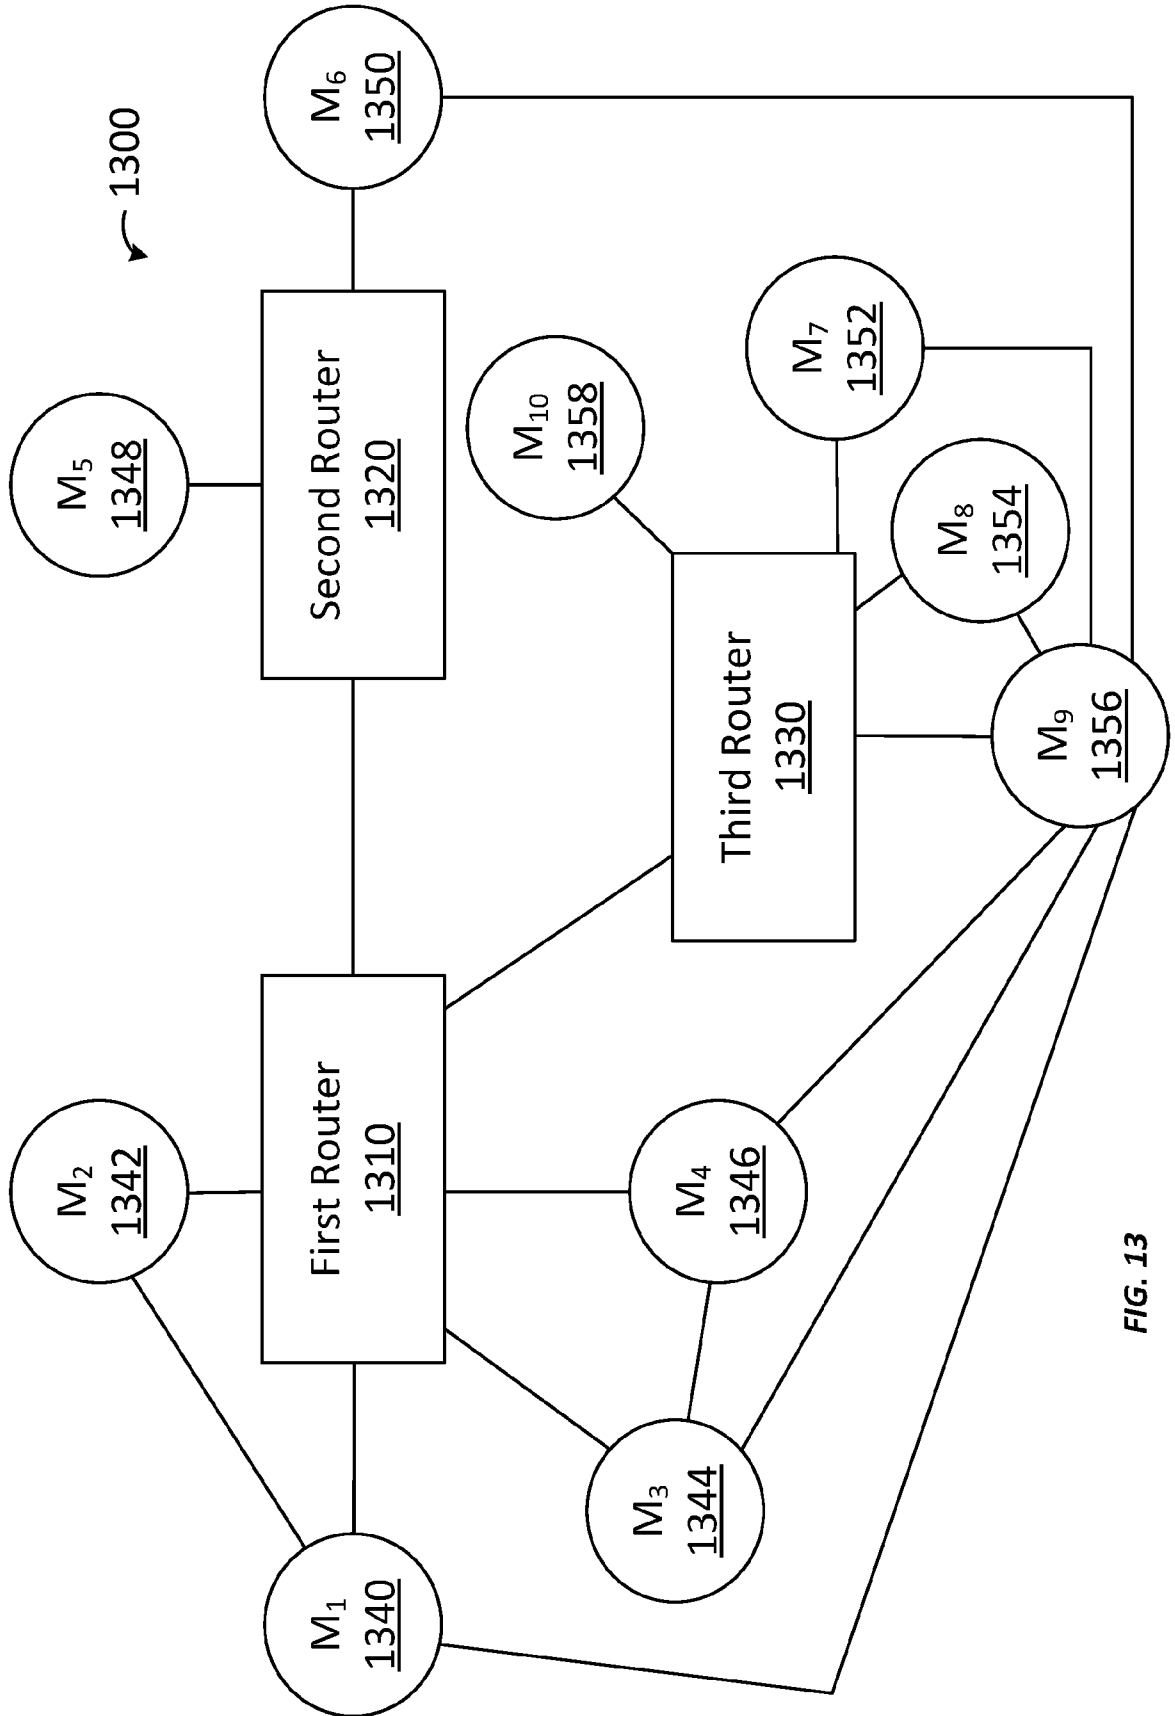


FIG. 13




FIG. 14B




FIG. 14C

1500

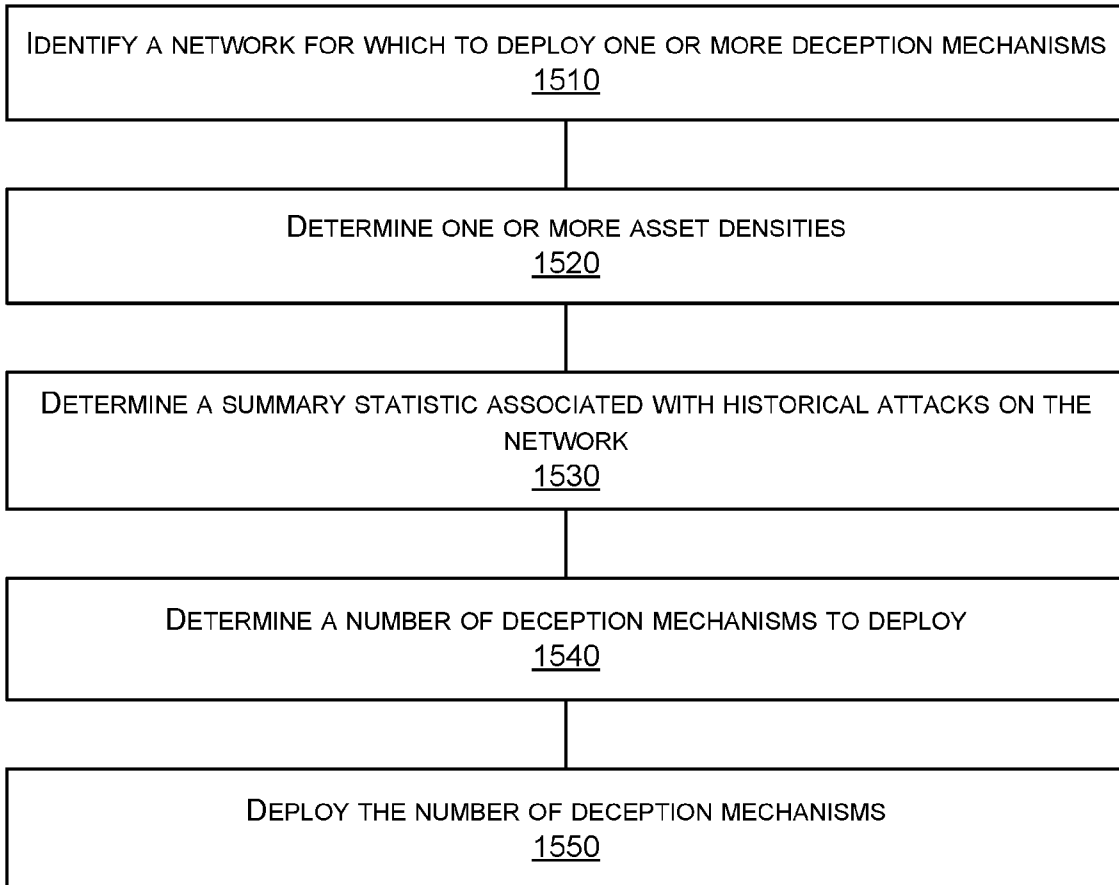


FIG. 15

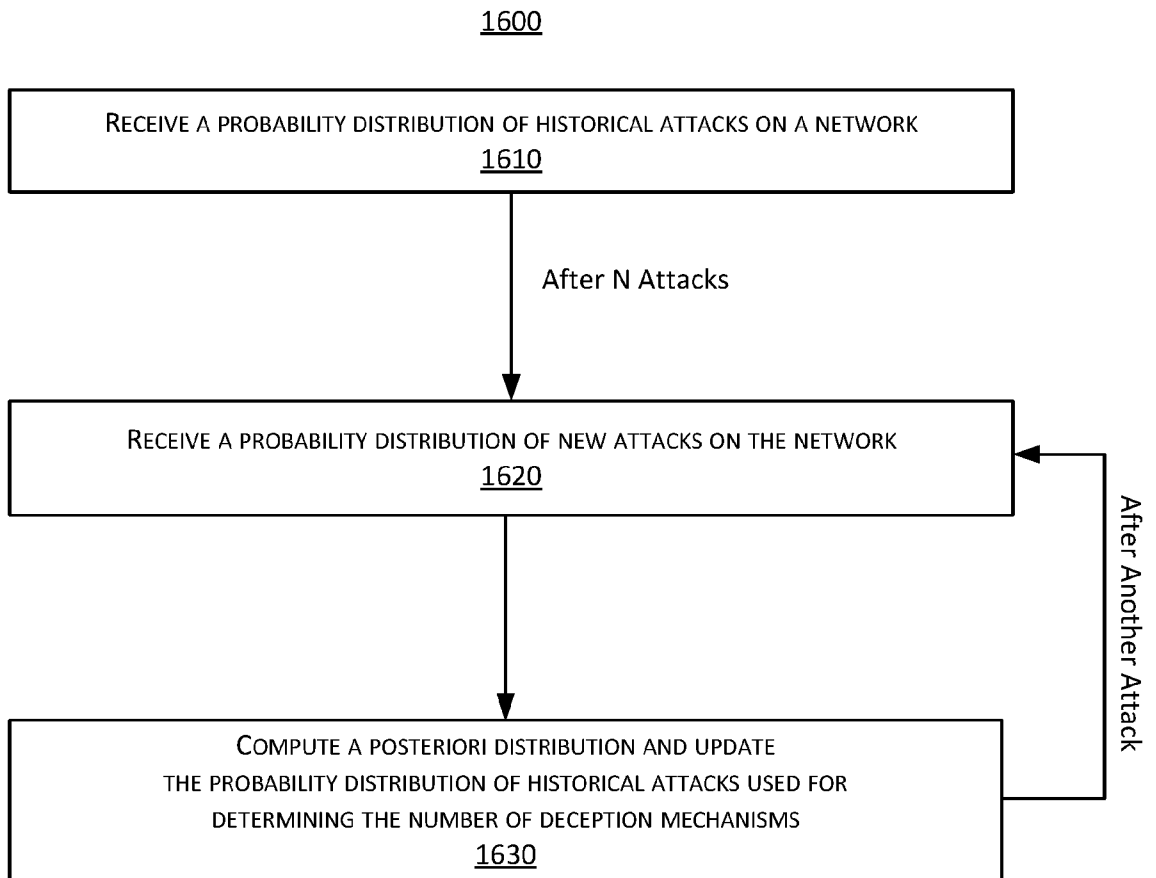


FIG. 16

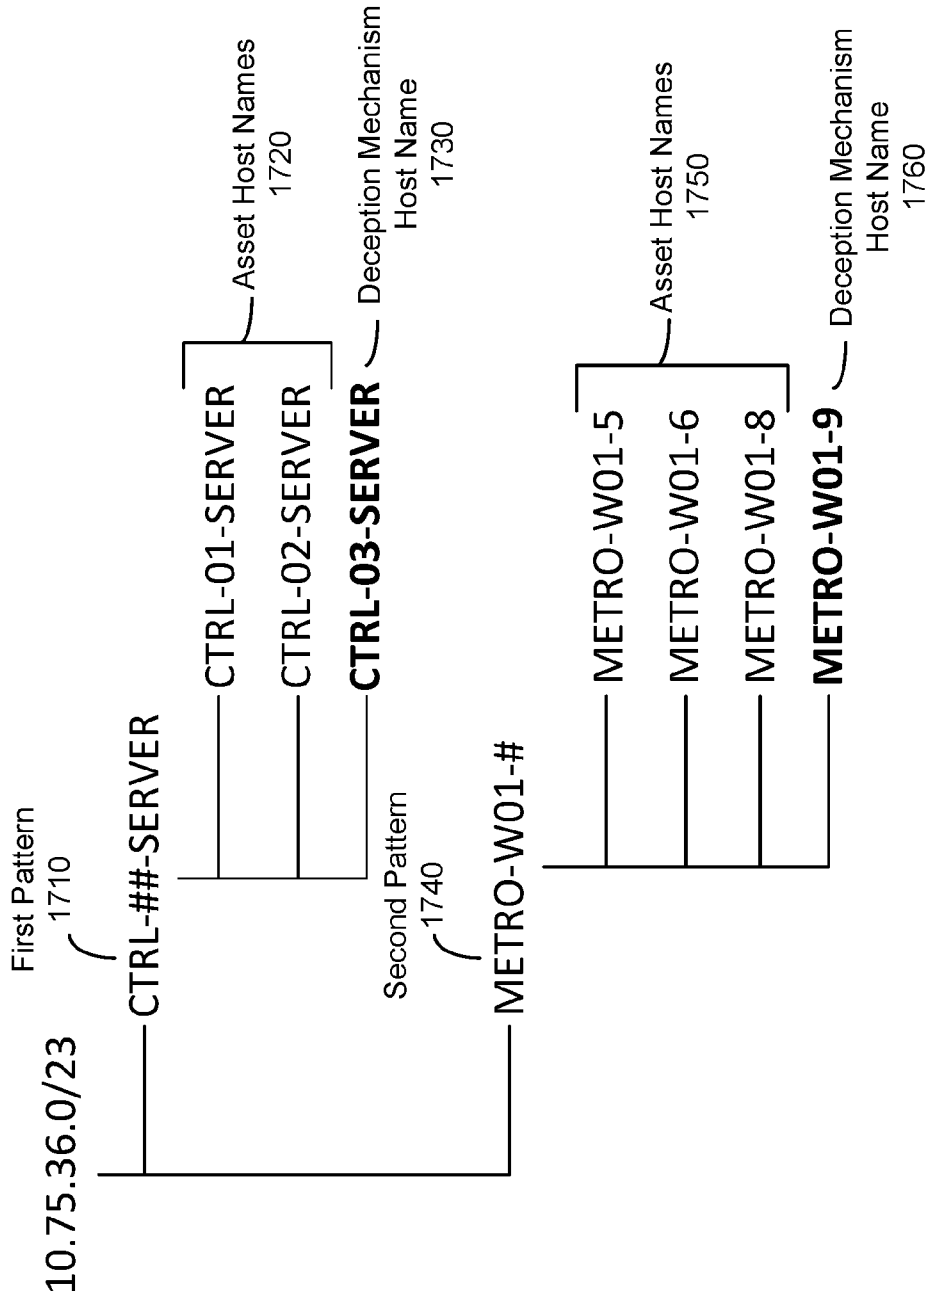


FIG. 17

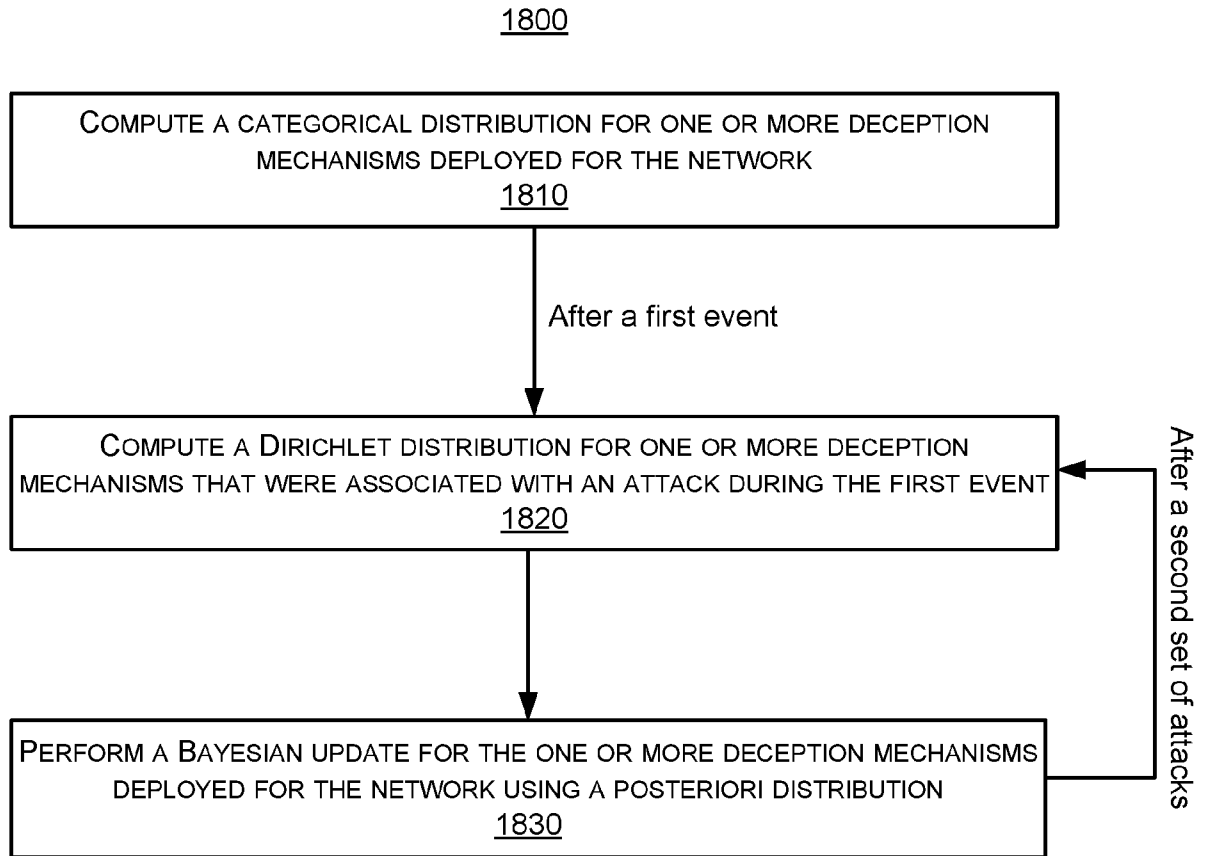


FIG. 18




FIG. 19A




FIG. 19B


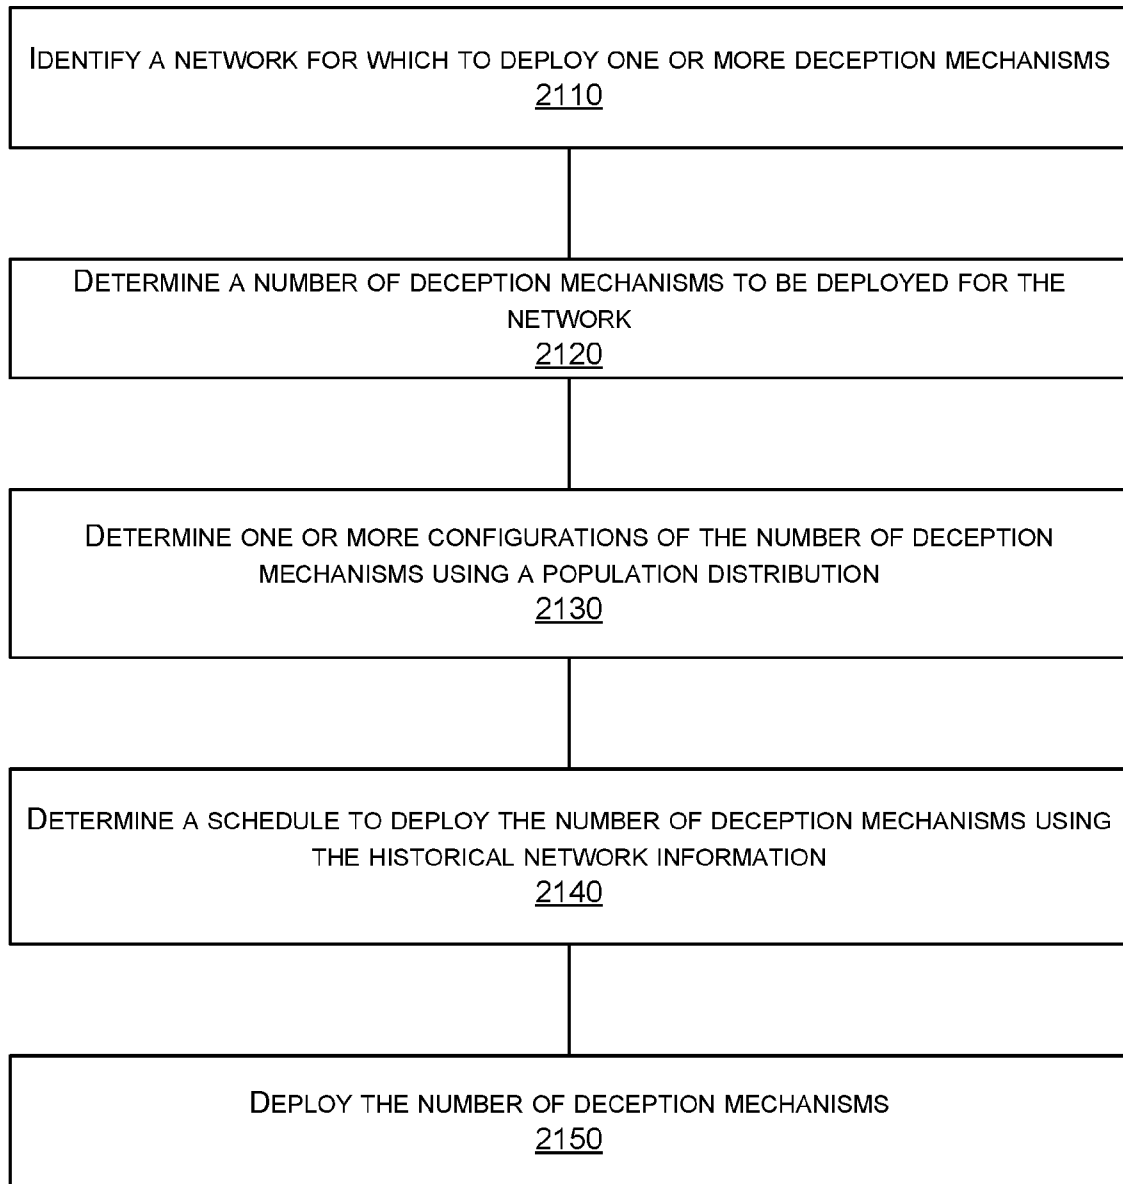


FIG. 20

27/27

2100**FIG. 21**

INTERNATIONAL SEARCH REPORT

International application No PCT/US2017/016535

A. CLASSIFICATION OF SUBJECT MATTER INV. H04L29/06 ADD.				
According to International Patent Classification (IPC) or to both national classification and IPC				
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) H04L				
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched				
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) EPO-Internal, WPI Data				
C. DOCUMENTS CONSIDERED TO BE RELEVANT				
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.		
X	US 2014/250524 A1 (MEYERS ADAM S [US] ET AL) 4 September 2014 (2014-09-04) figure 1 paragraph [0012] paragraph [0016] paragraph [0018] paragraph [0019] paragraph [0028] paragraph [0030] paragraph [0033] paragraph [0049] paragraph [0050]	1-20		
A	US 2004/128529 A1 (BLAKE KENNETH W [US] ET AL) 1 July 2004 (2004-07-01) paragraph [0052]	1-20		
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; border: none;"><input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C.</td> <td style="width: 50%; border: none;"><input checked="" type="checkbox"/> See patent family annex.</td> </tr> </table>			<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C.	<input checked="" type="checkbox"/> See patent family annex.
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C.	<input checked="" type="checkbox"/> See patent family annex.			
* Special categories of cited documents :				
"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family			
Date of the actual completion of the international search	Date of mailing of the international search report			
27 April 2017	11/05/2017			
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Caragata, Daniel			

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/016535

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2010/030169 A2 (MIMOS BHD [MY]; WIRA ZANORAMY ANSIRY ZAKARIA [MY]; SITI ROHAIDAH AHMAD) 18 March 2010 (2010-03-18) page 3, line 3 - line 14 page 3, line 26 - line 30 -----	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2017/016535

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014250524 A1	04-09-2014	AU 2014226405 A1 CA 2901302 A1 EP 2965256 A1 JP 2016514313 A SG 112015067190 A US 2014250524 A1 WO 2014137640 A1	10-09-2015 12-09-2014 13-01-2016 19-05-2016 29-09-2015 04-09-2014 12-09-2014

US 2004128529 A1	01-07-2004	US 2004128529 A1 US 2009328213 A1	01-07-2004 31-12-2009

WO 2010030169 A2	18-03-2010	CN 102216900 A EP 2327014 A2 MY 146995 A WO 2010030169 A2	12-10-2011 01-06-2011 15-10-2012 18-03-2010
