



US 20100138533A1

(19) **United States**

(12) **Patent Application Publication**
Pauly

(10) **Pub. No.: US 2010/0138533 A1**

(43) **Pub. Date: Jun. 3, 2010**

(54) **METHOD OF IDENTIFYING A ROOT CAUSE OF A NETWORK EVENT**

Publication Classification

(75) Inventor: **Duncan Gunther Pauly**,
Gloucestershire (GB)

(51) **Int. Cl.**
G06F 15/173 (2006.01)

(52) **U.S. Cl.** **709/224**

(57) **ABSTRACT**

Correspondence Address:
Husch Blackwell Sanders, LLP
Husch Blackwell Sanders LLP Welsh & Katz
120 S RIVERSIDE PLAZA, 22ND FLOOR
CHICAGO, IL 60606 (US)

A method of identifying a root cause of an event in a network of linked entities. The method comprises: defining a plurality of signatures, each signature comprising one or more classes, the classes being categorized by a taxonomy of classes; receiving a first message from the network; assigning one or more of the signatures to the first message; establishing a set of one or more hypotheses, each hypothesis identifying a respective root cause; and one of the signatures assigned to the first message; receiving one or more subsequent messages from the network; determining whether each subsequent message matches a signature identified by a previously hypothesis; and generating an output based on the determination.

(73) Assignee: **ZENULTA LIMITED**, Blinkworth,
Wiltshire (GB)

Typically the method involves an element of human interaction. Typically at least one of the signatures includes a class which encloses two or more subclasses categorized by the taxonomy of classes. The ability to define signatures including such higher level entity classes in the taxonomy enables a human user to generate, store and edit generic hypotheses quickly and intuitively, and/or provides an output which can be easily understood by a human user.

(21) Appl. No.: **12/596,274**

(22) PCT Filed: **Apr. 15, 2008**

(86) PCT No.: **PCT/GB2008/001326**

§ 371 (c)(1),
(2), (4) Date: **Feb. 15, 2010**

(30) **Foreign Application Priority Data**

Apr. 18, 2007 (GB) 0707513.8

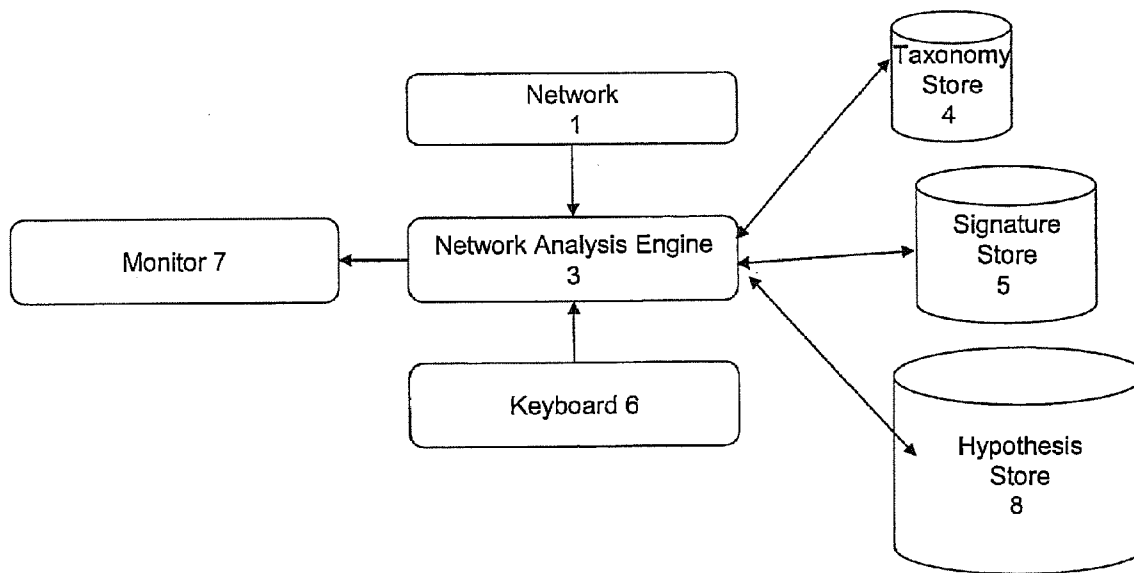
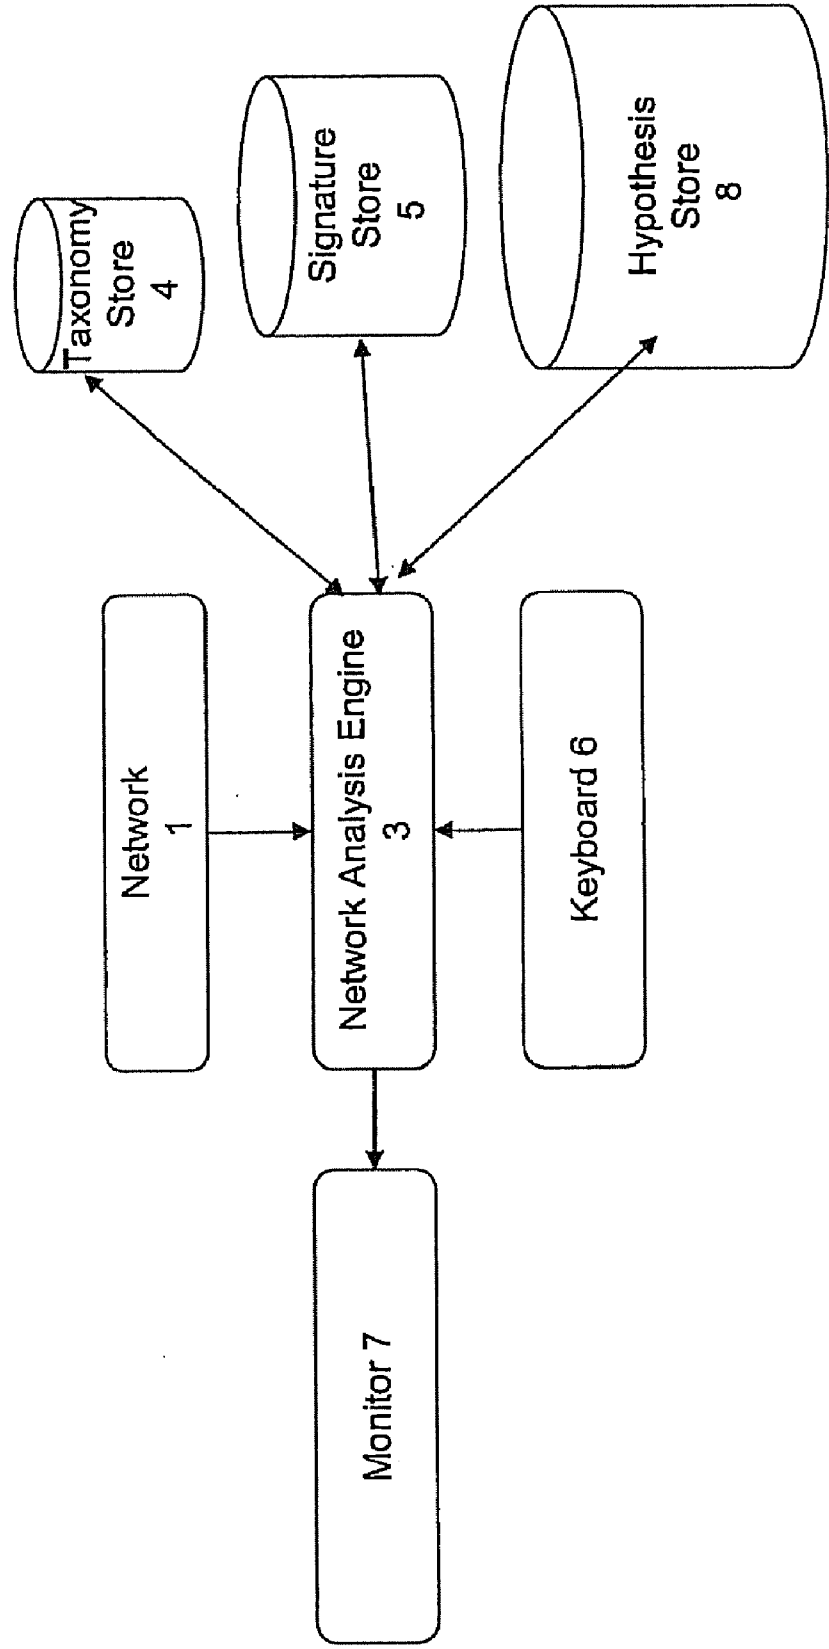


Fig. 1






Fig. 2

Fig. 3




Fig. 4

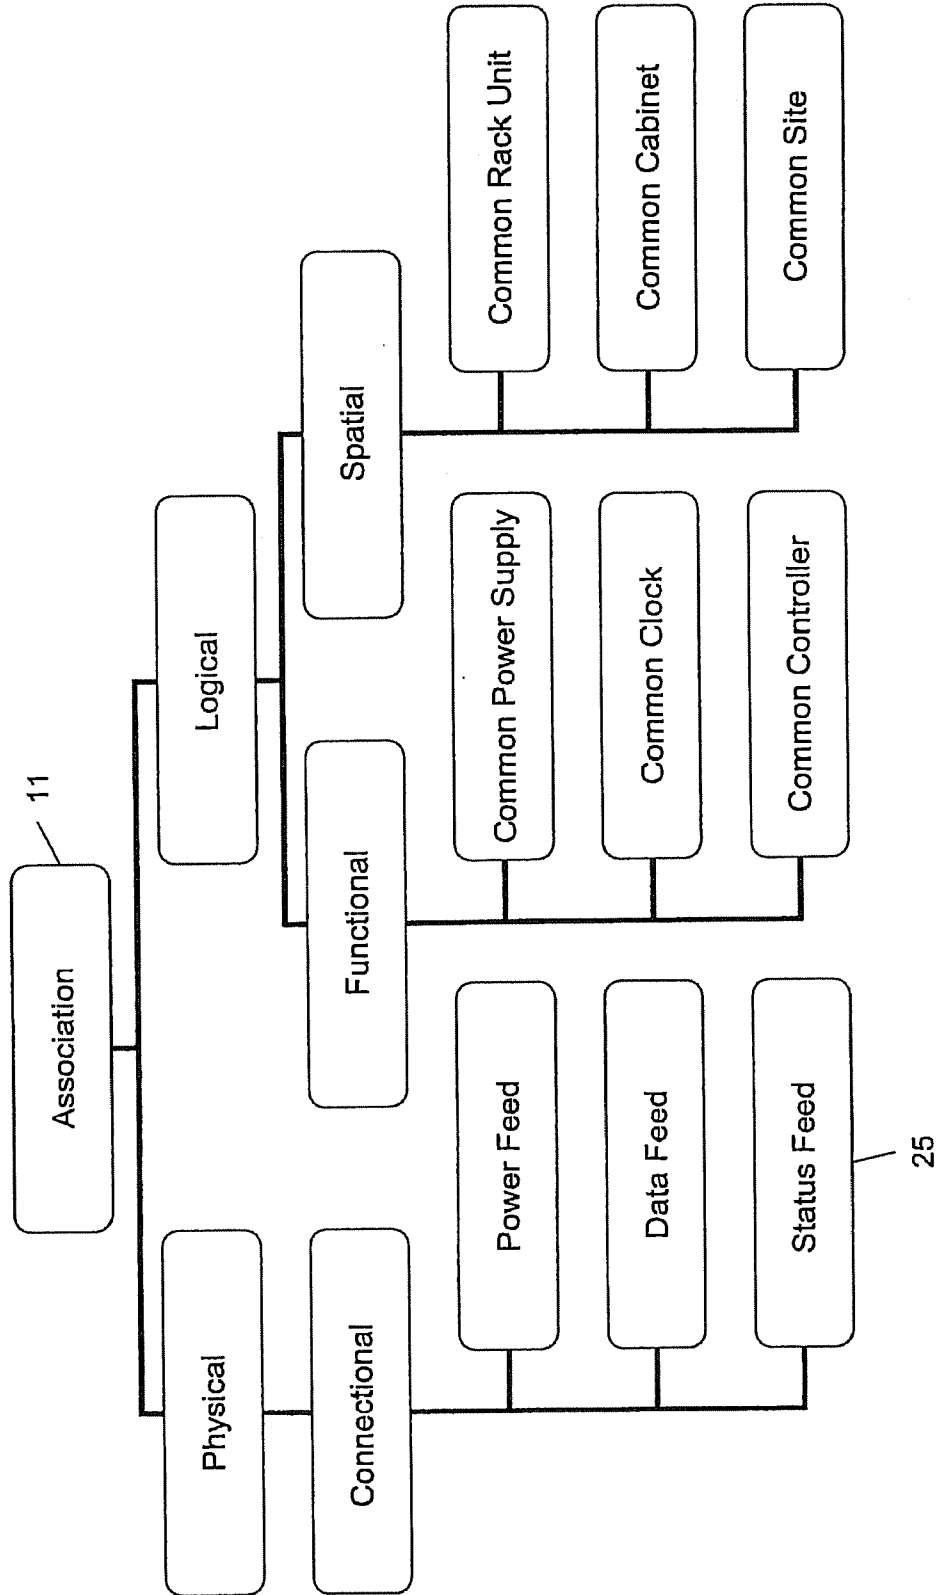


Fig. 5

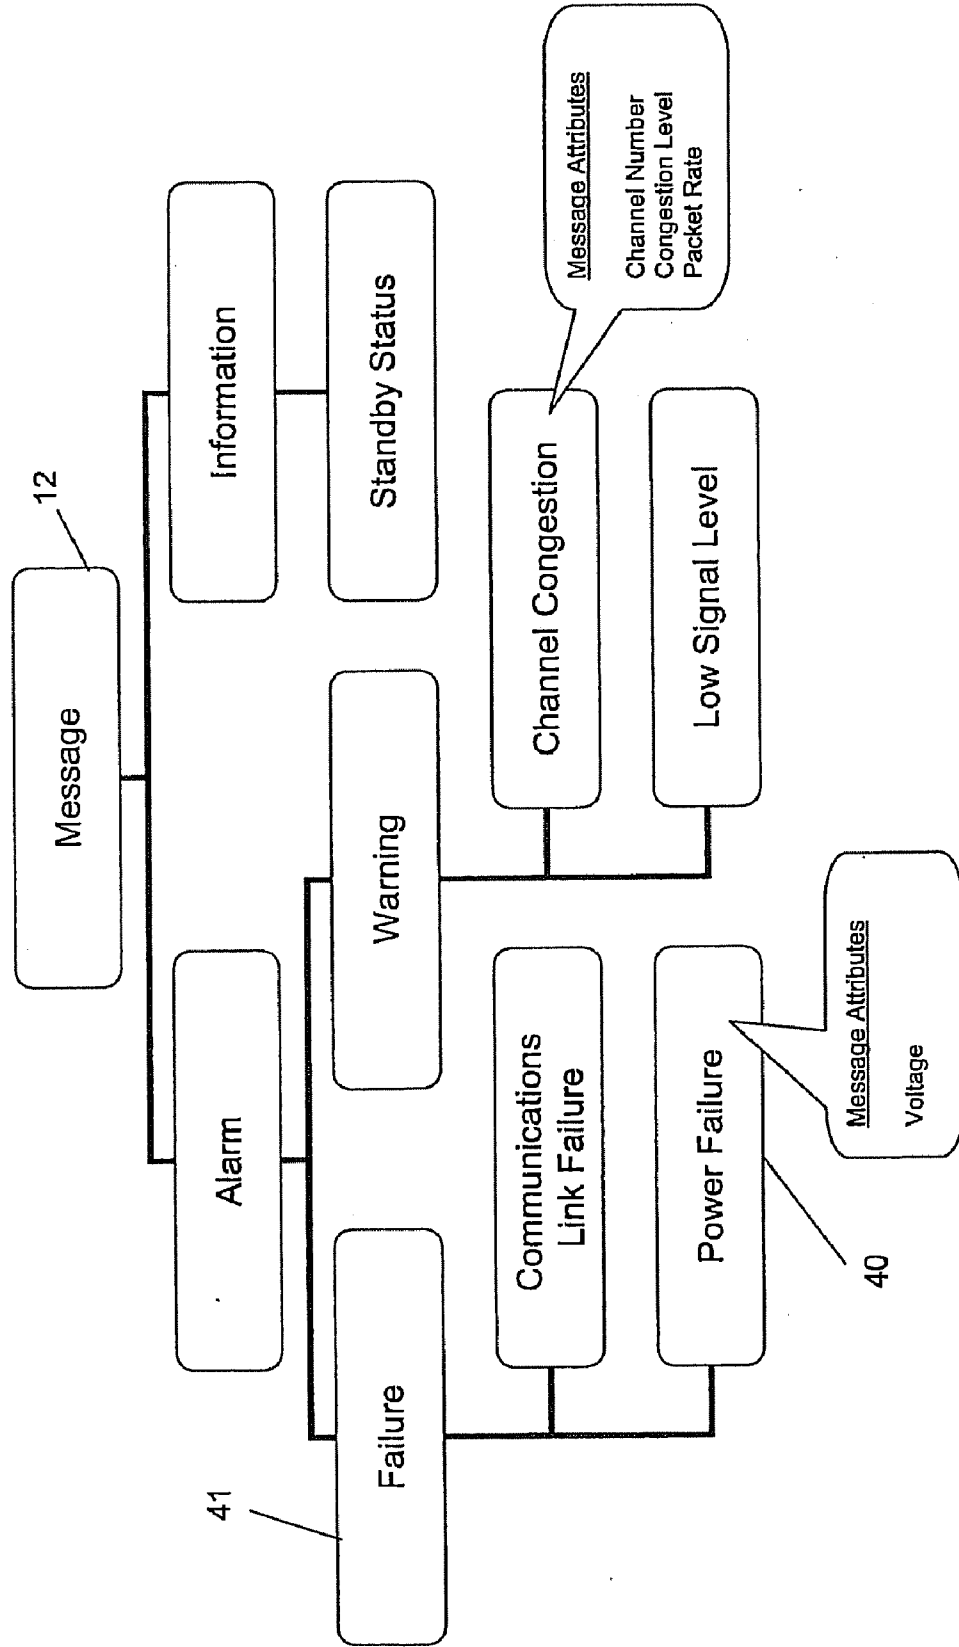
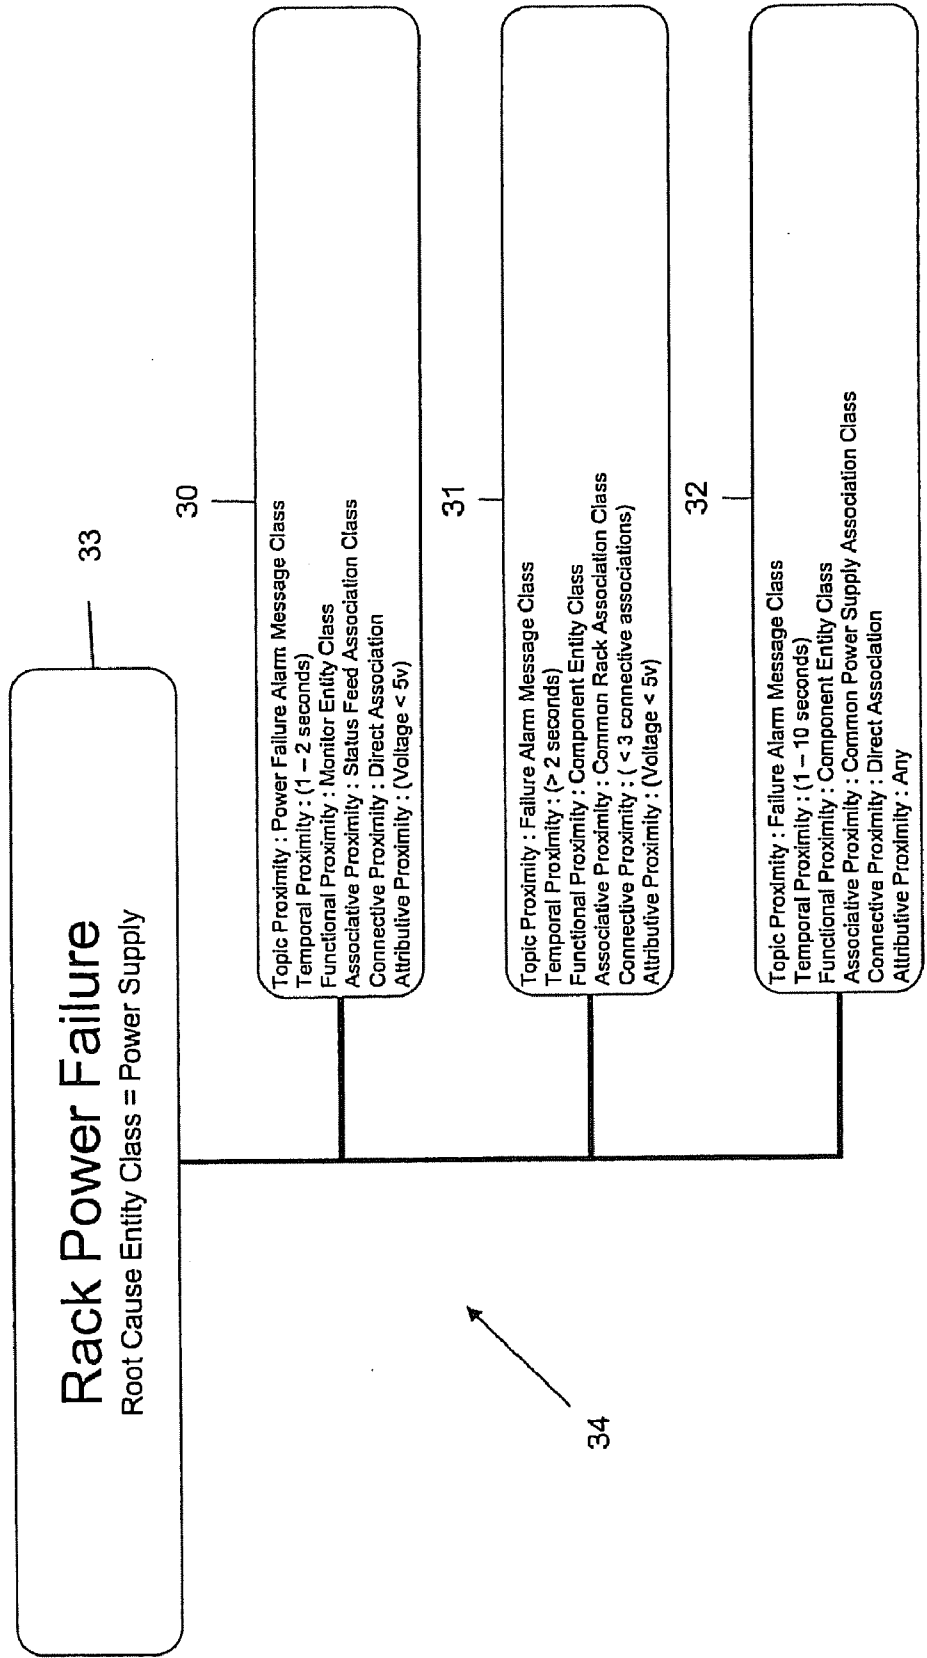


Fig. 6



METHOD OF IDENTIFYING A ROOT CAUSE OF A NETWORK EVENT

FIELD OF THE INVENTION

[0001] The present invention relates to a method of identifying a root cause of an event in a network of linked entities.

BACKGROUND OF THE INVENTION

[0002] A number of correlation techniques are currently available for identifying a root cause of an event in a network of linked entities. These include rule sequence, state transition, code-book and case based analysis. Each of these exhibits flaws that reduce the effectiveness and worth of the correlation process.

[0003] The correlation process may require precise and complete modelling of the physical network topology for a reasoned analysis of events. Any inaccuracies in the physical topology model arising from unavailable, stale, incomplete or incorrect information will undermine the accuracy of any analysis performed.

[0004] Moreover, the model used to depict the network structure may be rigid and unable to dynamically adapt to new information about network entity configuration and topology changes, creating inevitable errors in the correlation process when changes to the network occur.

[0005] The rules that guide the correlation process can be non-intuitive and make the analysis of the root cause event opaque and incomprehensible to human observers, which limits the trust that humans will place on the outcome of any analysis and restricts their ability to recognize why or where an outcome is inaccurate. This hinders the evolution of the correlation process in becoming a more consistent, accurate and complete process by restricting human interaction in the correlation process.

[0006] The context used for aggregating events under a common cause may be transient causing related events to be missed when they are delayed by communication lags and failures and arrive after the relevant correlation context has lapsed.

[0007] The overhead of the correlation process itself may enforce restrictions on the throughput of the events considered for correlation and the volume of events received for correlation may be attenuated by prior filtering and consolidation, which inevitably limit the information available to the correlation process and its accuracy.

[0008] The correlation model may recognize only narrow aspects of network functionality and associated operational activities and be unaware of all the information relevant to network status and performance. A partial view of the information relevant to a situation undermines the accuracy of the analysis performed around it.

SUMMARY OF THE INVENTION

[0009] A first aspect of the invention provides a method of identifying a root cause of an event in a network of linked entities, the method comprising:

[0010] a. defining a plurality of signatures, each signature comprising one or more classes, the classes being categorized by a taxonomy of classes;

[0011] b. receiving a first message from the network;

[0012] c. assigning one or more of the signatures to the first message;

[0013] d. establishing a set of one or more hypotheses, each hypothesis identifying:

[0014] i. a respective root cause; and

[0015] ii. one of the signatures assigned to the first message;

[0016] e. receiving one or more subsequent messages from the network;

[0017] f. determining whether each subsequent message matches a signature identified by a hypothesis established in step d; and

[0018] g. generating an output based on the determination in step f.

[0019] Typically the method involves an element of human interaction. That is, a human may provide input into step a. of the process, and/or may be presented with the output generated in step g. Typically at least one of the signatures defined in step a. may include a class which encloses two or more subclasses categorized by the taxonomy of classes. The ability to define signatures including such higher level entity classes in the taxonomy enables such a human user to generate, store and edit generic hypotheses quickly and intuitively, and/or provides an output which can be easily understood by such a human user.

[0020] Typically at least one of the signatures comprises a plurality of envelope sets each comprising a set of classes; wherein each envelope set is associated with one or more sub-events related to the root cause; and wherein a signature is assigned to the first message in step c if the first message matches a set of classes in at least one of that signature's envelope sets.

[0021] Typically at least one of the signatures comprises a message class categorized by a taxonomy of message classes and defining a message topic; and wherein steps c and f comprise matching a message class of the message with the message class of the signature.

[0022] Typically at least one of the signatures comprise one or more entity classes categorized by a taxonomy of entity classes, each entity class defining an entity type; and steps c and f comprise matching an entity class of the message with the entity class of the signature.

[0023] Typically at least one of the signatures comprise one or more association classes categorized by a taxonomy of association classes, each association class defining a type of connection between two directly communicating entities; and steps c and f comprise matching an association class of the message with the association class of the signature.

[0024] Typically the method further comprises giving the one or more hypotheses scores, and amending the scores when the signatures associated with them match subsequent received messages. In this case the method may further comprise promoting a hypothesis to an elected phase when its score reaches a predetermined level, increasing its score when the signature associated with it matches a message, and decreasing the score of any other co-existing hypotheses whose signature also matches the message.

[0025] Typically step f comprises comparing an observation time associated with the subsequent message with an event time associated with the hypothesis established in step d.

[0026] A second aspect of the invention provides a method of identifying a root cause of an event in a network of linked entities, the method comprising:

[0027] a. receiving a first message from the network;

[0028] b. establishing a first hypothesis associated with the first message;

- [0029] c. receiving one or more subsequent messages from the network;
 - [0030] d. for each subsequent message:
 - [0031] i. establishing a new hypothesis, and;
 - [0032] ii. determining whether the subsequent message matches any previously established hypothesis by comparing an observation time associated with the message with an event time associated with the hypothesis; and
 - [0033] e. generating an output based on the determination in step d.ii.
- [0034] By establishing a new hypothesis for each subsequent message as well as matching against previously established hypotheses, the method can provide an accurate estimation of the root cause. Note that some received messages may not establish new hypotheses, for instance if they do not match a previously defined signature.
- [0035] A further aspect of the invention provides a network analysis engine configured to identify a root cause of an event in a network of linked entities by the method of any preceding aspect.
- [0036] A further aspect of the invention provides a program or suite of programs for controlling the operating of a network analysis engine, the program(s) being arranged when executed to control the network analysis engine so as to operate in accordance with the method of any previous aspect.
- [0037] Although the embodiment below describes the invention in the context of a telecommunication network, this is for illustrative purposes only and the invention can be used to identify a root cause of an event in any network of linked entities.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0038] Embodiments of the invention will now be described with reference to the accompanying drawings, in which:
- [0039] FIG. 1 is a schematic view of a system architecture;
- [0040] FIG. 2 shows a portion of the network;
- [0041] FIG. 3 shows an entity taxonomy;
- [0042] FIG. 4 shows an association taxonomy;
- [0043] FIG. 5 shows a message taxonomy; and
- [0044] FIG. 6 shows a signature.

DETAILED DESCRIPTION OF EMBODIMENT(S)

[0045] FIG. 1 illustrates a telecommunication network 1, a portion 2 of the network being shown in FIG. 2. The network 1 is a graph of entities linked by associations, the associations being indicated by block arrows. Each entity provides a physical or logical function that contributes to the overall service provided by the network. Functional links between entities that do not directly represent physical assets of the network are known as “associations”.

[0046] Entities report messages to a network analysis engine 3 shown in FIG. 1. These messages notify the engine 3 of functional performance and significant events. When a single abnormality (or other event) occurs within the network 1, one or more entities may notify the engine 3 of alarm conditions associated directly or indirectly with the abnormality. To a network observer, the root cause of the abnormality may not be immediately apparent from the various

messages notified. Indeed, a large volume of indirectly symptomatic messages may serve to conceal the most indicative and relevant messages.

[0047] The engine 3 performs network event correlation to collect and aggregate the related symptomatic messages and discern a common aspect in order to suggest a root cause for the abnormality. The engine 3 performs event aggregation, as described in further detail below, to reduce the volume of information that has to be assimilated by network observers to promote clarity and assist in focusing effort at identifying the root cause.

[0048] The entities in the network are categorized by the engine 3 into a complete taxonomy of entity classes shown in FIG. 3. The taxonomy is rooted by a single generic and indistinct entity class 10, and each subclass inherits the properties of its enclosing class. The entity taxonomy of FIG. 3 is stored by the engine 3 in a taxonomy store 4 shown in FIG. 1.

[0049] An association symbolizes the connection between two directly communicating entities within the network. An association is passive, but is purposeful and every association in the network is categorized within a complete taxonomy of association classes shown in FIG. 4. The taxonomy is rooted by a single generic and indistinct association class 11 and each subclass inherits the properties of its enclosing class. The association taxonomy of FIG. 4 is stored by the engine 3 in the taxonomy store 4 shown in FIG. 1 along with the entity taxonomy.

[0050] A message conveys information about an event relevant to an entity associated with the network. A message may raise an event or may clear a pre-existing event. A message may be the event itself, in which case it is never cleared.

[0051] A message is issued by an observing entity at an observation time and contains information about an observed entity associated directly or indirectly with the observing entity.

[0052] Every message topic is categorized within a complete taxonomy of message classes shown in FIG. 5. The taxonomy is rooted by a single generic and indistinct message class 12 and each subclass inherits the properties of its enclosing class. The message taxonomy of FIG. 5 is stored along with the other taxonomies in the taxonomy store 4 shown in FIG. 1.

[0053] Every message has a perceived severity of critical, major, minor, information or unknown. This severity is specified for every message class in the taxonomy of message classes. For example, with reference to FIG. 5, all failure messages may have a severity of “major” and all warning messages may have a severity of “minor”.

[0054] Every message also has a priority of resolution which is a number from a bounded range of possible numbers. This priority is specified for every message class in the taxonomy of message classes. For example, with reference to FIG. 5, all failure messages may have a priority of 8/10 and all warning messages may have a priority of 3/10. The priority is used to order messages of a common severity to apply a relative scale of importance. For example one critical message may have a higher priority than another critical message to indicate it needs attention first—even though both indicate a major problem that needs urgent attention.

[0055] Attributes represent the aspects of the state of entities, and entities of a common class will be characterized by common attributes. For example, with reference to FIG. 3, entities 20,21 in the “GSM switch” class will have a common “GSM switch” attribute.

[0056] Messages impart information about attributes associated with an observed entity. For example, with reference to FIG. 5 the power failure message (40) has a voltage attribute that specifies the current voltage level provided by the power supply entity under the current failure condition.

[0057] The correlation method performed by the network analysis engine 3 simultaneously gathers symptomatic messages under multiple speculative hypotheses. Each hypothesis associates messages under a message signature centred on a candidate entity in the network proposed as the hypothetical root cause.

[0058] Each hypothesis is derived from a signature which is generated, stored and edited by user input, via keyboard 6 (or some other input device). The user generates the signatures by having knowledge and experience of the failure modes of the network. Furthermore, the user may refine the signatures over time by reviewing historical message data and analyzing its associated root causes. The signatures are stored in a signature store 5 shown in FIG. 1.

[0059] FIG. 6 gives an example of a signature 34. The signature comprises three envelope sets 30-32, each associated with a hypothetical root cause 33. In this case the hypothetical root cause 33 is Rack Power Failure. Note that the signature 34 is generic, in the sense that it does not refer to a specific candidate entity within the network. Each envelope set 30-32 comprises a set of proximity envelopes. Each envelope set collates observed messages under a conjunctive orthogonal proximity envelope relative to the hypothetical root cause 33. For example the set 30 comprises six proximity envelopes, each envelope defining a topical, temporal, functional, associative, connective or attributive proximity.

[0060] The stringency of each envelope for associating messages can be as strong or weak as the generality of the signature requires.

[0061] A topical proximity envelope applies a message classification envelope to a symptomatic message. For example the topical proximity envelope for the set 30 defines a "Power Failure Alarm" message class, indicated at 40 in FIG. 5. Any symptomatic messages that fall outside the message classification envelope (in other words any messages which are not "Power Failure Alarm" messages) are excluded from the hypothesis. A topical proximity envelope is able to refer to any entity class within the message class taxonomy shown in FIG. 5.

[0062] Note that the "Power Failure" message class 40 defined by the set 30 is at the lowest level of the taxonomy of FIG. 5, and is a subclass enclosed by the higher level "Failure Alarm" message class 41. In contrast, the "Failure Alarm" message class 41 defined by the sets 31 and 32 is a higher level class which encloses the subclass 40, as well as one other subclass (the "Communications Link Failure" subclass). The ability to define such higher level entity classes enables a user to generate, store and edit generic signatures quickly and intuitively.

[0063] A temporal proximity envelope applies a time envelope to the observation time of a symptomatic message relative to the event time of the hypothetical root cause. For example the temporal proximity envelope for the set 30 defines an observation time of 1-2 seconds relative to the event time of the hypothetical root cause. Any symptomatic messages with an observation time outside the time envelope are excluded from the hypothesis.

[0064] A functional proximity envelope applies an entity classification envelope to the observed entity notifying a

symptomatic message. For example the functional proximity envelope for the set 30 defines a "Monitor" entity class, indicated at 22 in FIG. 3, whereas the functional proximity envelopes for the sets 31 and 32 define a "Component" entity class, indicated at 23 in FIG. 3. Any symptomatic messages for an observed entity that fall outside of the classification envelope are excluded from the hypothesis. A functional proximity envelope is able to refer to any entity class within the entity class taxonomy shown in FIG. 3.

[0065] Note that the "Monitor" entity class 22 defined by the set 30 is at the lowest level of the taxonomy of FIG. 3, and is a subclass enclosed by the higher level "Component" entity class 23. In contrast, the "Component" entity class 23 defined by the sets 31 and 32 is a higher level class which encloses the subclass 22, as well as two other subclasses (the "Power Supply" and "Controller" subclasses). The ability to define such higher level entity classes in the taxonomy enables a user to generate, store and edit generic signatures quickly and intuitively.

[0066] An associative proximity envelope applies an association classification envelope to the associations between the hypothetical root cause entity and the observed entity identified in a symptomatic message. For example the associative proximity envelope for the set 30 defines a "Status Feed" association class, indicated at 25 in FIG. 4. Any symptomatic messages for an observed entity that fall outside of the association classification envelope are excluded from the hypothesis. An associative proximity envelope is able to refer to any association class within the association class taxonomy shown in FIG. 4.

[0067] Note that the association classes defined by the sets 30-32 are all at the lowest level of the taxonomy of FIG. 4. However, alternative signature sets may define higher level association classes such as "Connectional" or "Spatial", enabling a user to generate, store and edit generic signatures quickly and intuitively.

[0068] A connective proximity envelope applies an association distance envelope to the associations between the hypothetical root cause entity and the observed entity identified in a symptomatic message. The concept of "association distance" can be illustrated with reference to FIG. 2. For example the Controller has a "Direct" association with the 3G Nokia Switch, since they have a "data" association linking them directly. By contrast, the Controller has an "Indirect" association with the Monitor, since they are linked indirectly via the 3G Nokia Switch. The connective proximity envelope for the set 30 shown in FIG. 6 defines a "Direct" association between the observed entity and the hypothetical root cause entity. Any symptomatic messages for an observed entity that fall outside of the association distance envelope are excluded from the hypothesis.

[0069] An attributive proximity envelope applies a value envelope to an attribute notified in a symptomatic message. For example the attributive proximity envelope for the set 30 defines an attribute value of "Voltage <5V". Any symptomatic messages with an attribute value that falls outside of the attribute value envelope are excluded from the hypothesis.

[0070] The method of operation of the network analysis engine will now be described.

[0071] When a first message is received by the engine 3 from network 1, the engine 3 looks up the message in the taxonomy store 4 and assigns it a set of classes and subclasses.

[0072] The engine **3** then retrieves all signatures in signature store **5** which have at least one envelope set matching the set of classes and subclasses assigned to the first message. In this step, the engine may retrieve zero signatures, or more likely a large number of signatures. The engine **3** then establishes a hypothesis for each signature, and stores them in an established hypothesis store **8**. Each established hypothesis specifies: a respective root cause (identifying a particular entity and event time); and a signature associated with that root cause.

[0073] For example, if the received message is a power failure alarm message with a voltage attribute <5V, from a monitor with a direct association with a power supply, and a “status feed” association with that power supply, then it matches the envelope set **30** shown in FIG. **6**. The engine therefore establishes a new “Rack Power Failure” hypothesis which identifies a specific power supply entity (for instance the power supply **50** in FIG. **2**) and an event time relative to the observation time of the failure alarm message. In this case the event time will be **2** seconds before the observation time of the failure alarm message.

[0074] Each established hypothesis has a hypothesis score (initially set to zero).

[0075] The engine **3** then compares subsequent messages with all previously established hypotheses in the hypothesis store **8**. If a subsequent message falls within the classes identified by one or more of the envelope sets defined by a hypothesis, then the score of that hypothesis is increased. A message may be matched by more than one envelope, but this is more likely by accident than design, so the message should only contribute towards the score once, even if it matches two envelopes.

[0076] As well as comparing subsequent messages with all established hypotheses, the engine **3** also establishes a new set of hypotheses for each subsequent message having at least one matching signature. A message will not generate a new hypothesis if there are no matching signatures. A significant proportion of messages may never initiate a hypothesis and remain un-correlated.

[0077] Each message gathered under an established hypothesis increases the score of that hypothesis (or in some cases, decreases its score—see the section below describing the elected phase). As the score of a hypothesis changes, the hypothesis evolves through multiple phases. The possible phases for a hypothesis are: the anticipated phase, suspected phase, elected phase, cleared phase, completed phase, suspended phase, confirmed phase, refuted phase and the abandoned phase. Thus, as the score of a hypothesis increases, it moves up from the anticipated phase, to the suspected phase and so on.

[0078] The anticipation phase is an initial phase when there are very few messages associated with the hypothesis and there is little evidence to support the hypothesis.

[0079] The next phase is a suspected phase when there are enough messages gathered under the hypothesis to make it probable.

[0080] The elected phase is when the evidence is strong enough to make the hypothesis definite. At this point, each subsequent message gathered under the newly elected hypothesis is “revoked” for any co-existing un-elected hypothesis that shares the same message. In other words, any co-existing un-elected hypothesis that shares the same message has its score reduced instead of increased.

[0081] Each signature **34** is associated with a particular signature event. For instance, the signature **34** is associated with a rack power failure event. Also, each envelope **30-32** is associated with one or more sub-events, which are related in turn to the signature event. For instance the envelope **30** is associated with a monitor power failure event, occurring 1-2 seconds after the rack power failure event. A monitor power failure alarm (message) will be issued when the monitor power failure (event) occurs. Subsequently, the monitor power failure will be rectified and another message will be issued by the network **1** to the network analysis engine **3**. At this point, the event is “cleared” by the network analysis engine **3**. The cleared phase is when sufficient events have been cleared in an elected hypothesis. Thereafter a cleared hypothesis will only absorb messages related to events it has already cleared and will continue to absorb them until it expires. That is, once a hypothesis is cleared it will only collect messages that clear events it has already correlated. Note that a message may also indicate a change in circumstance for an event—eg change in severity or change in details for a power failure.

[0082] The completed phase is when all significant activity against a hypothesis has lapsed.

[0083] The engine simultaneously evaluates multiple speculative hypotheses, allowing each hypothesis to evolve until a single hypothesis can be resolved and elected as the root cause.

[0084] The suspended phase is when a hypothesis fails to progress to a subsequent phase in the required time. It is suspended pending further activity or until it expires.

[0085] A hypothesis enters the abandoned phase when at least one message matching the hypothesis is revoked by an elected hypothesis.

[0086] The confirmed phase is when a human operator manually confirms (via keyboard **6**) the hypothesis that has yet to become elected and is treated as an elected phase in all other respects.

[0087] The refuted phase is when a human operator manually refutes (via keyboard **6**) the hypothesis irrespective of whether it has become elected and is treated as an abandoned phase in all other respects.

[0088] A hypothesis expires when it completes, or when it has remained suspended and inactive for a sufficient period of time.

[0089] The phases of a hypothesis progress linearly as thresholds are reached. These thresholds may be expressed as: a) volumes of distinct messages, b) distinct events or c) distinct entities, as discussed in detail below.

[0090] a) Number of different messages. In this example, a hypothesis moves into a new phase when the number of messages matching that hypothesis exceeds a threshold (independently of the number of actual events). The threshold can be expressed as:

[0091] i) total number of messages accumulated over the life of the hypothesis;

[0092] ii) average number of messages accumulated per specified time period;

[0093] iii) average number of messages per envelope/event; or

[0094] iv) average number of messages per entity.

[0095] b) Number of separate events. In this example, a hypothesis moves into a new phase when the number of matched sub-events exceeds a threshold, regardless of the number of messages notified about each sub-event.

For example in the case of FIG. 6 there are three sub-events associated with the signature, and each sub-event has a particular envelope 30-32. When a message is received which matches the envelope 30, then that sub-event is “matched”. When the number of matched sub-events (or envelopes) exceeds a threshold then the hypothesis changes phase. The threshold can be expressed as:

- [0096] i) total number of events accumulated over the life of the hypothesis;
- [0097] ii) average number of events accumulated per specified time period;
- [0098] iii) average number of events per entity.
- [0099] c) Number of network entities. Each message is associated with a particular network entity. In this example, a hypothesis changes phase when messages are matched with that hypothesis which are associated with a number of network entities above a threshold. The threshold can be expressed as:
 - [0100] i) total number of entities accumulated over the life of the hypothesis; or
 - [0101] ii) average number of entities accumulated per specified time period.

[0102] Each phase has a timeout such that if it fails to progress to the next phase in that time, the hypothesis can either automatically progress to the next phase or can be suspended until it either expires or new relevant messages are observed later at which point it reverts back to its phase prior to suspension.

[0103] Each hypothesis has a defined severity of critical, major, minor, information or unknown and the severity remains unchanged throughout its lifetime, other than by optional assignment by human operator.

[0104] Each phase of a hypothesis has a defined priority within a bounded range of priority numbers. Each hypothesis has one priority defined. This is used to prioritize hypotheses with the same severity, similar to the message priority described above.

[0105] Every hypothesis notifies of a phase change by outputting one or more messages when the phase change occurs. These messages are of a class nominated by the signature from the complete taxonomy of message classes and can be issued singularly for the phase or in response to each distinct message, event or entity observed during a phase. Each message can be targeted at specific entities or specific entity classes near the candidate root cause entity. The observation time for the notification is the time of the phase change; the observed and observing entities are the correlation system itself; and the event notified is the specific phase change for that specific hypothesis. The severity and priority of the message notified are those of the hypothesis at the time of notification. Such messages may or may not be correlated under other signatures for additional hypotheses.

[0106] Each message collected under one or more hypotheses is only visible to human operators and other network management processes within the context of the hypotheses to which it has been associated.

[0107] Each message that fails to be collected under any hypothesis is made directly visible to human operators and other network management processes.

[0108] The output generated by the process may be displayed on monitor 7, printed, stored or transmitted electronically to a remote client. An example of the output generated by the process is given in Table 1 below:

TABLE 1

Signature	Candidate Root Cause Entity	Severity	Event Time	% Elected
Rack Power Failure	power supply 50	major	12:00:45	75%
Clock Low Signal Level	common clock 51	minor	08:14:26	100%

[0109] Table 1 lists two hypotheses. The first hypothesis is a power failure in the power supply 50 shown in FIG. 2, at 12:00:45. This is categorized as a major incident and is not yet elected. The second hypothesis is a clock failure in the common clock 51 shown in FIG. 2, at 08:14:26. This is categorized as a minor incident and is 100% elected.

[0110] Although the invention has been described above with reference to one or more preferred embodiments, it will be appreciated that various changes or modifications may be made without departing from the scope of the invention as defined in the appended claims.

1. A method of identifying a root cause of an event in a network of linked entities, the method comprising:
 - a. defining a plurality of signatures, each signature comprising one or more classes, the classes being categorized by a taxonomy of classes;
 - b. receiving a first message from the network;
 - c. assigning one or more of the signatures to the first message;
 - d. establishing a set of one or more hypotheses, each hypothesis identifying:
 - i. a respective root cause; and
 - ii. one of the signatures assigned to the first message;
 - e. receiving one or more subsequent messages from the network;
 - f. determining whether each subsequent message matches a signature identified by a hypothesis established in step d; and
 - g. generating an output based on the determination in step f.
2. The method of claim 1 wherein at least one of the signatures comprises a plurality of envelope sets each comprising a set of classes; wherein each envelope set is associated with one or more sub-events related to the root cause; and wherein a signature is assigned to the first message in step c if the first message matches a set of classes in at least one of that signature’s envelope sets.
3. The method of claim 1 wherein at least one of the signatures comprises a message class categorized by a taxonomy of message classes and defining a message topic; and steps c and f comprise matching a message class of the message with the message class of the signature.
4. The method of claim 1 wherein at least one of the signatures comprise one or more entity classes categorized by a taxonomy of entity classes, each entity class defining an entity type; and steps c and f comprise matching an entity class of the message with the entity class of the signature.
5. The method of claim 1 wherein at least one of the signatures comprise one or more association classes categorized by a taxonomy of association classes, each association class defining a type of connection between two directly

communicating entities; and wherein steps c and f comprise matching an association class of the message with the association class of the signature.

6. The method of claim 1 wherein at least one of the signatures includes a class which encloses two or more subclasses categorized by the taxonomy of classes.

7. The method of claim 1 wherein the signatures are defined by a human in step a.

8. The method of claim 1 further comprising giving the one or more hypotheses scores, and amending the scores when the signatures associated with them match subsequent received messages.

9. The method of claim 8 further comprising promoting a hypothesis to an elected phase when its score reaches a pre-determined level, increasing its score when the signature associated with it matches a message, and decreasing the score of any other co-existing hypotheses whose signature also matches the message.

10. The method of claim 1 wherein step f comprises comparing an observation time associated with the subsequent message with an event time associated with the hypothesis established in step d.

11. A method of identifying a root cause of an event in a network of linked entities, the method comprising:

- a. receiving a first message from the network;
- b. establishing a first hypothesis associated with the first message;
- c. receiving one or more subsequent messages from the network;

d. for each subsequent message:

- i. establishing a new hypothesis, and;
- ii. determining whether the subsequent message matches any previously established hypothesis by comparing an observation time associated with the message with an event time associated with the hypothesis; and

e. generating an output based on the determination in step d.ii.

12. The method of claim 11 further comprising defining a plurality of signatures, wherein any subsequent message which does not match one of the signatures does not establish a new hypothesis.

13. The method of claim 11 further comprising presenting the output to a user.

14. A network analysis engine configured to identify a root cause of an event in a network of linked entities by the method of claim 1.

15. A program or suite of programs for controlling the operating of a network analysis engine, the program(s) being arranged when executed to control the network analysis engine so as to operate in accordance with the method of claim 1.

16. A network analysis engine configured to identify a root cause of an event in a network of linked entities by the method of claim 11.

17. A program or suite of programs for controlling the operating of a network analysis engine, the program(s) being arranged when executed to control the network analysis engine so as to operate in accordance with the method of claim 11.

* * * * *