(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
3 June 2010 (03.06.2010)

PCT

(10) International Publication Number
WO 2010/062435 A1

(51) International Patent Classification:
G06F 11/00 (2006.01)

(21) International Application Number:
PCT/US2009/055958

(22) International Filing Date:
4 September 2009 (04.09.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/094,272 4 September 2008 (04.09.2008) US

(71) Applicant *(for all designated States except US)*: **TELCORDIA TECHNOLOGIES, INC.** [US/US]; One Telcordia Drive 5G116, Piscataway, NJ 08854-4157 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **BAKER, Donald** [US/US]; 8017 Baxter Springs Road, Austin, TX 78745 (US). **NODINE, Marian** [US/US]; 1800 Lemon Mint Court, Austin, TX 78733 (US).

(74) Agents: **FEIG, Phillip, J.** et al.; Telcordia Technologies, Inc., One Telcordia Drive 5G116, Piscataway, NJ 08854-4157 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

[Continued on next page]

(54) Title: COMPUTING DIAGNOSTIC EXPLANATIONS OF NETWORK FAULTS FROM MONITORING DATA

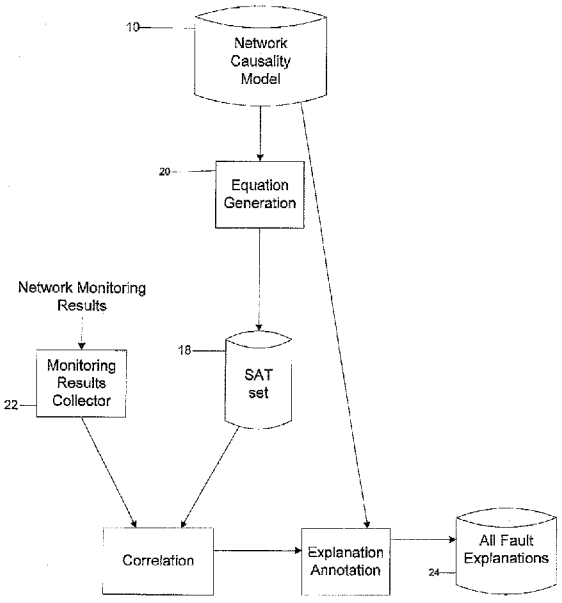

FIGURE 6

(57) Abstract: A system and method for network fault diagnosis in a network having network elements is presented The method comprises creating a network causality model, generating Boolean expressions from the network causality model, converting the Boolean expressions into SAT sets, receiving network monitoring results, correlating these monitoring results with the SAT sets, and enumerating all possible diagnostic explanations of potential faults, properly annotated. Creating a network causality model can comprise creating, for each network element, an element-specific causality model, stitching together all network elements using the element-specific causality models and a network topology, retrieving monitoring state and propagation information, and generating the network causality model using the stitched together network elements and the monitoring state and propagation information. Stitching together network elements can comprise adding causes and implies dependency between appropriate network elements and/or adding and connecting reachable and not-reachable states.

WO 2010/062435 A1

Published:

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

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

COMPUTING DIAGNOSTIC EXPLANATIONS OF NETWORK FAULTS FROM MONITORING DATA

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] The present invention claims the benefit of U.S. provisional patent application 61/094,272 filed September 4, 2008, the entire contents and disclosure of which are incorporated herein by reference as if fully set forth herein.

GOVERNMENT LICENSE RIGHTS

[0002] This invention was made with Government support under DAAE07-03-9-F001 awarded by the Department of the Army. The Government has certain rights in this invention.

FIELD OF THE INVENTION

[0003] The present invention relates generally to network management and to network fault diagnosis.

BACKGROUND OF THE INVENTION

[0004] Diagnosing the causes of network faults is the first step in their correction, either by restoring the network to full operation or by reconfiguring the network to mitigate the impact of a fault. Fault diagnosis is relatively straightforward if one has information about the state of all active elements in the network. Unfortunately, such monitoring information is usually transmitted through the network itself to a network management system (NMS) that processes the information for diagnosis. Thus, the existence of a fault can impede the gathering of information about the fault and other faults in the network. This problem can be reduced, but not solved, by distributing NMS functionality through the network. While distributing fault diagnosis among multiple NMSs allows for multiple perspectives, coordination between the NMSs also can be impacted by the same network faults they are trying to diagnose and repair. However, distributing the NMS capabilities also allows for various domains of the network to be managed autonomously, thus avoiding the problem of a single point of failure.

[0005] The most common approach to network fault diagnosis is known as the fault propagation approach, which leverages a model of how faults propagate through a network. For example, the failure of a network interface can effectively sever communication to its entire device, thus creating the appearance of secondary faults in the device. A fault propagation model of a complete network is often constructed from the modeled behavior of network elements and the network layout. Once constructed, the

propagation model is used during live fault diagnosis to reason about the network monitoring data that are available to an NMS, such as SNMP queries and traps, periodic polling, and other mechanisms that allow monitoring of the state of various network elements. Based on the results of the reasoning over the fault propagation model, the fault is localized.

[0006] Various approaches to solve the diagnosis problem have included expert systems, neural networks, fuzzy logic, max product algorithm, petri-nets, and so on. A number of groups have used variants of the fault propagation approach with a variety of model types including dependency graphs, causality graphs, coding approaches, Bayesian reasoning, and even phase structured grammars. Also, Boolean variables have been used, but they focus primarily on dealing with reasoning about reachability from the observer, i.e., the NMS. However, these techniques are quite complex. A simpler fault propagation approach is needed, such as one that enables enumeration of all possible alternative diagnostic explanations of possible faults that would give rise to the communication network monitoring results being considered.

SUMMARY OF THE INVENTION

[0007] An inventive method and apparatus for determining and enumerating all possible/valid alternative diagnostic explanations for the root cause(s) of communication network faults along with their relative likelihoods given appropriate network and device models and network monitoring results is presented. The problem is to determine the root causes of communication network faults based on network monitoring results, either given as a snapshot of the network state or as a continuous stream. Given the network monitoring results, an inventive method and apparatus diagnose these monitoring results and enumerate all possible combinations of root causes that could give rise to the communication network faults, additionally determining the relative likelihoods of those explanations. In the case of continuous operation, the resulting explanations are periodically revised.

[0008] Note that fault determination is the first step in recovering from the fault, which is outside the scope of this invention. Additionally, the proposed solution is not fully general in that it does not deal with the issue of possibly faulty monitoring results.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The invention is further described in the detailed description that follows, by reference to the noted drawings by way of non-limiting illustrative embodiments of the invention, in which like reference numerals represent similar parts throughout the drawings. As should be understood, however, the invention is not limited to the precise arrangements and instrumentalities shown. In the drawings:

Figure 1 is illustrates the flow of the pre-processing performed in accordance with the inventive system;

Figure 2 is an example causality model for a switch in accordance with the inventive system;

Figure 3 illustrates stitching network element causality models in accordance with inventive system;

Figure 4 illustrates stitching for network connectivity in accordance with inventive system;

Figure 5 is a flow diagram of pre-processing;

Figure 6 illustrates the flow of the run time processing performed in accordance with the inventive system; and

Figure 7 is a flow diagram of run-time processing.

DETAILED DESCRIPTION

[0010] Diagnosis of network faults is a problem of reasoning in the face of missing information concerning the state of the network. With incomplete information, a network fault diagnosis is likely to have some ambiguity, which leads to the secondary challenge of how to react appropriately in the face of the ambiguity. A proper network fault diagnosis capability must therefore have two requirements. First, the diagnoses must be *complete* with respect to the NMS's knowledge, meaning that the diagnoses must accurately describe the space of possibilities of all root causes that lead to the network state known to the NMS and its understanding of the network. Most known network fault diagnosis mechanisms produce either the single most likely cause or the single simplest cause that might explain the observed network state. While this approach might arrive at the correct diagnosis in many cases, it will necessarily give incorrect results in more complex cases. The ambiguity of the information available plays into this complexity, so an increase in ambiguity generally results in a greater likelihood of incorrect diagnosis in existing, known fault diagnosis system.

[0011] Complete diagnostic information, while valuable, is unwieldy. Therefore, the second requirement in dealing with complete but ambiguous diagnoses is to express the ambiguity in a useful, actionable representation. To address this issue, network fault diagnoses can be represented in terms of a number of competing independent diagnostic explanations, where each explanation provides a possible set of root causes leading to the observed network state and the entire set of explanations cover all possible failures. In such a representation, each root cause failure may be associated with a relative likelihood of occurrence.

[0012] The NMS computes a failure rate, such as the probability of a failure in an hour of operation, for each independent diagnostic explanation, such as those described above. The failure rate for an explanation is based on failure rate information for known possible root causes, which are assumed to be independent. The NMS uses this information to compute the relative likelihood of each explanation and to rank the explanations by likelihood. Finally, the NMS projects alternative diagnostic explanations onto the network layout so that alternative explanations of root causes can be computed per network element or element locale. Such element-centric explanations are useful in determining where attention should be focused and in creating actionable plans for addressing the root cause failures.

[0013] Another essential element for diagnosing network faults is updating the set of diagnostic explanations as the network state changes over time. Fault diagnosis must therefore be performed in a continuous fashion, possibly revising diagnoses as new network monitoring data becomes available.

[0014] The classic Boolean satisfiability (SAT) problem is simply stated: given a set of Boolean expressions over a number of Boolean variables, discover a truth assignment to those variables that will allow all the expressions to be satisfied, e.g. evaluate to true, or alternatively, determine that no such solution exists. For example, given the expressions "x or y" and "x $\neq$y", the assignment x=false, y=true satisfies both expressions. While solving the general Boolean satisfiability problem is inherently expensive, a great deal of research has enabled the construction of SAT solvers, which solve this class of problems, to be as efficient as possible.

[0015] For network fault diagnosis, a variant of the Boolean satisfiability can be used which, instead of finding a solution to a given satisfiability, determines the space of all possible solutions, i.e. all possible assignments that satisfy the equations. In the problem

above, the complete set of solutions would be {(x=false, y=true), (x=true, y=false)}. Many of the techniques employed within SAT solvers can also be applied to this related problem.

[0016] In order to reduce the network fault diagnosis problem to a problem involving Boolean expressions, an artifact called a *causality model*, that is, a fault propagation model that represents the behavior of network elements in terms of Boolean variables with constraints represented as Boolean expressions, can be used. Network monitoring information is cast into assignments to observable variables, that is, Boolean variables in the causality model that are associated with aspects of network elements that can be observed by the NMS.

[0017] For example, there might be an observable variable that indicates that the most recent poll of a particular switch interface timed out, therefore the interface is unreachable from the NMS. Each resulting diagnostic explanation contains assertions about whether each possible root cause failure is deemed to have occurred or not, or whether the root cause is indeterminate within that explanation. Each root cause variable in the model attempts to capture an independent cause of network device non-normality. A root cause variable, for example, would indicate whether the network cable has become disconnected from a particular interface on a particular network switch. While root causes are independent sources of failures, each root cause variable is related to internal state variables within the model that attempt to capture the network behavior implications of a root cause's occurrence or non-occurrence. Internal variables do not directly indicate observable or causal states in the network.

[0018] The causality model captures the salient aspects of the network's behavior in terms of Boolean variables representing primitive aspects of the state of devices within the network. The causality model is a design artifact that is constructed once per network, during pre-processing, and reused at run time each time faults are diagnosed on that network. Not surprisingly, the quality of the model has a large impact on the quality of the diagnostic explanations it generates. Issues surrounding the desired granularity of the causality model are discussed below.

[0019] The causality model for the network can be constructed automatically from element-specific causality model fragments for network elements that are stitched together, much like parts of a quilt, based on the connections between those elements in

the network layout. This overall approach can be generalized to include other aspects of the network layout that can provide observations to the NMS.

[0020] The causality model for a network element is comprised of Boolean state variables linked by Boolean dependencies that capture constraints among the states associated with the variables. State variables model well-defined aspects of the network element's behavior, whether, for example, the element has power supplied, the element is up, or the element has the ability to respond to an SNMP request. Each network element will also have root cause variables indicating externally-caused failures, such as whether a device has failed, whether it is off, etc. Ultimately, within each network element causality model, these root cause variables must be related directly or indirectly to internal variables that represent behaviors of a network element that are possible to observe or that have some impact on other network elements. An interface on a switch, for example, may be non-communicative (due to a variety of causes), which will have an impact on the network connectivity to physically connected devices. The causality model mechanism provides flexibility as to the level of detail of the modeling for a device.

[0021] Figure 1 shows a flow of pre-processing in accordance with the inventive system and method. As shown in Figure 1, in pre-processing, three models are used to construct the Network Causality Model **10**. The Communication Network Topology **12** records the devices used in the communication network, their types, and their connectivity. The Communication Network Topology **12** also records the placement of the NMS where the Correlation Process will occur. The Device Type System **14** records the possible types of devices, their characteristics, and behaviors. Device behaviors are captured as possible device states, the relationships between those states ("causes", "implies", "mutually exclusive", "or", "and", etc.), the independent root causes of failures and the mean time between failures (MTBF) of those root causes. The Monitoring Results Catalog **16** records the states of particular devices in the network that can be monitored by the NMS, and how long each such monitoring result takes to propagate to the network management system. The Causality Model Generation step uses the Communication Network Topology **12**, Device Type System **14**, and Monitoring Results Catalog **16** to construct a Causality Model **10**.

[0022] Accordingly, the Causality Model **10** contains the states for all devices in the network, all possible root cause failures, and possible observations of network state through monitoring results. These states are interconnected with relationships described in

the Device Type System **14** and generated based on the network connectivity as described in the Communication Network Topology **12** and Monitoring Results Catalog **16**.

[0023] Figure 2 shows a simple causality model for a network switch. Each switch interface would be considered a separate component with its own causality model. In Figure 2, "Switch Fail" is the only root cause of failures for the switch. The "xor" and "implies" dependencies have the standard logical meaning. The causes dependencies are converted into a term of an "or" expression, thus "Switch Down" is equivalent to "Switch Fail" or "Switch Unpowered." The states "Switch Up" and "Switch Down" might be observable in some network configurations.

[0024] Combining causality models of network elements can be automated readily if network element types and their causality model fragments are organized into a taxonomy and the possible connectivity relationships between network element types are properly characterized. This taxonomy can be called the Device Type System **14** or *element type catalog*. In practice, only a small number of such relationship types are needed in the catalog, such as: a power supplier "supplies power" to a power consumer; a component may "require container" to be up for the component to be up; or two 100BaseT network devices/cable ends are related by the "connected" relationship. The presence of a relationship in the network layout for a given relationship type indicates how the causality model fragments will be combined into the causality model for the entire network.

[0025] In constructing the causality model for a complete network, as shown in Figure 1 for example, each network element has instantiated for it the causality model for its type, thus creating an element-specific causality model fragment. This construction proceeds by adding dependencies among related state variables across connected network elements, a process called "stitching." For example, if a container supplies power to its parts, such as in the case of a host computer supplying power to its network interfaces, the container's "does supply power" variable will have a dependency with the component's "does receive power" variable with a Boolean equality dependency (assuming the simplest case). In the case of network reachability from the NMS, a more complex stitching mechanism will create observable variables concerning the reachability of the network element from the NMS based on variables for whether the particular element can communicate and whether the next closer network element to the NMS on the communication path is reachable. This technique can be generalized to network topologies with loops and redundant components.

[0026] Figure 3 shows an example of stitching the network element causality models for a switch and its interface which are elements related in such a way that the interface requires its containing switch to be up for the interface to be up. Stitching in this case involves adding a causes dependency and an "implies" dependency between the appropriate variables in the respective network element causality models.

[0027] Stitching related to network connectivity is more complex because reachable and not-reachable states (from the NMS) must be added and connected based on the network layout. In Figure 4, the causality models of a host interface, cable, switch interface, and switch are stitched based on the network layout and the location of the NMS.

[0028] The final step in generating the causality model for the entire network is the creation of observable variables for aspects of the network that will actually be observed by the NMS. For example, while it is possible to observe network reachability for a particular switch interface from the NMS, if the NMS does not poll that interface, no such observation will be generated. Whether a possible observable state in a causality model actually generates an observation to the NMS depends on the actual configuration of the network element and the NMS in particular. The network causality model **10** also keeps track of upper bounds on the time it takes for the network state change to be noted by the NMS. This latency information is needed to create a practical fault diagnosis component that operates continuously.

[0029] Hence, the network causality model **10** is rapidly generated during pre-processing from the network layout or Communication Network Topology **12**, the element type catalog or Device Type System **14**, and information about the NMS configuration, such as the Monitoring Results Catalog **16**. Thus the network causality model **10** can be treated as an internal data structure to the NMS that is constructed at NMS start up or when the network topology is known to have changed.

[0030] Figure 5 is a flow diagram of the pre-processing steps including generation of the network causality model **10**. There are two parts to the start-up or pre-processing; the network causality model **10** can be truly generated as a separate process (steps S1-S5) and saved to be read at run-time. The other steps (S5, S6) must be done during the initialization of the correlation process. In step S1, a causality model is generated for each network element. In step S2, these causality models are stitched together using the topology and placement information, obtained, for example from the Communication Network Topology **12**. In step S3, information regarding the states of the network

elements that can be monitored by the network management system, and how long each such monitoring result takes to propagate to the network management system is retrieved, for example from the Monitoring Results Catalog **16**. In step S4, population of a graph or model of combined network elements with monitoring states and propagation information occurs. In step S5, the network causality model **10** is generated based on the information obtained in steps S1 through S4. In step S6, Boolean expressions **20** are generated based on the network causality model **10**. In step S7, the Boolean expressions **20** are converted into SAT sets **18**.

[0031] One example of this approach is a model of a simple network with a single switch, 6 hosts with applications running on the hosts, necessary network connectors, and a power supply. The resulting causality model can contain around 650 states, including 57 root cause variables and 64 observable variables, with around 1,100 dependencies.

[0032] As shown in Figure 6, the correlation process occurs within the NMS, the Causality Model **10** generated during pre-processing can be converted into a SAT Set **18** by generating Boolean expressions **20** from the Causality Model **10**. The SAT Set **18** represents the truth value of every state in the Causality Model **10**, for example as true/false/unknown, and the Boolean conditions that must hold between such states. Expressions within the SAT Set **18** are represented in conjuctive normal form. At run-time, as Network Monitoring Results as described by the Monitoring Results Catalog **16** arrive, they are collected by the Monitoring Results Collector **22** until the collector determines that it is likely to have all results relating to a particular network fault event. Once that has occurred, the Monitoring Results Collector **22** sends the collected monitoring results to the Correlation step where the monitoring results are applied to the SAT Set **18**, yielding a simpler SAT Set (not shown) where more states are known and fewer constraints exist.

[0033] The SAT Set is further simplified to remove mention of non-root-cause states in the constrained equations. The remaining equations involve only root cause states. The simplified SAT Set expressions are then converted into disjunctive normal form (DNF). DNF represents a union (disjunction) of alterative clauses where each alternative makes a determination about each root cause: whether it has occurred, has not occurred, or is unknown. From the combined expression, each disjunction corresponds to a distinct possible explanation of the faults in the network. These possible explanations are then sent to the Explanation Annotation step where the mean time between failures of each

such explanation can be computed from the root cause failures in each. Once the mean time between failures of each explanation is computed, the likelihood of each alternative explanation can be readily determined by proportionate weighting by failure rate of each explanation. The Explanation Annotation step also annotates each explanation with relevant information about devices and failure states contained in the Causality Model **10** but not needed in the earlier steps. The resulting Fault Explanations **24** enumerate all possible combinations of faults that will result in the observations seen by the network management system and the likelihood of each explanation. The result is suitable for use by human operators or automatic fault mitigation mechanisms to address the possible failures.

[0034] A flow diagram of the run time procedure is shown in Figure 7. In step S8, which can occur prior to or simultaneously with steps S6 and/or S7 shown in Figure 5, network monitoring results are received and collected in the Monitoring Results Collector **22**. In step S9, the monitoring results are correlated with the SAT sets **18**. In step S10, explanations are annotated by combining the correlated information and the network causality model **10**. Step S11 performs enumerating all possible diagnostic explanations of potential faults, properly annotated. Hence, all fault explanations are produced and output as a ranked fault list. The output can be on a computer monitor or other computer or display device.

[0035] Hence, once the network causality model **10** has been created, it is treated as a system of Boolean expressions over Boolean variables to be satisfied, that is, an instance of a Boolean satisfiability problem. Internal state variables in the expressions are removed without impacting the remaining variables by applying a series of techniques, such as substituting an internal variable with an equivalent expression. This simplified SAT set is created during NMS initialization and copied for reuse each time explanations are generated from a new set of observations. When simplified, the causality model in the above example of a network with a single switch results in a SAT set with 120 variables and 230 elementary constraints. Note that simplification has the side effect of removing redundancy from the causality model.

[0036] At any given time, the NMS can generate an explanation for network failures based on the most recent values for the observations known to the NMS. Each such observation creates a Boolean binding to the corresponding observable variable which can then be applied to a copy of the simplified SAT set. The variable elimination techniques

previously described are then applied to remove any unbound observable variables (corresponding to unknown observations). The remaining reduced SAT set describes a set of constraints over root cause variables only and is a compact representation of all possible diagnostic causes given the observations applied.

[0037] To generate alternative diagnostic explanations, the NMS converts the SAT set into an equivalent Boolean expression in DNF. Each DNF clause corresponds to an alternative hypothesis over root causes that explain the observations seen. The resulting hypothesis set is complete with respect to the causality model, in that no other explanations are possible. Thus, the Boolean satisfiability approach leads to a complete set of alternative diagnostic explanations for the network monitoring data seen by the NMS.

[0038] Exemplary network monitoring results can be:

- Interface 1 of switch 3 is not reachable.
- Interface 2 of switch 3 is reachable,
- Interface 3 of switch 3 is reachable.
- Interface 1 of switch 3's administrative status is up
- Host 5 is not reachable
- Interface 1 of host 5 is not reachable
- Interface 2 of host 5 is not reachable.
- ...
- Interface 1 of switch 3's operational status is unknown.
- The web application running on host 5 is answering queries.
- ...
- The DHCP service on host 12 is not answering queries

[0039] Many devices, network elements, etc. can be monitored. The raw data will have a large stream of information, with large swaths of things that may be unknown or non-nominal because of one or more failures.

[0040] Exemplary output of the inventive system can be:

- Explanation 1 (75% likelihood)
    - The cable between host 1 and host 2 is disconnected and
    - The DHCP service on host 12 is down
- Explanation 2 (25% likelihood)
    - Interface 1 of host 3 is down and
    - The DHCP service on host 12 is down

[0041] In this situation, there are two competing explanations that might independently explain all of the monitoring data seen. Of the two explanations, the first one is the most likely. Each explanation has two independent failures that have both occurred. Because the DHCP failure appears in all explanations, its likelihood is 100%. That is, it is known that one thing failed but the other source of problem has some ambiguity. These explanations can be provided in various formats.

[0042] It is possible to have a single explanation indicating that nothing is wrong. Generally, however, there will be some number of explanations, each with a list of independent root causes or real-world problems that can then be remedied. It is important to recognize that based on this output, there are no other combination(s) of failures that would explain the given network monitoring data.

[0043] In the above example of a network with a single switch, single failures actually existing in the simulated, single switch network give rise to observations presented to the NMS that result in one to three hypotheses, with more showing up only in the most unusual cases. An implementation written in the Java programming language can generate explanations for that network in around 100ms running on a personal computer. These results illustrate the reasonable efficiency offered by the inventive Boolean satisfiability approach.

[0044] To make the diagnostic explanations more useful, the NMS associates each explanation with a relative likelihood so that corrective actions can be taken in likelihood order. The relative likelihood of a hypothesis in a hypothesis set is directly computed if the failure rate for each hypothesis is known. Specifically, the likelihood of a hypothesis is the failure rate of the hypothesis divided by the sum of the failure rates of all hypotheses in the set. The failure rate of a hypothesis is the product of the failure rates of the root causes in the hypothesis known to have occurred. Even if failure rates for root causes are estimated, the likelihood computation can still yield useful results because the relative ordering of hypotheses is not sensitive to such errors. Note that, if some failure is directly observed by the NMS without ambiguity, the explanations produced will all have some root cause failure. The alternative, which is that no failure is being observed, will produce a single explanation that all is well. Any single explanation will be assigned a likelihood of 1.

[0045] A number of issues arise when using this approach in the implementation of a practical system. These include model granularity, continuous operation of the NMS, dynamic network, acting on diagnostic explanations, and distributing the algorithm across

NMSs. The inventive system and method deals with these issues as it applied to a variety of simulated networks.

[0046] As with all model-based approaches, there is an inherent tradeoff to be made between the accuracy of the model and its complexity. In practical situations, the causality model can be designed to have the appropriate level of detail necessary for the application at hand. For example, at design time, root causes can be related to actions that might be available to remedy those causes. It does not make sense to model fine detail of device states when remedying actions might only include actions such as restarting or replacing the network element. A related issue is the variety of types of observations available to distinguish independent root causes. To aid with this design issue, it is possible to use the SAT set **18** to identify root causes that are indistinguishable by the observations available to the NMS. Indistinguishable root causes have the undesirable effect of creating alternative hypotheses when the failure is indicated in the explanations. In one embodiment, indistinguishable causes can be combined into a compound cause that stands for either.

[0047] How to generate explanations for a given snapshot of network monitoring data known to an NMS is described above. Observations related to a particular root cause may arrive at the NMS over a relatively long time period. If one assumes that network state changes occur relatively infrequently, the observation latency information described above can be used to estimate a window in which the root cause state change that ultimately caused the observation is likely to have occurred. The longest observational latency in the network causality model **10** can therefore be used to compute an upper bound on the expected arrival time of the last observation caused by a root cause state change. Thus, the NMS can predict a time at which it is likely to have seen all observations related to a root cause state change and therefore an appropriate time to generate a new set of diagnostic explanations.

[0048] Two situations must be accounted for in the above approach. First, it is possible to have multiple root causes occurring close enough in time so that the NMS has an inconsistent snapshot of the network state. In such a situation, the observations will most likely be inconsistent with the assumptions embodied in the network causality model **10** and the simplified SAT set will be reduced to an infeasible solution. In practice, this is not a serious issue as it is an indication that more observations are expected with the explanations generated again after a suitable delay.

[0049] The second troublesome issue has to do with transient observation changes. In one embodiment, an observation preprocessing module that either filters or marks transient observations is used. Transient observations can be related to the appropriate variables in a network element causality model **10** just as any other observation. The details of these steps depend, of course, on the specific observation being considered.

[0050] Fault diagnosis in dynamic is challenging because the NMS makes sense of the network in the face of rapid change. The inventive system and method can be readily used with dynamic networks if the changes to the network layout and configuration are known to the NMS along with relevant information about the timing of the reconfiguration. Such information might be given to the NMS in the form of notifications from the network's configuration management mechanism. Given these simplifying assumptions, diagnosing faults in a dynamic network requires that the NMS stop generating explanations during the network reconfiguration and recreate the causality model based on the new configuration once it is known. Fortunately, recreating the causality model is a straightforward process as was described above.

[0051] Upon completion of the network reconfiguration, a new set of diagnostic explanations must be generated, but the most recent observations of the network state cannot necessarily be reused. This is because many observable variables in a causality model capture topology-dependent information about the network. For example, an observation about the reachability of a network component from the NMS, e.g. via polling, has something to say about all components on the network paths to the destination. Thus, when the topology changes, the simplest approach is to drop any lingering topology-dependent observations and reacquire them in the new network configuration, possibly by actively probing the network elements involved. Other analytical approaches to inferring new observations from old observations using the old and new causality models are impractical, given the complexity involved and the fact that acquiring new observations is relatively inexpensive.

While the details of acting on diagnostic explanations is outside the scope of this disclosure, it is important to emphasize the fact that actionable explanations go a long way toward the stated goal of providing useful explanations. Diagnostic explanations become actionable in two ways, either by resolving ambiguity or by taking corrective action. Resolving ambiguity is greatly aided by having a complete set of diagnostic explanations. It is straightforward to compare alternative explanations to discover the nature of the

ambiguity. In many cases a set of conflicting causes can be associated with a probing action that will distinguish among them by initiating the gathering of new information. Ideally, the probe should generate new observations that would be considered by the next generation of explanations.

[0052] The diagnostic explanations generated in accordance with the inventive techniques are also useful for corrective action because the explanations are partitioned by locales in the network such as all the causes related to a switch, its interfaces, and its connections. A set of diagnostic explanations can thus be created for each locale, enabling a human or automated actor to focus attention on the area(s) of the network that need corrective action, perhaps performing corrective action for different locales in parallel.

[0053] The inventive system and method can be used in actions related to resolving ambiguity or resolving the faults and in environments with continuous operation. Furthermore, the methodology can be adapted to a wide variety of fault diagnosis applications where the system under consideration has a static topology and discrete behaviors. The inventive mechanisms for distribution are well suited to many military networks with locally-managed autonomous domains. This approach may be applicable to dynamic networks such as mobile ad hoc networks (MANETS) where network element behavior has an inherently analog nature and the distinction between failures and network topology changes need to be inferred by the NMS.

[0054] The inventive system and method provides numerous advantages over previous solutions. For example, the use of a variant of Boolean Satisfiability enables the space of all possible solutions to be efficiently represented in accordance with the classic Boolean Satisfiability approach which is to find any one single solution that satisfies the Boolean expressions. This approach allows the rapid enumeration of all possible network failure modes. The invention advantageously gives all possible explanations for the observations seen, enhancing network fault diagnosis because sometimes the actual failure is one that is unlikely and thus might be omitted from an incomplete list. Previous approaches generate a single, presumed "best ", explanation for the observations.

[0055] In addition, the invention is tolerant of incomplete information, such as from missing monitoring results, which sometimes occur in practical networks. In such a case, the invention may generate more explanations or explanations with fewer definitive failure assessments. Previous methods produce no results in the face of missing information.

[0056] Further, the invention can be extended to include states and monitoring results that go beyond network elements, thus extending the solution to the network and its greater environment. For example, the state of power components can also be represented and reasoned about.

[0057] The invention allows flexibility in what device states are to be considered important by the designer and what faults are to be considered. Moreover, the models that must be generated by a designer, e.g. Communication Network Topology **12**, Device Type System **14**, and Monitoring Results Catalog **16** shown in Figure 1, are simple to specify and do not require knowledge of the inventive correlation process.

[0058] Various aspects of the present disclosure may be embodied as a program, software, or computer instructions embodied in a computer or machine usable or readable medium, which causes the computer or machine to perform the steps of the method when executed on the computer, processor, and/or machine. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform various functionalities and methods described in the present disclosure is also provided.

[0059] The system and method of the present disclosure may be implemented and run on a general-purpose computer or special-purpose computer system. The computer system may be any type of known or will be known systems and may typically include a processor, memory device, a storage device, input/output devices, internal buses, and/or a communications interface for communicating with other computer systems in conjunction with communication hardware and software, etc.

[0060] The terms "computer system" and "computer network" as may be used in the present application may include a variety of combinations of fixed and/or portable computer hardware, software, peripherals, and storage devices. The computer system may include a plurality of individual components that are networked or otherwise linked to perform collaboratively, or may include one or more stand-alone components. The hardware and software components of the computer system of the present application may include and may be included within fixed and portable devices such as desktop, laptop, and server. A module may be a component of a device, software, program, or system that implements some "functionality", which can be embodied as software, hardware, firmware, electronic circuitry, or etc.

[0061] The embodiments described above are illustrative examples and it should not be construed that the present invention is limited to these particular embodiments. Thus, various changes and modifications may be effected by one skilled in the art without departing from the spirit or scope of the invention as defined in the appended claims.

What is claimed is:

1. A method for network fault diagnosis in a network having a plurality of network elements, comprising steps of:

creating a network causality model;

generating Boolean expressions from said network causality model;

converting said Boolean expressions into SAT sets;

receiving network monitoring results;

correlating said network monitoring results with said SAT sets; and

enumerating all possible diagnostic explanations of potential faults and displaying a ranked fault list having said explanations with annotations on a computer.

2. The method according to claim 1, the step of creating a network causality model comprising steps of:

for each network element of the plurality of network elements, creating an element-specific causality model;

stitching together all network elements using the plurality of element-specific causality models and a network topology;

retrieving monitoring state and propagation information; and

generating the network causality model using the stitched together network elements and the monitoring state and propagation information.

3. The method according to claim 2, the step of stitching together network elements comprises at least one of:

adding a causes dependency and an implies dependency between appropriate network elements; and

adding and connecting reachable and not-reachable states.

4. The method according to claim 2, wherein the element-specific causality model instantiates a type of network element.

5. The method according to claim 2, wherein the plurality of network elements and element-specific causality models are organized into a taxonomy and categorized.

6. The method according to claim 1, wherein the network causality model comprises states for all of the network elements.

7. A computer readable medium having computer readable program for operating on a computer for network fault diagnosis in a network having a plurality of network elements, said program comprising instructions that cause the computer to perform steps of:

creating a network causality model;

generating Boolean expressions from said network causality model;

converting said Boolean expressions into SAT sets;

receiving network monitoring results;

correlating said network monitoring results with said SAT sets; and

enumerating all possible diagnostic explanations of potential faults and displaying a ranked fault list having said explanations with annotations on a computer.

8. The computer program according to claim 7, the step of creating a network causality model comprising steps of:

for each network element of the plurality of network elements, creating an element-specific causality model;

stitching together all network elements using the plurality of element-specific causality models and a network topology;

retrieving monitoring state and propagation information; and

generating the network causality model using the stitched together network elements and the monitoring state and propagation information.

9. The computer program according to claim 8, the step of stitching together network elements comprises at least one of:

adding a causes dependency and an implies dependency between appropriate network elements; and

adding and connecting reachable and not-reachable states.

10. The computer program according to claim 8, wherein the element-specific causality model instantiates a type of network element.

11. The computer program according to claim 8, wherein the plurality of network elements and element-specific causality models are organized into a taxonomy and categorized.

12. The computer program according to claim 7, wherein the network causality model comprises states for all of the network elements.

13. A system for network fault diagnosis in a network having a plurality of network elements, comprising:

a network causality model;

SAT sets; and

network monitoring results, wherein the network causality model generates the SAT sets using Boolean expressions, the network monitoring results and the SAT sets are correlated to enumerate all possible diagnostic explanations of potential faults, and a ranked fault list having said explanations with annotations are displayed on a computer.

14. The system according to claim 14, further comprising:

a plurality of element-specific causality models, each element-specific causality model created from one of the plurality of network elements;

a network topology; and

monitoring state and propagation information, wherein the plurality of network elements are stitched together using the plurality of element-specific causality models and the network topology, and the network causality model is generated using the stitched together network elements and the monitoring state and propagation information.

15. The system according to claim 14, wherein the stitching together of the network elements comprises at least one of:

adding a causes dependency and an implies dependency between appropriate network elements; and

adding and connecting reachable and not-reachable states.

16. The system according to claim 14, wherein the element-specific causality model instantiates a type of network element.

17. The system according to claim 14, wherein the plurality of network elements and element-specific causality models are organized into a taxonomy and categorized.

18. The system according to claim 13, wherein the network causality model comprises states for all of the network elements.

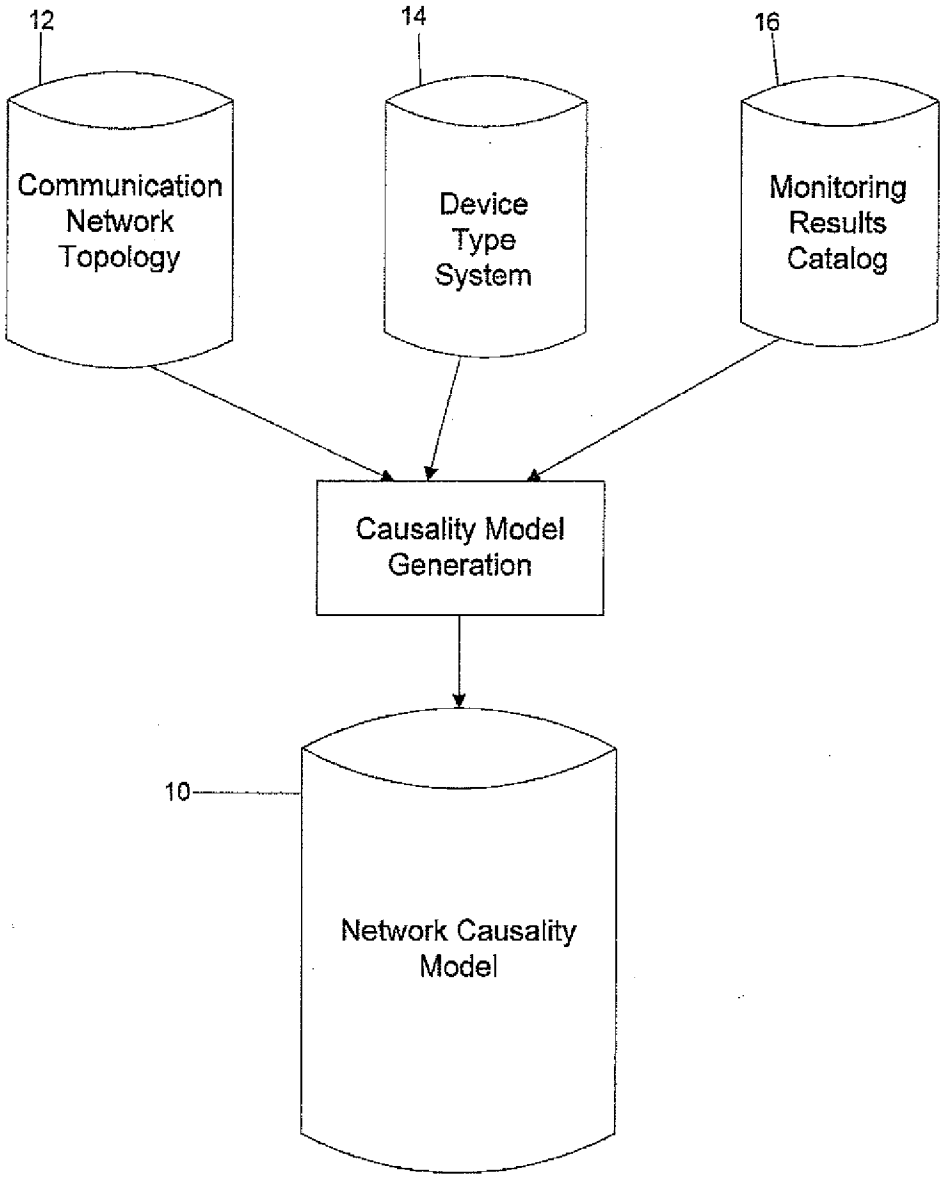
12
14
16
Communication Network Topology
Device Type System
Monitoring Results Catalog
Causality Model Generation
10
Network Causality Model

FIGURE 1

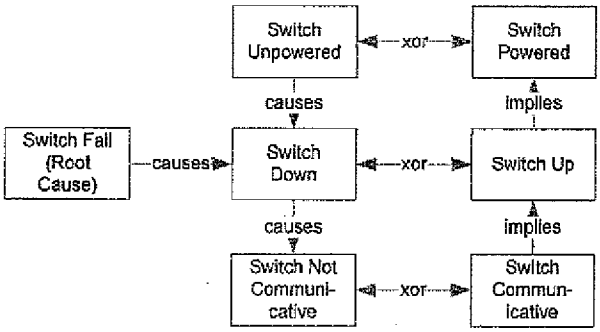

FIGURE 2

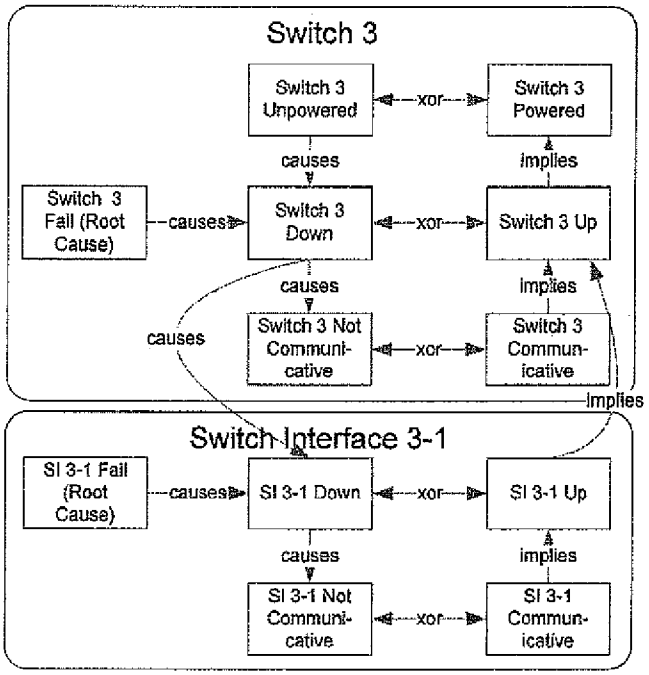

FIGURE 3

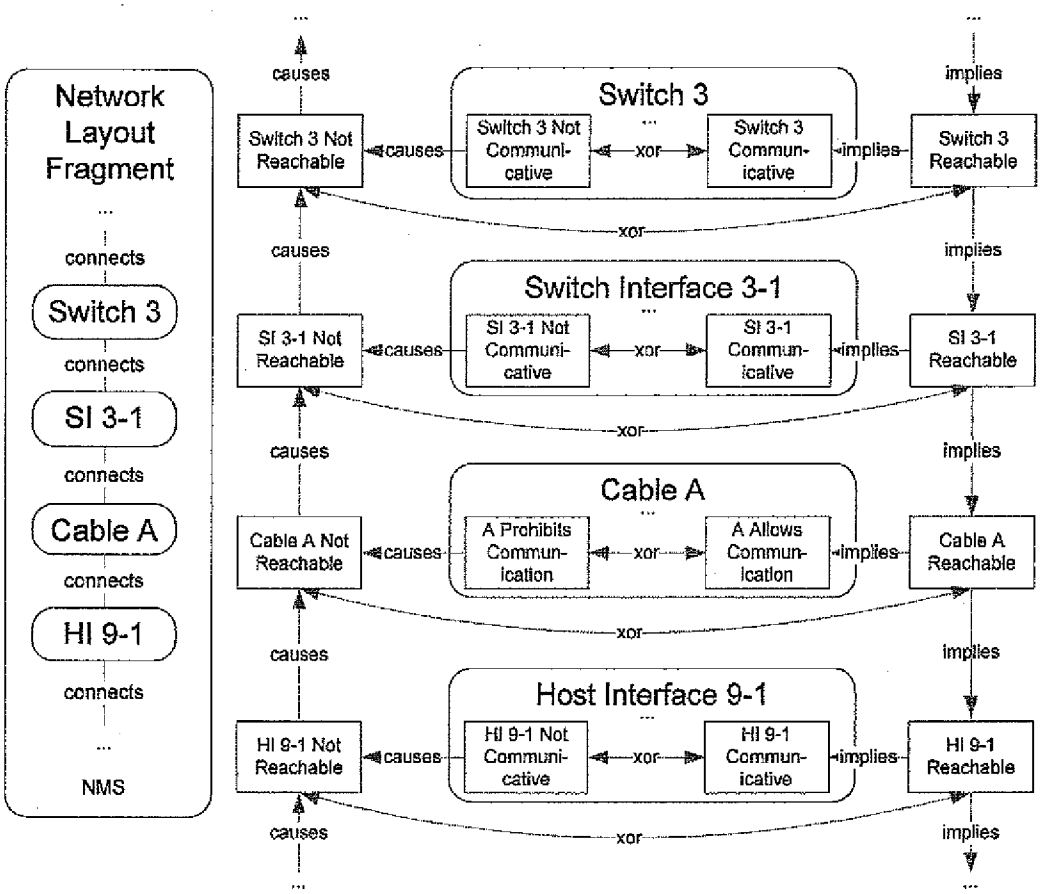
Network Layout Fragment
...
connects
Switch 3
connects
SI 3-1
connects
Cable A
connects
HI 9-1
connects
...
NMS
causes
Switch 3
Switch 3 Not Reachable
causes
Switch 3 Not Communi-cative
xor
Switch 3 Commun-icative
implies
Switch 3 Reachable
xor
causes
implies
Switch Interface 3-1
SI 3-1 Not Reachable
causes
SI 3-1 Not Communi-cative
xor
SI 3-1 Commun-icative
implies
SI 3-1 Reachable
xor
causes
implies
Cable A
Cable A Not Reachable
causes
A Prohibits Commun-ication
xor
A Allows Commun-ication
implies
Cable A Reachable
xor
causes
implies
Host Interface 9-1
HI 9-1 Not Reachable
causes
HI 9-1 Not Communi-cative
xor
HI 9-1 Commun-icative
implies
HI 9-1 Reachable
xor
causes
implies

FIGURE 4

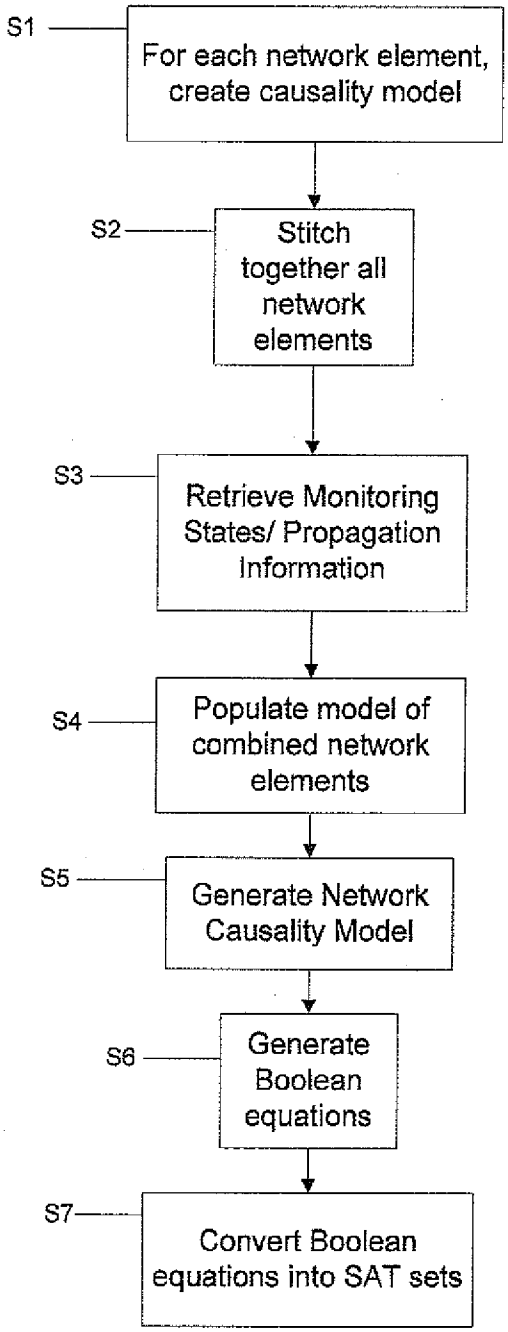
S1
For each network element, create causality model
S2
Stitch together all network elements
S3
Retrieve Monitoring States/ Propagation Information
S4
Populate model of combined network elements
S5
Generate Network Causality Model
S6
Generate Boolean equations
S7
Convert Boolean equations into SAT sets

FIGURE 5

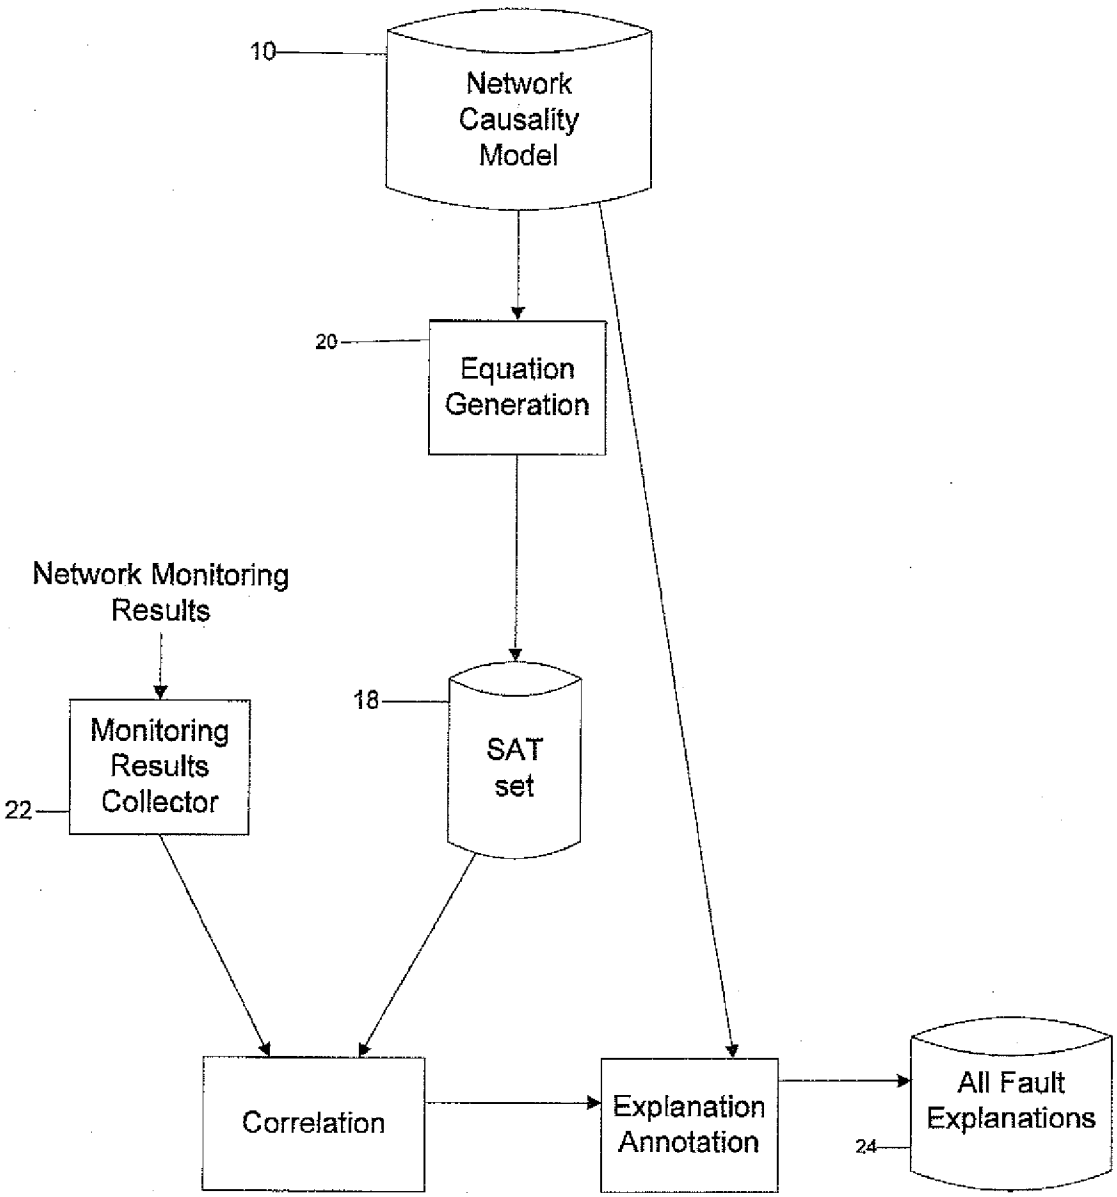
10
Network Causality Model
20
Equation Generation
Network Monitoring Results
18
SAT set
Monitoring Results Collector
22
Correlation
Explanation Annotation
All Fault Explanations
24

FIGURE 6

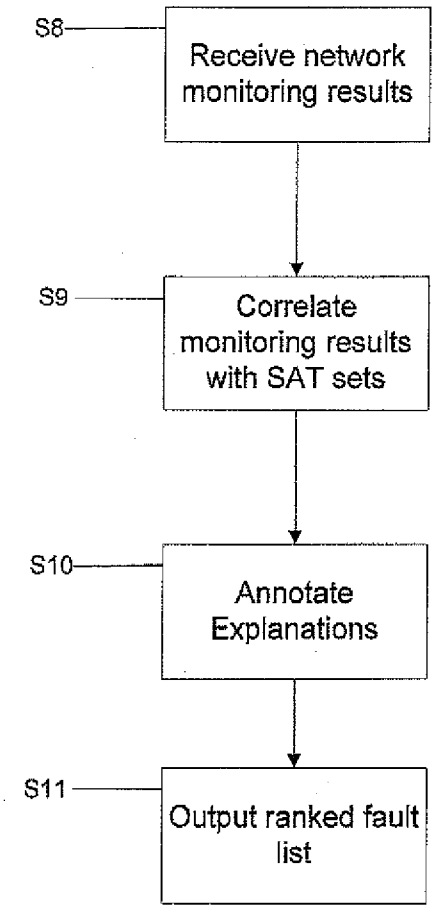
S8
Receive network monitoring results
S9
Correlate monitoring results with SAT sets
S10
Annotate Explanations
S11
Output ranked fault list

FIGURE 7

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 09/55958

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8) - G06F 11/00 (2010.01)
USPC - 714/26
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC(8): G06F 11/00 (2010.01)
USPC: 714/26

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
IPC(8): G06F 11/00 (2010.01) (keyword limited; terms below)
USPC: 714/1,4,25-26,43,48; 709/223,224,235; 706/45 (keyword limited; terms below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
PubWEST(PGPB, USPT, EPAB, JPAB); Google; Search Terms: fault diagnosis error failure defect wrong shortcoming examination interpretation determin network component member causal causality antecedent causation determinant system diagram Bolean expression formula equation monitor result outcome output correlat result outcome output display rank prio

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 2005/0096854 A1 (LARSSON et al.) 05 May 2005 (05.05.2005), entire document, especially abstract, para. [0002], [0004], [0008], [0012], [0014], [0017], [0020]-[0021], [0025], [0030], [0051]-[0052], [0055], [0062], [0068], [0077]-[0078], [0082]-[0083], [0086], [0102]-[0103], and [0115].	1-4, 6-10, 12-16, 18 ---------- 5, 11, 17
Y	US 2008/0039993 A1 (CLEARY et al.) 14 February 2008 (14.02.2008), entire document, especially para. [0006]-[0007].	5, 11, 17
A	US 5,661,668 A (YEMINI et al.) 26 August 1997 (26.08.1997), entire document	1-18
A	US 2008/0209269 A1 (BRODIE et al.) 28 August 2008 (28.08.2008), entire document	1-18

☐ Further documents are listed in the continuation of Box C. ☐

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
10 April 2010 (10.04.2010)	22 APR 2010

Name and mailing address of the ISA/US	Authorized officer:
Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774