



US 20080175266A1

(19) **United States**

(12) **Patent Application Publication**
Alperovitch et al.

(10) **Pub. No.: US 2008/0175266 A1**

(43) **Pub. Date: Jul. 24, 2008**

(54) **MULTI-DIMENSIONAL REPUTATION
SCORING**

(75) **Inventors:** **Dmitri Alperovitch**, Atlanta, GA (US); **Tomo Foote-Lennox**, Maple Grove, MN (US); **Jeremy Gould**, Atlanta, GA (US); **Paula Greve**, Lino Lakes, MN (US); **Alejandro Manuel Hernandez**, Cartersville, GA (US); **Paul Judge**, Atlanta, GA (US); **Sven Krasser**, Atlanta, GA (US); **Tim Lange**, Bielefeld (DE); **Phyllis Adele Schneck**, Atlanta, GA (US); **Martin Stecher**, Paderborn (DE); **Yuchun Tang**, Roswell, GA (US); **Aarjav Jyotindra Neeta Trivedi**, Atlanta, GA (US); **Lamar Lorenzo Willis**, Woodstock, GA (US); **Weilai Yang**, Alpharetta, GA (US); **Jonathan Alexander Zdziarski**, Milledgeville, GA (US)

Correspondence Address:
FISH & RICHARDSON P.C.
P.O BOX 1022
Minneapolis, MN 55440-1022

(73) **Assignee:** **SECURE COMPUTING CORPORATION**, Saint Paul, MN (US)

(21) **Appl. No.:** **11/626,603**

(22) **Filed:** **Jan. 24, 2007**


Publication Classification

(51) **Int. Cl.**
H04J 3/22 (2006.01)

(52) **U.S. Cl.** **370/465; 370/252**

(57) **ABSTRACT**

Methods and systems for assigning reputation to communications entities include collecting communications data from distributed agents, aggregating the communications data, analyzing the communications data and identifying relationships between communications entities based upon the communications data.






FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10




FIG. 11A




FIG. 11B




FIG. 11C




FIG. 12





FIG. 13

**FIG. 14**




FIG. 15D




FIG. 15C




FIG. 15B




FIG. 15A

MULTI-DIMENSIONAL REPUTATION SCORING

CROSS-REFERENCE

[0001] This application incorporates by reference, in their entirety and for all purposes, commonly assigned U.S. Patent Applications:

Application No.	Title	Filing Date
10/094,211	"Systems and Methods for Enhancing Electronic Communication Security"	Mar. 8, 2002
10/361,067	"Systems and Methods for Automated Whitelisting in Monitored Communications"	Feb. 7, 2003
10/373,325	"Systems and Methods for Upstream Threat Pushback"	Feb. 24, 2003
10/384,924	"Systems and Methods for Secure Communication Delivery"	Mar. 6, 2003
11/173,941	"Message Profiling Systems and Methods"	Jun. 2, 2005
11/142,943	"Systems and Methods for Classification of Messaging Entities"	Jun. 2, 2005
11/388,575	"Systems and Methods for Message Threat Management"	Mar. 24, 2006
11/456,803	"Systems And Methods For Adaptive Message Interrogation Through Multiple Queues"	Jul. 11, 2006
11/456,765	"Systems and Methods For Anomaly Detection in Patterns of Monitored Communications"	Jul. 11, 2006
11/423,313	"Systems and Methods for Identifying Potentially Malicious Messages"	Jun. 9, 2006
11/456,954	"Systems and Methods For Message Threat Management"	Jul. 12, 2006
11/456,960	"Systems and Methods For Message Threat Management"	Jul. 12, 2006
11/423,308	"Systems and Methods for Graphically Displaying Messaging Traffic"	Jun. 9, 2006
11/383,347	"Content-Based Policy Compliance Systems and Methods"	May 15, 2006
11/423,329	"Methods and Systems for Exposing Messaging Reputation to an End User"	Jun. 9, 2006

[0002] This application incorporates by reference, in their entirety and for all purposes, commonly assigned U.S. Patents:

U.S. Pat. No.	Title	Filing Date
6,941,467	"Systems and Methods for Adaptive Message Interrogation through Multiple Queues"	Mar. 8, 2002
7,089,590	"Systems and Methods for Adaptive Message Interrogation through Multiple Queues"	Sep. 2, 2005
7,096,498	"Systems and Methods for Message Threat Management"	Feb. 7, 2003
7,124,438	"Systems and Methods for Anomaly Detection in Patterns of Monitored Communications"	Mar. 8, 2002

TECHNICAL FIELD

[0003] This document relates generally to systems and methods for processing communications and more particularly to systems and methods for classifying entities associated with communications.

BACKGROUND

[0004] In the anti-spam industry, spammers use various creative means for evading detection by spam filters. As such,

the entity from which a communication originated can provide another indication of whether a given communication should be allowed into an enterprise network environment.

[0005] However, current tools for message sender analysis include internet protocol (IP) blacklists (sometimes called real-time blacklists (RBLs)) and IP whitelists (real-time whitelists (RWLs)). Whitelists and blacklists certainly add

value to the spam classification process; however, whitelists and blacklists are inherently limited to providing a binary-type (YES/NO) response to each query. Moreover, blacklists and whitelists treat entities independently, and overlook the evidence provided by various attributes associated with the entities.

SUMMARY

[0006] Systems and methods for a distributed reputation architecture are provided. A distributed reputation system can include a communications interface, a data aggregation engine, an analyzer, a correlation engine and a reputation engine. The communications interface can communicate with a plurality of agents arranged within a global network. The data aggregation engine can aggregate the collected data via the communications interface. The analyzer can analyze the data to identify attributes respectively associated with entities originating the received communications. The correlation engine can correlate the attributes of the entities and to identify relationships between the entities. The reputation engine can identify relationships between the entities and to update reputations associated with one or more entities based upon its relationship to one or more other entities. The communications interface can also to communicate updated reputation information to devices operating on the global network.

[0007] Other systems operable to derive and assign reputations can include a communications interface, a data aggregation engine, an analyzer, a correlation engine, a reputation engine and a traffic control engine. The communications interface can receive information from agents or a central server within a global network. The data aggregation engine can aggregate the received information from the communications interface. The analyzer can analyze the received information to identify attributes respectively associated with entities originating the received communications. The correlation engine can correlate the attributes of the entities and identify relationships between the entities. The reputation engine can identify relationships between the entities and to update reputations associated with one or more entities based upon its relationship to one or more other entities. The traffic control engine can determine handling associated with a communication based upon the updated reputations.

[0008] Methods of assigning reputation to communications entities can include: arranging a plurality of agents within a network, the plurality of agents being associated with a security device operable to protect an associated network from communications that violate a policy associated with the associated network; collecting data associated with entities originating communications, wherein collecting data comprises using the plurality of agents to collect data associated with the communications; aggregating the collected data; analyzing the aggregated data to identify attributes respectively associated with entities originating communications; correlating the attributes to identify relationships between entities; updating a reputation associated with one or more entities based upon the relationship to one or more other entities identified by correlating the attributes; and, communicating updated reputation information to one or more of the plurality of agents.

[0009] Methods of assigning reputation to a communications entity can include: collecting data associated with entities originating communications, wherein collecting data comprises receiving data from a plurality of agents to collect data associated with the communications; aggregating the collected data; analyzing the aggregated data to identify attributes respectively associated with entities originating communications; correlating the attributes to identify relationships between entities; updating a reputation associated with one or more entities based upon the relationship to one or more other entities identified by correlating the attributes; and, handling, communications based upon the updated reputation information.

DESCRIPTION OF DRAWINGS

[0010] FIG. 1 is a block diagram depicting an example network in which systems and methods of this disclosure can operate.

[0011] FIG. 2 is a block diagram depicting an example network architecture of this disclosure.

[0012] FIG. 3 is a block diagram depicting an example of communications and entities including identifiers and attributes used to detect relationships between entities.

[0013] FIG. 4 is a flowchart depicting an operational scenario used to detect relationships and assign risk to entities.

[0014] FIG. 5 is a block diagram illustrating an example network architecture including local reputations stored by local security agents and a global reputation stored by one or more servers.

[0015] FIG. 6 is a block diagram illustrating a determination of a global reputation based on local reputation feedback.

[0016] FIG. 7 is a flow diagram illustrating an example resolution between a global reputation and a local reputation.

[0017] FIG. 8 is an example graphical user interface for adjusting the settings of a filter associated with a reputation server.

[0018] FIG. 9 is a block diagram illustrating reputation based connection throttling for voice over internet protocol (VoIP) or short message service (SMS) communications.

[0019] FIG. 10 is a block diagram illustrating a reputation based load balancer.

[0020] FIG. 11A is a flowchart illustrating an example operational scenario for geolocation based authentication.

[0021] FIG. 11B is a flowchart illustrating another example operational scenario for geolocation based authentication.

[0022] FIG. 11C is a flowchart illustrating another example operational scenario for geolocation based authentication.

[0023] FIG. 12 is a flowchart illustrating an example operational scenario for a reputation based dynamic quarantine.

[0024] FIG. 13 is an example graphical user interface display of an image spam communication.

[0025] FIG. 14 is a flowchart illustrating an example operational scenario for detecting image FIG. 15A is a flowchart illustrating an operational scenario for analyzing the structure of a communication.

[0026] FIG. 15B is a flowchart illustrating an operational scenario for analyzing the features of an image.

[0027] FIG. 15C is a flowchart illustrating an operational scenario for normalizing the an image for spam processing.

[0028] FIG. 15D is a flowchart illustrating an operational scenario for analyzing the fingerprint of an image to find common fragments among multiple images.

DETAILED DESCRIPTION

[0029] FIG. 1 is a block diagram depicting an example network environment in which systems and methods of this disclosure can operate. Security agent 1100 can typically reside between a firewall system (not shown) and servers (not shown) internal to a network 110 (e.g., an enterprise network). As should be understood, the network 110 can include a number of servers, including, for example, electronic mail servers, web servers, and various application servers as may be used by the enterprise associated with the network 110.

[0030] The security agent 100 monitors communications entering and exiting the network 110. These communications are typically received though the internet 120 from many entities 130a-f that are connected to the internet 120. One or more of the entities 130a-f can be legitimate originators of communications traffic. However, one or more of the entities 130a-f can also be non-reputable entities originating unwanted communications. As such, the security agent 100 includes a reputation engine. The reputation engine can inspect a communication and to determine a reputation associated with an entity that originated the communication. The security agent 100 then performs an action on the communication based upon the reputation of the originating entity. If the reputation indicates that the originator of the communication is reputable, for example, the security agent can forward the communication to the recipient of the communication. However, if the reputation indicates that the originator of the communication is non-reputable, for example, the security agent can quarantine the communication, perform more tests on the message, or require authentication from the mes-

sage originator, among many others. Reputation engines are described in detail in United States Patent Publication No. 2006/0015942, which is hereby incorporated by reference.

[0031] FIG. 2 is a block diagram depicting an example network architecture of this disclosure. Security agents **100a-n** are shown logically residing between networks **110a-n**, respectively, and the internet **120**. While not shown in FIG. 2, it should be understood that a firewall may be installed between the security agents **100a-n** and the internet **120** to provide protection from unauthorized communications from entering the respective networks **110a-n**. Moreover, intrusion detection systems (IDS) (not shown) can be deployed in conjunction with firewall systems to identify suspicious patterns of activity and to signal alerts when such activity is identified.

[0032] While such systems provide some protection for a network they typically do not address application level security threats. For example, hackers often attempt to use various network-type applications (e.g., e-mail, web, instant messaging (IM) etc.) to create a pre-textual connection with the networks **110a-n** in order to exploit security holes created by these various applications using entities **130a-e**. However, not all entities **130a-e** imply threats to the network **110a-n**. Some entities **130a-e** originate legitimate traffic, allowing the employees of a company to communicate with business associates more efficiently. While examining the communications for potential threats is useful, it can be difficult to maintain current threat information because attacks are being continually modified to account for the latest filtering techniques. Thus, security agents **100a-n** can run multiple tests on a communication to determine whether the communication is legitimate.

[0033] Furthermore, sender information included in the communication can be used to help determine whether or not a communication is legitimate. As such, sophisticated security agents **100a-n** can track entities and analyze the characteristics of the entities to help determine whether to allow a communication to enter a network **110a-n**. The entities **110a-n** can then be assigned a reputation. Decisions on a communication can take into account the reputation of an entity **130a-e** that originated the communication. Moreover, one or more central systems **200** can collect information on entities **120a-e** and distribute the collected data to other central systems **200** and/or the security agents **100a-n**.

[0034] Reputation engines can assist in identifying the bulk of the malicious communications without extensive and potentially costly local analysis of the content of the communication. Reputation engines can also help to identify legitimate communications and prioritize their delivery and reduce the risk of misclassifying a legitimate communication. Moreover, reputation engines can provide a dynamic and predictive approaches to the problem of identifying malicious, as well as legitimate, transactions in physical or virtual worlds. Examples include the process of filtering malicious communications in an email, instant messaging, VoIP, SMS or other communication protocol system using analysis of the reputation of sender and content. A security agent **100a-n** can then apply a global or local policy to determine what action to perform with respect to the communication (such as deny, quarantine, load balance, deliver with assigned priority, analyze locally with additional scrutiny) to the reputation result.

[0035] However, the entities **130a-e** can connect to the internet in a variety of methods. As should be understood, an entity **130a-e** can have multiple identifiers (such as, for example, e-mail addresses, IP addresses, identifier documen-

tation, etc) at the same time or over a period of time. For example, a mail server with changing IP addresses can have multiple identities over time. Moreover, one identifier can be associated with multiple entities, such as, for example, when an IP address is shared by an organization with many users behind it. Moreover, the specific method used to connect to the internet can obscure the identification of the entity **130a-e**. For example, an entity **130b** may connect to the internet using an internet service provider (ISP) **200**. Many ISPs **200** use dynamic host configuration protocol (DHCP) to assign IP addresses dynamically to entities **130b** requesting a connection. Entities **130a-e** can also disguise their identity by spoofing a legitimate entity. Thus, collecting data on the characteristics of each entity **130a-e** can help to categorize an entity **130a-e** and determine how to handle a communication.

[0036] The ease of creation and spoofing of identities in both virtual and physical world can create an incentive for users to act maliciously without bearing the consequences of that act. For example, a stolen IP address on the Internet (or a stolen passport in the physical world) of a legitimate entity by a criminal can enable that criminal to participate in malicious activity with relative ease by assuming the stolen identity. However, by assigning a reputation to the physical and virtual entities and recognizing the multiple identities that they can employ, reputation systems can influence reputable and non-reputable entities to operate responsibly for fear of becoming non-reputable, and being unable to correspond or interact with other network entities.

[0037] FIG. 3 is a block diagram depicting an example of communications and entities including using identifiers and attributes used to detect relationships between entities. Security agents **100a-b** can collect data by examining communications that are directed to an associated network. Security agents **100a-b** can also collect data by examining communications that are relayed by an associated network. Examination and analysis of communications can allow the security agents **100a-b** to collect information about the entities **300a-c** sending and receiving messages, including, transmission patterns, volume, or whether the entity has a tendency to send certain kinds of message (e.g., legitimate messages, spam, virus, bulk mail, etc.), among many others.

[0038] As shown in FIG. 3, each of the entities **300a-c** is associated with one or more identifiers **310a-c**, respectively. The identifiers **310a-c** can include, for example, IP addresses, universal resource locator (URL), phone number, IM username, message content, domain, or any other identifier that might describe an entity. Moreover, the identifiers **310a-c** are associated with one or more attributes **320a-c**. As should be understood, the attributes **320a-c** are fitted to the particular identifier **310a-c** that is being described. For example, a message content identifier could include attributes such as, for example, malware, volume, type of content, behavior, etc. Similarly, attributes **320a-c** associated with an identifier, such as IP address, could include one or more IP addresses associated with an entity **300a-c**.

[0039] Furthermore, it should be understood that this data can be collected from communications **330a-c** (e.g., e-mail) typically include some identifiers and attributes of the entity that originated the communication. Thus, the communications **330a-c** provide a transport for communicating information about the entity to the security agents **100a**, **100b**. These attributes can be detected by the security agents **100a**, **100b** through examination of the header information included in the message, analysis of the content of the message, as well as

through aggregation of information previously collected by the security agents 100a, 100b (e.g., totaling the volume of communication is received from an entity).

[0040] The data from multiple security agents 100a, 100b can be aggregated and mined. For example, the data can be aggregated and mined by a central system which receives identifiers and attributes associated with all entities 300a-c for which the security agents 100a, 100b have received communications. Alternatively, the security agents 100a, 100b can operate as a distributed system, communicating identifier and attribute information about entities 300a-c with each other. The process of mining the data can correlate the attributes of entities 300a-c with each other, thereby determining relationships between entities 300a-c (such as, for example, correlations between an event occurrence, volume, and/or other determining factors).

[0041] These relationships can then be used to establish a multi-dimensional reputation "vector" for all identifiers based on the correlation of attributes that have been associated with each identifier. For example, if a non-reputable entity 300a with a known reputation for being non-reputable sends a message 330a with a first set of attributes 350a, and then an unknown entity 300b sends a message 330b with a second set of attributes 350b, the security agent 100a can determine whether all or a portion of the first set of attributes 350a matched all or a portion of the second set of attributes 350b. When some portion of the first set of attributes 350a matches some portion of the second set of attributes 330b, a relationship can be created depending upon the particular identifier 320a, 320b that included the matching attributes 330a, 330b. The particular identifiers 340a, 340b which are found to have matching attributes can be used to determine a strength associated with the relationship between the entities 300a, 300b. The strength of the relationship can help to determine how much of the non-reputable qualities of the non-reputable entity 300a are attributed to the reputation of the unknown entity 300b.

[0042] However, it should also be recognized that the unknown entity 300b may originate a communication 330c which includes attributes 350c that match some attributes 350d of a communication 330d originating from a known reputable entity 300c. The particular identifiers 340c, 340d which are found to have matching attributes can be used to determine a strength associated with the relationship between the entities 300b, 300c. The strength of the relationship can help to determine how much of the reputable qualities of reputable entity 300c are attributed to the reputation of the unknown entity 300b.

[0043] A distributed reputation engine also allows for real-time collaborative sharing of global intelligence about the latest threat landscape, providing instant protection benefits to the local analysis that can be performed by a filtering or risk analysis system, as well as identify malicious sources of potential new threats before they even occur. Using sensors positioned at many different geographical locations information about new threats can be quickly and shared with the central system 200, or with the distributed security agents 100a, 100b. As should be understood, such distributed sensors can include the local security agents 100a, 100b, as well as local reputation clients, traffic monitors, or any other device suitable for collecting communication data (e.g., switches, routers, servers, etc.).

[0044] For example, security agents 100a, 100b can communicate with a central system 200 to provide sharing of

threat and reputation information. Alternatively, the security agents 100a, 100b can communicate threat and reputation information between each other to provide up to date and accurate threat information. In the example of FIG. 3, the first security agent 100a has information about the relationship between the unknown entity 300b and the non-reputable entity 300a, while the second security agent 100b has information about the relationship between the unknown entity 300b and the reputable entity 300c. Without sharing the information, the first security agent 100a may take a particular action on the communication based upon the detected relationship. However, with the knowledge of the relationship between the unknown entity 300b and the reputable entity 300c, the first security agent 100a might take a different action with a received communication from the unknown entity 300b. Sharing of the relationship information between security agents, thus provides for a more complete set of relationship information upon which a determination will be made.

[0045] The system attempts to assign reputations (reflecting a general disposition and/or categorization) to physical entities, such as individuals or automated systems performing transactions. In the virtual world, entities are represented by identifiers (ex. IPs, URLs, content) that are tied to those entities in the specific transactions (such as sending a message or transferring money out of a bank account) that the entities are performing. Reputation can thus be assigned to those identifiers based on their overall behavioral and historical patterns as well as their relationship to other identifiers, such as the relationship of IPs sending messages and URLs included in those messages. A "bad" reputation for a single identifier can cause the reputation of other neighboring identifiers to worsen, if there is a strong correlation between the identifiers. For example, an IP that is sending URLs which have a bad reputation will worsen its own reputation because of the reputation of the URLs. Finally, the individual identifier reputations can be aggregated into a single reputation (risk score) for the entity that is associated with those identifiers.

[0046] It should be noted that attributes can fall into a number of categories. For example, evidentiary attributes can represent physical, digital, or digitized physical data about an entity. This data can be attributed to a single known or unknown entity, or shared between multiple entities (forming entity relationships). Examples of evidentiary attributes relevant to messaging security include IP (internet protocol) address, known domain names, URLs, digital fingerprints or signatures used by the entity, TCP signatures, and etcetera.

[0047] As another example, behavioral attributes can represent human or machine-assigned observations about either an entity or an evidentiary attribute. Such attributes may include one, many, or all attributes from one or more behavioral profiles. For example, a behavioral attribute generically associated with a spammer may be a high volume of communications being sent from that entity.

[0048] A number of behavioral attributes for a particular type of behavior can be combined to derive a behavioral profile. A behavioral profile can contain a set of predefined behavioral attributes. The attributive properties assigned to these profiles include behavioral events relevant to defining the disposition of an entity matching the profile. Examples of behavioral profiles relevant to messaging security might include, "Spammer", "Scammer", and "Legitimate Sender". Events and/or evidentiary attributes relevant to each profile

define appropriate entities to which a profile should be assigned. This may include a specific set of sending patterns, blacklist events, or specific attributes of the evidentiary data. Some examples include: Sender/Receiver Identification; Time Interval and sending patterns; Severity and disposition of payload; Message construction; Message quality; Protocols and related signatures; Communications medium

[0049] It should be understood that entities sharing some or all of the same evidentiary attributes have an evidentiary relationship. Similarly, entities sharing behavioral attributes have a behavioral relationship. These relationships help form logical groups of related profiles, which can then be applied adaptively to enhance the profile or identify entities slightly more or less standard with the profiles assigned.

[0050] FIG. 4 is a flowchart depicting an operational scenario 400 used to detect relationships and assign risk to entities. The operational scenario begins at step 410 by collecting network data. Data collection can be done, for example, by a security agent 100, a client device, a switch, a router, or any other device operable to receive communications from network entities (e.g., e-mail servers, web servers, IM servers, ISPs, file transfer protocol (FTP) servers, gopher servers, VoIP equipments, etc.).

[0051] At step 420 identifiers are associated with the collected data (e.g., communication data). Step 420 can be performed by a security agent 100 or by a central system 200 operable to aggregate data from a number of sensor devices, including, for example, one or more security agents 100. Alternatively, step 420 can be performed by the security agents 100 themselves. The identifiers can be based upon the type of communication received. For example, an e-mail can include one set of information (e.g., IP) address of originator and destination, text content, attachment, etc.), while a VoIP communication can include a different set of information (e.g., originating phone number (or IP address if originating from a VoIP client), receiving phone number (or IP address if destined for a VoIP phone), voice content, etc.). Step 420 can also include assigning the attributes of the communication with the associated identifiers.

[0052] At step 430 the attributes associated with the entities are analyzed to determine whether any relationships exist between entities for which communications information has been collected. Step 430 can be performed, for example, by a central system 200 or one or more distributed security agents 100. The analysis can include comparing attributes related to different entities to find relationships between the entities. Moreover, based upon the particular attribute which serves as the basis for the relationship, a strength can be associated with the relationship.

[0053] At step 440 a risk vector is assigned to the entities. As an example, the risk vector can be assigned by the central system 200 or by one or more security agents 100. The risk vector assigned to an entity 130 (FIGS. 1-2), 300 (FIG. 3) can be based upon the relationship found between the entities and on the basis of the identifier which formed the basis for the relationship.

[0054] At step 450, an action can be performed based upon the risk vector. The action can be performed, for example, by a security agent 100. The action can be performed on a received communication associated with an entity for which a risk vector has been assigned. The action can include any of allow, deny, quarantine, load balance, deliver with assigned priority, or analyze locally with additional scrutiny, among

many others. However, it should be understood that a reputation vector can be derived separately

[0055] FIG. 5 is a block diagram illustrating an example network architecture including local reputations 500a-e derived by local reputation engines 510a-e and a global reputation 520 stored by one or more servers 530. The local reputation engines 510a-e, for example, can be associated with local security agents such as security agents 100. Alternatively, the local reputation engines 510a-e can be associated, for example, with a local client. Each of the reputation engines 510a-e includes a list of one or more entities for which the reputation engine 510a-e stores a derived reputation 500a-e.

[0056] However, these stored derived reputations can be inconsistent between reputation engines, because each of the reputation engines may observe different types of traffic. For example, reputation engine 1 510a may include a reputation that indicates a particular entity is reputable, while reputation engine 2 510b may include a reputation that indicates that the same entity is non-reputable. These local reputational inconsistencies can be based upon different traffic received from the entity. Alternatively, the inconsistencies can be based upon the feedback from a user of local reputation engine 1 510a indicating a communication is legitimate, while a user of local reputation engine 2 510b provides feedback indicating that the same communication is not legitimate.

[0057] The server 530 receives reputation information from the local reputation engines 510a-e. However, as noted above, some of the local reputation information may be inconsistent with other local reputation information. The server 530 can arbitrate between the local reputations 500a-e to determine a global reputation 520 based upon the local reputation information 500a-e. In some examples, the global reputation information 520 can then be provided back to the local reputation engines 510a-e to provide these local engines 510a-e with up-to-date reputational information. Alternatively, the local reputation engines 510a-e can be operable to query the server 530 for reputation information. In some examples, the server 530 responds to the query with global reputation information 520.

[0058] In other examples, the server 530 applies a local reputation bias to the global reputation 520. The local reputation bias can perform a transform on the global reputation to provide the local reputation engines 510a-e with a global reputation vector that is biased based upon the preferences of the particular local reputation engine 510a-e which originated the query. Thus, a local reputation engine 510a with an administrator or user(s) that has indicated a high tolerance for spam messages can receive a global reputation vector that accounts for an indicated tolerance. The particular components of the reputation vector returns to the reputation engine 510a might include portions of the reputation vector that are deemphasized with relationship to the rest of the reputation vector. Likewise, a local reputation engine 510b that has indicated, for example, a low tolerance communications from entities with reputations for originating viruses may receive a reputation vector that amplifies the components of the reputation vector that relate to virus reputation.

[0059] FIG. 6 is a block diagram illustrating a determination of a global reputation based on local reputation feedback. A local reputation engine 600 is operable to send a query through a network 610 to a server 620. In some examples, the local reputation engine 600 originates a query in response to receiving a communication from an unknown entity. Alterna-

tively, the local reputation engine 600 can originate the query responsive to receiving any communications, thereby promoting use of more up-to-date reputation information.

[0060] The server 620 is operable to respond to the query with a global reputation determination. The central server 620 can derive the global reputation using a global reputation aggregation engine 630. The global reputation aggregation engine 630 is operable to receive a plurality of local reputations 640 from a respective plurality of local reputation engines. In some examples, the plurality of local reputations 640 can be periodically sent by the reputation engines to the server 620. Alternatively, the plurality of local reputations 640 can be retrieved by the server upon receiving a query from one of the local reputation engines 600.

[0061] The local reputations can be combined using confidence values related to each of the local reputation engines and then accumulating the results. The confidence value can indicate the confidence associated with a local reputation produced by an associated reputation engine. Reputation engines associated with individuals, for example, can receive a lower weighting in the global reputation determination. In contrast, local reputations associated with reputation engines operating on large networks can receive greater weight in the global reputation determination based upon the confidence value associated with that reputation engine.

[0062] In some examples, the confidence values 650 can be based upon feedback received from users. For example, a reputation engine that receives a lot of feedback indicating that communications were not properly handled because local reputation information 640 associated with the communication indicated the wrong action can be assigned low confidence values 650 for local reputations 640 associated with those reputation engines. Similarly, reputation engines that receive feedback indicating that the communications were handled correctly based upon local reputation information 640 associated with the communication indicated the correct action can be assigned a high confidence value 650 for local reputations 640 associated with the reputation engine. Adjustment of the confidence values associated with the various reputation engines can be accomplished using a tuner 660, which is operable to receive input information and to adjust the confidence values based upon the received input. In some examples, the confidence values 650 can be provided to the server 620 by the reputation engine itself based upon stored statistics for incorrectly classified entities. In other examples, information used to weight the local reputation information can be communicated to the server 620.

[0063] In some examples, a bias 670 can be applied to the resulting global reputation vector. The bias 670 can normalize the reputation vector to provide a normalized global reputation vector to a reputation engine 600. Alternatively, the bias 670 can be applied to account for local preferences associated with the reputation engine 600 originating the reputation query. Thus, a reputation engine 600 can receive a global reputation vector matching the defined preferences of the querying reputation engine 600. The reputation engine 600 can take an action on the communication based upon the global reputation vector received from the server 620.

[0064] FIG. 7 is a block diagram illustrating an example resolution between a global reputation and a local reputation. The local security agent 700 communicates with a server 720 to retrieve global reputation information from the server 720. The local security agent 700 can receive a communication at 702. The local security agent can correlate the communica-

tion to identity attributes of the message at 704. The attributes of the message can include, for example, an originating entity, a fingerprint of the message content, a message size, etc. The local security agent 700 includes this information in a query to the server 720. In other examples, the local security agent 700 can forward the entire message to the server 720, and the server can perform the correlation and analysis of the message.

[0065] The server 720 uses the information received from the query to determine a global reputation based upon a configuration 725 of the server 720. The configuration 725 can include a plurality of reputation information, including both information indicating that a queried entity is non-reputable 730 and information indicating that a queried entity is reputable 735. The configuration 725 can also apply a weighting 740 to each of the aggregated reputations 730, 735. A reputation score determinator 745 can provide the engine for weighting 740 the aggregated reputation information 730, 735 and producing a global reputation vector.

[0066] The local security agent 700 then sends a query to a local reputation engine at 706. The local reputation engine 708 performs a determination of the local reputation and returns a local reputation vector at 710. The local security agent 700 also receives a response to the reputation query sent to the server 720 in the form of a global reputation vector. The local security agent 700 then mixes the local and global reputation vectors together at 712. An action is then taken with respect to the received message at 714.

[0067] FIG. 8 is an example graphical user interface 800 for adjusting the settings of a filter associated with a reputation server. The graphical user interface 800 can allow the user of a local security agent to adjust the settings of a local filter in several different categories 810, such as, for example, "Virus," "Worms," "Trojan Horse," "Phishing," "Spyware," "Spam," "Content," and "Bulk." However, it should be understood that the categories 810 depicted are merely examples, and that the disclosure is not limited to the categories 810 chosen as examples here.

[0068] In some examples, the categories 810 can be divided into two or more types of categories. For example, the categories 810 of FIG. 8 are divided into a "Security Settings" type 820 of category 810, and a "Policy Settings" type 830 of category. In each of the categories 810 and types 820, 830, a mixer bar representation 840 can allow the user to adjust the particular filter setting associated with the respective category 810 of communications or entity reputations.

[0069] Moreover, while categories 810 of "Policy Settings" type 830 can be adjusted freely based upon the user's own judgment, categories of "Security Settings" type 820 can be limited to adjustment within a range. This distinction can be made in order to prevent a user from altering the security settings of the security agent beyond an acceptable range. For example, a disgruntled employee could attempt to lower the security settings, thereby leaving an enterprise network vulnerable to attack. Thus, the ranges 850 placed on categories 810 in the "Security Settings" type 820 are operable to keep security at a minimum level to prevent the network from being compromised. However, as should be noted, the "Policy Settings" type 830 categories 810 are those types of categories 810 that would not compromise the security of a network, but might only inconvenience the user or the enterprise if the settings were lowered.

[0070] Furthermore, it should be recognized that in various examples, range limits 850 can be placed upon all of the

categories **810**. Thus, the local security agent would prevent users from setting the mixer bar representation **840** outside of the provided range **850**. It should also be noted, that in some examples, the ranges may not be shown on the graphical user interface **800**. Instead, the range **850** would be abstracted out of the graphical user interface **800** and all of the settings would be relative settings. Thus, the category **810** could display and appear to allow a full range of settings, while transforming the setting into a setting within the provided range. For example, the "Virus" category **810** range **850** is provided in this example as being between level markers **8** and **13**. If the graphical user interface **800** were set to abstract the allowable range **850** out of the graphical user interface **800**, the "Virus" category **810** would allow setting of the mixer bar representation **840** anywhere between 0 and 14. However, the graphical user interface **800** could transform the 0-14 setting to a setting within the 8 to 13 range **850**. Thus, if a user requested a setting of midway between 0 and 14, the graphical user interface could transform that setting into a settings of midway between 8 and 13.

[0071] FIG. 9 is a block diagram illustrating reputation based connection throttling for voice over internet protocol (VoIP) or short message service (SMS) communications. As should be understood, an originating IP phone **900** can place a VoIP call to a receiving IP phone **910**. These IP phones **900**, **910** can be, for example, computers executing soft-phone software, network enabled phones, etc. The originating IP phone **900** can place a VoIP call through a network **920** (e.g., the internet). The receiving IP phone **910** can receive the VoIP call through a local network **930** (e.g., an enterprise network).

[0072] Upon establishing a VoIP call, the originating IP phone has established a connection to the local network **930**. This connection can be exploited similarly to the way e-mail, web, instant messaging, or other internet applications can be exploited for providing unregulated connect to a network. Thus, a connection to a receiving IP phone can be exploited, thereby putting computers **940**, **950** operating on the local network **930** at risk for intrusion, viruses, trojan horses, worms, and various other types of attacks based upon the established connection. Moreover, because of the time sensitive nature of VoIP communications, these communications are typically not examined to ensure that the connection is not being misused. For example, voice conversations occur in real-time. If a few packets of a voice conversation are delayed, the conversation becomes stilted and difficult to understand. Thus, the contents of the packets typically cannot be examined once a connection is established.

[0073] However, a local security agent **960** can use reputation information received from a reputation engine or server **970** to determine a reputation associated with the originating IP phone. The local security agent **960** can use the reputation of the originating entity to determine whether to allow a connection to the originating entity. Thus, the security agent **960** can prevent connections to non-reputable entities, as indicated by reputations that do not comply with the policy of the local security agent **960**.

[0074] In some examples, the local security agent **960** can include a connection throttling engine operable to control the flow rate of packets being transmitted using the connection established between the originating IP phone **900** and the receiving IP phone **910**. Thus, an originating entities **900** with a non-reputable reputation can be allowed to make a connection to the receiving IP phone **910**. However, the packet throughput will be capped, thereby preventing the originating

entity **900** from exploiting the connection to attack the local network **930**. Alternatively, the throttling of the connection can be accomplished by performing a detailed inspection of any packets originating from non-reputable entities. As discussed above, the detailed inspection of all VoIP packets is not efficient. Thus, quality of service (QoS) can be maximized for connections associated with reputable entities, while reducing the QoS associated with connections to non-reputable entities. Standard communication interrogation techniques call be performed on connections associated with non-reputable entities in order to discover whether any of the transmitted packets received from the originating entity comprise a threat to the network **930**. Various interrogation techniques and systems are described in U.S. Pat. Nos. 6,941,467, 7,089, 590, 7,096,498, and 7,124,438 and in U.S. Patent Application Nos. 2006/0015942, 2006/0015563, 2003/0172302, 2003/0172294, 2003/0172291, and 2003/0172166, which are hereby incorporated by reference.

[0075] FIG. 10 is a block diagram illustrating an operation of a reputation based load balancer **1000**. The load balancer **1000** is operable to receive communications from reputable and non-reputable entities **1010**, **1020** (respectively) through a network **1030** (e.g. the internet). The load balancer **1000** communicates with a reputation engine **1040** to determine the reputation of entities **1010**, **1020** associated with incoming or outgoing communications.

[0076] The reputation engine **1030** is operable to provide the load balancer with a reputation vector. The reputation vector can indicate the reputation of the entity **1010**, **1020** associated with the communication in a variety of different categories. For example, the reputation vector might indicate a good reputation for an entity **1010**, **1020** with respect to the entity **1010**, **1020** originating spam, while also indicating a poor reputation for the same entity **1010**, **1020** with respect to that entity **1010**, **1020** originating viruses.

[0077] The load balancer **1000** can use the reputation vector to determine what action to perform with respect to a communication associated with that entity **1010**, **1020**. In situations where a reputable entity **1010** is associated with the communication, the message is sent to a message transfer agent (MTA) **1050** and delivered to a recipient **1060**.

[0078] In situations where a non-reputable entity **1020** has a reputation for viruses, but does not have a reputation for other types of non-reputable activity, the communication is forwarded to one of a plurality of virus detectors **1070**. The load balancer **1000** is operable to determine which of the plurality of virus detectors **1070** to use based upon the current capacity of the virus detectors and the reputation of the originating entity. For example, the load balancer **1000** could send the communication to the least utilized virus detector. In other examples, the load balancer **1000** might determine a degree of non-reputability associated with the originating entity and send slightly non-reputable communications to the least utilized virus detectors, while sending highly non-reputable communications to a highly utilized virus detector, thereby throttling the QoS of a connection associated with a highly non-reputable entity.

[0079] Similarly, in situations where a non-reputable entity **1020** has a reputation for originating spam communications, but no other types of non-reputable activities, the load balancer can send the communication to specialized spam detectors **1080** to the exclusion of other types of testing. It should be understood that in situations where a communication is associated with a non-reputable entity **1020** that originates

multiple types of non-reputable activity, the communication can be sent to be tested for each of the types of non-reputable activity that the entity 1020 is known to display, while avoiding tests associated with non-reputable activity that the entity 1020 is not known to display.

[0080] In some examples, every communication can receive routine testing for multiple types of non-legitimate content. However, when an entity 1020 associated with the communication shows a reputation for certain types of activity, the communication can also be quarantined for detailed testing for the content that the entity shows a reputation for originating.

[0081] In yet further examples, every communication may receive the same type of testing. However, communications associated with reputable entities 1010 is sent to the testing modules with the shortest queue or to testing modules with spare processing capacity. On the other hand, communications associated with non-reputable entities 1020 is sent to testing modules 1070, 1080 with the longest queue. Therefore, communications associated with reputable entities 1010 can receive priority in delivery over communications associated with non-reputable entities. Quality of service is therefore maximized for reputable entities 1010, while being reduced for non-reputable entities 1020. Thus, reputation based load balancing can protect the network from exposure to attack by reducing the ability of a non-reputable entity to connect to the network 930.

[0082] FIG. 11A is a flowchart illustrating an example operational scenario for collection of geolocation based data for authentication analysis. At step 1100 the operational scenario collects data from various login attempts. Step 1100 can be performed for example by a local security agent, such as the security agent 100 of FIG. 1. The collected data can include IP address associated with the login attempt, time of the login attempt, number of login attempts before successful, or the details of any unsuccessful passwords attempted, among many other types of information. The collected data is then analyzed in step 1105 to derive statistical information such as, for example, a geographical location of the login attempts. Step 1105 can be performed, for example, by a reputation engine. The statistical information associated with the login attempts is then stored at step 1110. The storing can be performed, for example, by a system data store.

[0083] FIG. 11B is a flowchart illustrating an example operational scenario for geolocation based authentication. A login attempt is received at step 1115. The login attempt can be received for example, by a secure web server operable to provide secure financial data over a network. It is then determined whether the login attempt matches a stored username and password combination at step 1120. Step 1120 can be performed, for example, by a secure server operable to authenticate login attempts. If the username and password do not match a stored username/password combination, the login attempt is declared a failure at step 125.

[0084] However, if the username and password do match a legitimate username/password combination, the origin of the login attempt is ascertained at step 1130. The origin of the login attempt can be determined by a local security agent 100 as described in FIG. 1. Alternatively, the origin of the login attempt can be determined by a reputation engine. The origin of the login attempt can then be compared with the statistical information derived in FIG. 11A, as shown in step 1135. Step 1135 can be performed, for example, by a local security agent 100 or by a reputation engine. It is determined whether the

origin matches statistical expectations at step 1140. If the actual origin matches statistical expectations, the user is authenticated at step 1145.

[0085] Alternatively, if the actual origin does not match statistical expectations for the origin, further processing is performed in step 1150. It should be understood that further processing can include requesting further information from the user to verify his or her authenticity. Such information can include, for example, home address, mother's maiden name, place of birth, or any other piece of information known about the user (e.g., secret question). Other examples of additional processing can include searching previous login attempts to determine whether the location of the current login attempt is truly anomalous or merely coincidental. Furthermore, a reputation associated with the entity originating the login attempt can be derived and used to determine whether to allow the login.

[0086] FIG. 11C is a flowchart illustrating another example operational scenario for geolocation based authentication using reputation of an originating entity to confirm authentication. A login attempt is received at step 1155. The login attempt can be received for example, by a secure web server operable to provide secure financial data over a network. It is then determined whether the login attempt matches a stored username and password combination at step 1160. Step 1160 can be performed, for example, by a secure server operable to authenticate login attempts. If the username and password do not match a stored username/password combination, the login attempt is declared a failure at step 1165.

[0087] However, if the username and password do match a legitimate username/password combination, the origin of the login attempt is ascertained at step 1170. The origin of the login attempt can be determined by a local security agent 100 as described in FIG. 1. Alternatively, the origin of the login attempt can be determined by a reputation engine. A reputation associated with the entity originating the login attempt can then be retrieved, as shown in step 1175. Step 1175 can be performed, for example, by a reputation engine. It is determined whether the reputation of the originating entity is reputable at step 1180. If the originating entity is reputable, the user is authenticated at step 1185.

[0088] Alternatively, if the originating entity is non-reputable, further processing is performed in step 1190. It should be understood that further processing can include requesting further information from the user to verify his or her authenticity. Such information can include, for example, home address, mother's maiden name, place of birth, or any other piece of information known about the user (e.g., secret question). Other examples of additional processing can include searching previous login attempts to determine whether the location of the current login attempt is truly anomalous or merely coincidental.

[0089] Thus, it should be understood that reputation systems can be applied to identifying fraud in financial transactions. The reputation system can raise the risk score of a transaction depending on the reputation of the transaction originator or the data in the actual transaction (source, destination, amount, etc). In such situations, the financial institution can better determine the probability that a particular transaction is fraudulent based upon the reputation of the originating entity.

[0090] FIG. 12 is a flowchart illustrating an example operational scenario for a reputation based dynamic quarantine. Communications are received at step 1200. The communica-

tions are then analyzed to determine whether they are associated with an unknown entity at step 1205. It should be noted, however, that this operational scenario could be applied to any communications received, not merely communications received from previously unknown entities. For example, communications received from a non-reputable entity could be dynamically quarantined until it is determined that the received communications do not pose a threat to the network. Where the communications are not associated with a new entity, the communications undergo normal processing for incoming communications as shown in step 1210.

[0091] If the communications are associated with a new entity, a dynamic quarantine counter is initialized in step 1215. Communications received from the new entity are then sent to a dynamic quarantined at step 1220. The counter is then checked to determine whether the counter has elapsed in step 1225. If the counter has not elapsed, the counter is decremented in step 1230. The behavior of the entity as well as the quarantined communications can be analyzed in step 1235. A determination is made whether the quarantined communications or behavior of the entity is anomalous in step 1240. If there is no anomaly found, the operational scenario returns to step 1220, where new communications are quarantined.

[0092] However, if the communications or behavior of the entity are found to be anomalous in step 1240, a non-reputable reputation is assigned to the entity in step 1245. The process ends by sending notification to an administrator or recipients of communications sent by the originating entity.

[0093] Returning to step 1220, the process of quarantining and examining communications and entity behavior continues until anomalous behavior is discovered, or until the dynamic quarantine counter elapses in step 1225. If the dynamic quarantine counter elapses, a reputation is assigned to the entity at step 1255. Alternatively, in situations where the entity is not an unknown entity, the reputation would be updated in steps 1245 or 1255. The operational scenario ends at step 1260 by releasing the dynamic quarantine where the dynamic quarantine counter has elapsed without discovery of an anomaly in the communications or in the originating entity behavior.

[0094] FIG. 13 is an example graphical user interface 1300 display of an image spam communication which can be classified as an unwanted image or message. As should be understood, image spam poses a problem for traditional spam filters. Image spam bypasses the traditional textual analysis of spam by converting the text message of the spam into an image format. FIG. 13 shows an example of image spam. The message shows an image 1310. While the image 1300 appears to be textual, it is merely the graphic encoding of a textual message. Image spam also typically includes a textual message 1320 comprising sentences which are structured correctly, but make no sense in the context of the message. The message 1320 is designed to elude spam filters that key on communications that only include an image 1310 within the communication. Moreover, the message 1320 is designed to trick filters that apply superficial testing to the text of a communication that includes an image 1310. Further, while these messages do include information about the origination of the message in the header 1330, an entity's reputation for originating image spam might not be known until the entity is caught sending image spam.

[0095] FIG. 14 is a flowchart illustrating an example operational scenario for detecting unwanted images (e.g., image

spam). It should be understood that many of the steps shown in FIG. 14 can be performed alone or in combination with any or all of the other steps shown in FIG. 14 to provide some detection of image spam. However, the use of each of the steps in FIG. 14 provides a comprehensive process for detecting image spam.

[0096] The process begins at step 1400 with analysis of the communication. Step 1400 typically includes analyzing the communication to determine whether the communication includes an image that is subject to image spam processing. At step 1410, the operational scenario performs a structural analysis of the communication to determine whether the image comprises spam. The header of the image is then analyzed in step 1420. Analysis of the image header allows the system to determine whether anomalies exist with respect to the image format itself (e.g., protocol errors, corruption, etc.). The features of the image are analyzed in step 1430. The feature analysis is intended to determine whether any of the features of the image are anomalous. The image can be normalized in step 1440. Normalization of an image typically includes removal of random noise that might be added by a spammer to avoid image fingerprinting techniques. Image normalization is intended to convert the image into a format that can be easily compared among images. A fingerprint analysis can be performed on the normalized image to determine whether the image matches images from previously received known image spam.

[0097] FIG. 15A is a flowchart illustrating an operational scenario for analyzing the structure of a communication. The operational scenario begins at step 1500 with analysis of the message structure. At step 1505 the hypertext markup language (HTML) structure of the communication is analyzed to introduce n-gram tags as additional tokens to a Bayesian analysis. Such processing can analyze the text 1320 that is included in an image spam communication for anomalies. The HTML structure of the message can be analyzed to define meta-tokens. Meta-token is the HTML content of the message, processed to discard any irrelevant HTML tags and compressed by removing white space to create a "token" for Bayesian analysis. Each of the above described tokens can be used as input to a Bayesian analysis for comparison to previously received communications.

[0098] The operational scenario then includes image detection at step 1515. The image detection can include partitioning the image into a plurality of pieces and performing fingerprinting on the pieces to determine whether the fingerprints match pieces of previously received images.

[0099] FIG. 15B is a flowchart illustrating an operational scenario for analyzing the features of an image to extract features of the message for input into a clustering engine to identify components of the image which align with known image spam. The operational scenario begins at step 1520 where a number of high level features of the image are detected for use in a machine learning algorithm. Such features can include values such as the number of unique colors, number of noise black pixels, number of edges in horizontal direction (sharp transitions between shapes), etc.

[0100] One of the features extracted by the operational scenario can include the number of histogram modes of the image, as shown at step 1525. The number of modes is yielded by an examination of spectral intensity of the image. As should be understood, artificial images will typically include fewer modes than natural images, because natural image colors are typically spread through a broad spectrum.

[0101] As described above, the features extracted from the image can be used to identify anomalies. In some examples, anomalies can include analyzing the characteristics of a message to determine a level of similarity of a number of features to the features of stored unwanted images. Alternatively, in some examples, the image features can also be analyzed for comparison with known reputable images to determine similarity to reputable images. It should be understood that none of the extracted features alone are determinative of a classification. For example, a specific feature might be associated with 60% of unwanted messages, while also being associated with 40% of wanted messages. Moreover, as the value associated with the feature changed, there might be a change in the probability that the message is wanted or unwanted. There are many features that can indicate a slight tendency. If each of these features are combined the image spam detection system can make classification decision.

[0102] The aspect ratio is then examined in step 1530 to determine whether there are any anomalies with respect to the image size or aspect. Such anomalies in the aspect ratio could be indicated by similarity of the image size or aspect ratio to known sizes or aspect ratios which are common to known image spam. For example, image spam can come in specific sizes to make the image spam look more like common e-mail. Messages that include images which share a common size with known spam images are more likely to be spam themselves. Alternatively, there are image sizes which are not conducive to spam (e.g., a 1"x1" square image might be difficult to read if a spammer inserted a message into the image). Messages that include images which are known to be non-conducive to spam insertion are less likely to be image spam. Thus, the aspect ratio of a message can be compared to common aspect ratios used in image spam to determine a probability that the image is an unwanted image or that the image is a reputable image.

[0103] At step 1535, the frequency distribution of the image is examined. Typically, natural pictures have uniform frequency distribution with a relative scarcity of sharp frequency gradations. On the other hand, image spam typically includes a choppy frequency distribution as a result of black letters being placed on a dark background. Thus, such non-uniform frequency distribution can indicate image spam.

[0104] At step 1540, the signal to noise ratio can be analyzed. A high signal to noise ratio might indicate that a spammer may be trying to evade fingerprinting techniques by introducing noise into the image. Increasing noise levels can thereby indicate an increasing probability that the image is an unwanted image.

[0105] It should be understood that some features can be extracted on the scale of the entire image, while other features can be extracted from subparts of the image. For example the image can be subdivided into a plurality of subparts. Each of the rectangles can be transformed into a frequency domain using a fast Fourier transform (FFT). In the transformed image, the predominance of frequencies in a plurality of directions can be extracted as features. These subparts of the transformed image can also be examined to determine the amount of high frequencies and low frequencies. In the transformed image, the points that are further away from the origin represent higher frequencies. Similarly to the other extracted features, these features can then be compared to known legitimate and unwanted images to determine which characteristics the unknown image shares with each type of known image. Moreover, the transformed (e.g., frequency domain)

image can also be divided into subparts (e.g., slices, rectangles, concentric circles, etc.) and compared against data from known images (e.g., both known unwanted images and known legitimate images).

[0106] FIG. 15C is a flowchart illustrating an operational scenario for normalizing the an image for spam processing. At step 1545, obfuscation and noise is removed from the image. As discussed previously, these can be introduced by spammers to evade fingerprinting techniques such as hashing by varying the sum of the hash such that it does not match any previously received hash fingerprints of known image spam. Obfuscation and noise removal can describe several techniques for removing artificial noise introduced by spammers. It should be understood that artificial noise can include techniques used by spammers such as banding (where a font included in the image is varied to vary the hash of the image).

[0107] An edge detection algorithm can be run on the normalized image at step 1550. In some examples, the edge detected image can be used provided to an optical character recognition engine to convert the edge detected image to text. The edge detection can be used to remove unnecessary detail from the picture which can cause inefficiency in processing the image again other images.

[0108] At step 1555, median filtering can be applied. The median filtering is applied to remove random pixel noise. Such random pixels can cause problems to content analysis of the image. The median filtering can help to remove single pixel type of noise introduced by spammers. It should be understood that single pixel noise is introduced by spammers using an image editor to alter one or more pixels in the image, which can make the image appear grainy in some areas, thereby making the image more difficult to detect.

[0109] At step 1560, the image is quantized. Quantizing of the image remove unnecessary color information. The color information typically requires more processing and is unrelated to the attempted propagation of the spam. Moreover, spammers could vary the color scheme in an image slightly and again vary the hash such that known image spam hashes would not match the derived hash from the color variant image spam.

[0110] At step 1565, contrast stretching is performed. Using contrast stretching the color scale in the image is maximized from black to white, even if the colors only vary through shades of gray. The lightest shade of the image is assigned a white value, while the darkest shade in the image is assigned a black value. All other shades are assigned their relative position in the spectrum in comparison to the lightest and darkest shades in the original image. Contrast stretching helps to define details in an image that may not make full use of the available spectrum and therefore can help to prevent spammers from using different pieces of the spectrum to avoid fingerprinting techniques. Spammers sometimes intentionally shift the intensity range of an image to defeat some types of feature identification engines. Contrast stretching can also help normalize an image such that it can be compared to other images to identify common features contained in the images.

[0111] FIG. 15D is a flowchart illustrating an operational scenario for analyzing the fingerprint of an image to find common fragments among multiple images. The operational scenario begins a step 1570 by defining regions within an image. A winnowing algorithm is then performed on the defined regions to identify the relevant portions of the image upon which fingerprints should be taken at step 1575. At step

1580, the operational scenario fingerprints the resulting fragments from the winnowing operation and determines whether there is a match between the fingerprints of the received image and known spam images. A similar winnowing fingerprint approach is described in United States Patent Application Publication No. 2006/0251068, which is hereby incorporated by reference.

[0112] As used in the description herein and throughout the claims that follow, the meaning of “a,” “an,” and “the” includes plural reference unless the context clearly dictates otherwise. Also, as used in the description herein and throughout the claims that follow, the meaning of “in” includes “in” and “on” unless the context clearly dictates otherwise. Finally, as used in the description herein and throughout the claims that follow, the meanings of “and” and “or” include both the conjunctive and disjunctive and may be used interchangeably unless the context clearly dictates otherwise.

[0113] Ranges may be expressed herein as from “about” one particular value, and/or to “about” another particular value. When such a range is expressed, another embodiment includes from the one particular value and/or to the other particular value. Similarly, when values are expressed as approximations, by use of the antecedent “about,” it will be understood that the particular value forms another embodiment. It will be further understood that the endpoints of each of the ranges are significant both in relation to the other endpoint, and independently of the other endpoint.

[0114] A number of embodiments of the invention have been described. Nevertheless it will be understood that various modifications may be made without departing from the spirit and scope of the invention. Accordingly, other embodiments are within the scope of the following claims.

1. A computer implemented method operable to assign a reputation to a communications entity associated with a received communication, comprising:

- arranging a plurality of agents within a network, the plurality of agents being associated with a security device operable to protect an associated network from communications that violate a policy associated with the associated network;
- collecting data associated with entities originating communications, wherein collecting data comprises using the plurality of agents to collect data associated with the communications;
- aggregating the collected data;
- analyzing the aggregated data to identify attributes respectively associated with entities originating communications;
- correlating the attributes to identify relationships between entities;
- updating a reputation associated with one or more entities based upon the relationship to one or more other entities identified by correlating the attributes; and
- communicating updated reputation information to one or more of the plurality of agents.

2. The method of claim 1, wherein the reputation comprises a reputation vector indicating a reputation of an entity in a plurality of categories, the categories representing types of activity in which the entity or related entities have engaged.

3. The method of claim 2, further comprising using the communicated updated reputation information to determine whether to allow a voice over internet protocol connection or a short message service connection.

4. The method of claim 3, further comprising comparing the reputation information to a policy to determine whether to allow a voice over internet protocol connection or a short message service connection.

5. The method of claim 4, wherein the policy is defined by an administrator of the associated network to which the policy belongs.

6. The method of claim 2, further comprising using the communicated reputation information to determine one or more selected interrogation engines selected from a plurality of interrogation engines, wherein the selected interrogation engines will interrogate a received communication.

7. The method of claim 6, wherein the one or more selected interrogation engines omit any of the plurality of interrogation engines that test for risks that are not implicated by the reputation vector.

8. The method of claim 2, further comprising using the communicated reputation information to determine which of a plurality of agents to use for interrogation of a received communication.

9. The method of claim 8, wherein a communication associated with an entity that does not comply with policy for the associated network is assigned to an agent with an interrogation queue experiencing the highest interrogation volume.

10. The method of claim 8, wherein a communication associated with an entity that complies with policy for the associated network is assigned to an agent with an interrogation queue experiencing the lowest interrogation volume.

11. The method of claim 2, wherein the reputation comprises one or more of a reputation for originating spam, a reputation for originating image spam, a reputation for originating attacks, a reputation for intrusion, a reputation for originating phishing communications, a geolocation-based reputation, or a reputation for originating bulk mail.

12. The method of claim 2, wherein the plurality of agents include local reputation information, the local reputation information comprising one or more reputations respectively associated with one or more entities.

13. The method of claim 12, further comprising: aggregating the local reputation information; and deriving a global reputation vector based upon the aggregation of the local reputation information.

14. The method of claim 13, further comprising: receiving a reputation query from an agent; deriving the global reputation vector responsive to the reputation query; and communicating the global reputation vector to the agent originating the reputation query.

15. The method of claim 14, wherein the local reputation information comprises a plurality of local reputations and confidence values associated with the local reputations, wherein aggregating the local reputation information comprises adjusting each of the local reputations by an associated confidence value and combining the adjusted local reputations to produce a native global reputation vector.

16. The method of claim 15, further comprising applying a transform to the native global reputation vector to produce the global reputation vector.

17. The method of claim 16, wherein the transform comprises a local bias associated with the agent originating the reputation query.

18. The method of claim 17, wherein the local bias is operable to adjust the native global bias based upon preferences associated with the associated network.

19. The method of claim 1, wherein the entities comprise web entities operable to receive hypertext transfer protocol requests and to communicate one or more web pages responsive to the hypertext transfer protocol requests, wherein the reputation associated with the web entities comprises a reputation for communication of non-reputable web pages.

20. A computer implemented method operable to assign a reputation to a communications entity associated with a received communication, comprising:

- collecting data associated with entities originating communications, wherein collecting data comprises receiving data from a plurality of agents to collect data associated with the communications;

- aggregating the collected data;

- analyzing the aggregated data to identify attributes respectively associated with entities originating communications;

- correlating the attributes to identify relationships between entities;

- updating a reputation associated with one or more entities based upon the relationship to one or more other entities identified by correlating the attributes; and

- handling communications based upon the updated reputation information.

21. A distributed system operable to derive and communicate a reputation associated with a communications entity, comprising:

- a communications interface operable to communicate with a plurality of agents arranged within a global network, the plurality of agents being operable to derive local reputations associated with entities from which communications are received, wherein the plurality of agents are further operable to collect data associated with received communications;

- one or more data aggregation engines operable to aggregate the collected data via the communications interface;

- an analyzer operable to analyze the data to identify attributes respectively associated with entities originating the received communications;

- a correlation engine operable to correlate the attributes of the entities and to identify relationships between the entities;

- a reputation engine operable to identify relationships between the entities and to update reputations associated with one or more entities based upon its relationship to one or more other entities; and

- wherein the communications interface is further operable to communicate updated reputation information to devices operating on the global network.

22. The system of claim 21, wherein the reputation comprises a reputation vector indicating a reputation of an entity in a plurality of categories, the categories representing types of activity in which the entity has engaged.

23. The system of claim 22, wherein the devices operating on the global network are operable to use the communicated updated reputation information to determine whether to allow a voice over internet protocol connection or a short message service connection.

24. The system of claim 23, wherein the devices are operable to compare the reputation information to a policy to determine whether to allow a voice over internet protocol connection or a short message service connection.

25. The system of claim 24, wherein the policy is defined by an administrator of the associated network to which the device belongs.

26. The system of claim 22, further comprising a load balancer operable to use the communicated reputation information to determine one or more selected interrogation engines selected from a plurality of interrogation engines, wherein the one or more selected interrogation engines will interrogate a received communication.

27. The system of claim 26, wherein the one or more selected interrogation engines comprise those of the plurality of interrogation engines that test for risks that are implicated by the reputation vector.

28. The system of claim 22, further comprising a load balancer operable to use the communicated reputation information to determine which of a plurality of agents to use for interrogation of a received communication.

29. The system of claim 28, wherein a communication associated with an entity that does not comply with policy for the associated network is assigned to an agent with an interrogation queue experiencing the highest interrogation volume.

30. The system of claim 28, wherein a communication associated with an entity that complies with policy for the associated network is assigned to an agent with an interrogation queue experiencing the lowest interrogation volume.

31. The system of claim 22, wherein the reputation comprises one or more of a reputation for originating spam, a reputation for originating image spam, a reputation for originating attacks, a reputation for intrusion, a reputation for originating phishing communications, or a reputation for originating bulk mail.

32. The system of claim 22, wherein the plurality of agents include local reputation information, the local reputation information comprising one or more reputations respectively associated with one or more entities.

33. The system of claim 32, further comprising a reputation aggregation engine operable to aggregate local reputation information retrieved from the plurality of agents, and to derive a global reputation vector based upon the aggregation of the local reputation information.

34. The system of claim 33, wherein the communications interface is operable to receive a reputation query from an agent, and the reputation aggregation engine is operable to derive the global reputation vector responsive to the reputation query, whereby the communications interface can communicate the global reputation vector to the agent originating the reputation query.

35. The system of claim 34, wherein the local reputation information comprises a plurality of local, reputations and confidence values associated with the local reputations, wherein reputation aggregation engine is operable to adjust each of the local reputations by an associated confidence value and combining the adjusted local reputations to produce a native global reputation vector.

36. The system of claim 35, wherein the reputation aggregation engine is further operable to apply a transform to the native global reputation vector to produce the global reputation vector.

37. The system of claim 36, wherein the transform comprises a local bias associated with the agent originating the reputation query.

38. The system of claim **37**, wherein the local bias is operable to adjust the native global bias based upon preferences associated with the associated network.

39. The system of claim **21**, wherein the entities comprise web entities operable to receive hypertext transfer protocol requests and to communicate one or more web pages responsive to the hypertext transfer protocol requests, wherein the reputation associated with the web entities comprises a reputation for communication of non-reputable web pages.

40. A system operable to derive a reputation associated with a communications entity, comprising:

- a communications interface operable to receive information from a plurality of agents or a central server within a global network, the plurality of agents being operable to collect data associated with received communications;

- one or more data aggregation engines operable to aggregate the received information from the communications interface;

- an analyzer operable to analyze the received information to identify attributes respectively associated with entities originating the received communications;

- a correlation engine operable to correlate the attributes of the entities and to identify relationships between the entities;

- a reputation engine operable to identify relationships between the entities and to update reputations associated with one or more entities based upon its relationship to one or more other entities; and

- a traffic control engine operable to determine handling associated with a communication based upon the updated reputations.

41. A system comprising:

- a security control interface operable to produce a plurality of security control representations, each of the plurality of security control representations being operable to control a plurality of security settings associated with a protected entity; and

- a policy control interface operable to produce a plurality of policy control representations, each of the plurality of policy control representations being operable to control a plurality of policy settings associated with a protected entity;

- a filtering module operable to filter one or more communication streams based upon the plurality of security settings and based upon the plurality of policy settings.

42. The system of claim **41**, wherein the security control representations comprise a plurality of security slider representations in a plurality of security categories operable to control the security settings associated with the protected entity.

43. The system of claim **42**, wherein the plurality of security categories comprise two or more of a virus category, a phishing category, a worms category, or a trojan horse category.

44. The system of claim **43**, wherein the plurality of security control representations are each operable to adjust a threshold sensitivity for an associated one of the security settings.

45. The system of claim **44**, wherein the threshold sensitivity comprises a level of similarity between communication stream characteristics and characteristics associated with the security categories.

46. The system of claim **42**, wherein the policy control representations comprise a plurality of policy slider representations in a plurality of policy categories operable to control the policy settings associated with the protected entity.

47. The system of claim **46**, wherein the plurality of policy categories comprise two or more of a spam category, a content category, a spyware category, or a bulk mail category.

48. The system of claim **47**, wherein the plurality of policy control representations are each operable to adjust a threshold sensitivity for an associated one of the policy settings.

49. The system of claim **48**, wherein the threshold sensitivity comprises a level of similarity between communication stream characteristics and characteristics associated with the policy categories.

50. The system of claim **41**, wherein the protected entity is one of a computing device, a communications device, a mobile device, or a network.

51. The system of claim **41**, wherein the security control interface and the policy control interface are controlled by an administrator.

52. The system of claim **41**, wherein the security control interface and the policy control interface are controlled by an end user.

53. The system of claim **52**, wherein the security control interface include a plurality of ranges associated with the plurality of security control representations, the security control settings operable to be adjusted within the range.

54. The system of claim **53**, wherein the plurality of ranges are controlled by an administrator.

55. A computer implemented method comprising:

- receiving a plurality of ranges from an administrator;

- providing a security control interface to a user, the security control interface comprising a plurality of security control representations associated with security controls, each of the security control mechanisms including an associated range from among the plurality of ranges, the associated range defining a minimum and maximum setting associated with the respective security controls;

- receiving a plurality of security control settings from the user through the security control interface;

- adjusting a plurality of thresholds related to plurality of control settings received from the user, the plurality of thresholds being associated with tolerance for a classification of potential security violation; and

- filtering communications streams from a protected entity associated with the user based upon the plurality of thresholds.

56. The method of claim **55**, wherein the security control representations comprise a plurality of security slider representations in a plurality of security categories operable to control the security settings associated with the protected entity.

57. The method of claim **56**, wherein the plurality of security categories comprise two or more of a virus category, a phishing category, a worms category, a trojan horse category, a spam category, a content category, a spyware category, or a bulk mail category.

58. The method of claim **57**, wherein the plurality of security control representations are each operable to adjust a threshold sensitivity for an associated one of the security settings.

59. The method of claim 58, wherein the threshold sensitivity comprises a level of similarity between communication stream characteristics and characteristics associated with the security categories.

60. The method of claim 55, wherein the protected entity is one of a computing device, a communications device, a mobile device, or a network.

61. The method of claim 55, wherein the user comprises an administrator.

62. One or more computer readable media having software program code operable to enable filter adjustments for incoming and outgoing communications streams, comprising:

- receiving a plurality of ranges from an administrator;
- providing a security control interface to a user, the security control interface comprising a plurality of security control representations associated with a plurality of secu-

- rity control settings, each of the security control mechanisms including an associated range from among the plurality of ranges, the associated range defining a minimum and maximum setting associated with the respective security controls;

- receiving input from the user through the security control interface, the input requesting adjustment of the security control settings;

- adjusting a plurality of thresholds related to plurality of control settings received from the user, the plurality of thresholds being associated with tolerance for a classification of potential security violation; and

- filtering communications streams from a protected entity associated with the user based upon the plurality of thresholds.

* * * * *