



- (51) International Patent Classification:
H04L 12/00 (2006.01) H04L 1/24 (2006.01)
H04L 12/02 (2006.01)
- (21) International Application Number:
PCT/US2015/042966
- (22) International Filing Date:
30 July 2015 (30.07.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/031,009 30 July 2014 (30.07.2014) US
14/813,841 30 July 2015 (30.07.2015) US
- (71) Applicant: FORWARD NETWORKS, INC. [US/US];
5050 El Camino Real, Suite 226, Los Altos, California
94022 (US).
- (72) Inventors: ERICKSON, David; 1175 Glin Ter.,
Sunnyvale, California 94089 (US). HELLER, Brandon;
830 Fremont Street, Apt. 2, Menlo Park, California 94025
(US). HANDIGOL, Nikhil; 180 Elm Court, Apt. 508,
Sunnyvale, California 94086 (US). KAZEMIAN, Pey-
man; 745 Mediterranean Lane, Redwood City, California
94065 (US). RADHAKRISHNAN, Sivasankar; 718 Old
San Francisco Road, Apt. 345, Sunnyvale, California
94086 (US).

- (74) Agent: FELD, Nathan L.; Hogan Lovells US LLP, Three
Embarcadero Center, 15th Floor, San Francisco, California
94111 (US).
- (81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,
TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE,
DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LI, LU,
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to the identity of the inventor (Rule 4.17(i))

Published:

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

(54) Title: SYSTEMS AND METHODS FOR NETWORK MANAGEMENT

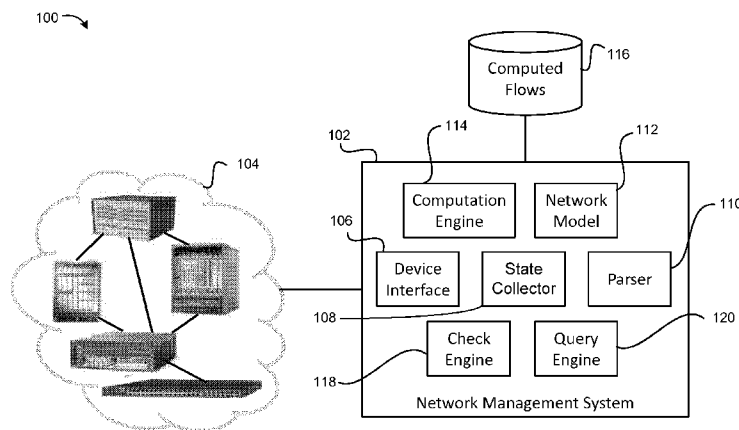


FIG. 1

(57) Abstract: Embodiments relate generally to network hardware, network software and methods for network management and testing. In some embodiments, state information (e.g., configuration data, forwarding states, IP tables, rules, network topology information, etc.) can be received from devices in a network. The state information can be parsed and used to generate a network model, which describes how data is processed by the network. Using the model, possible flow paths of data through the network can be identified and used to analyze the network and identify network behavior, such as types of traffic, frequency of rule matches, what kind of transformation occurs as traffic flows through the network, and where the traffic gets dropped, etc. Policies can be verified against the network model to ensure compliance, and in the event of non-compliance, a report or interface can indicate the cause and/or allow a user to explore specific details about the cause.

WO 2016/019172 A1

SYSTEMS AND METHODS FOR NETWORK MANAGEMENT

CROSS-REFERENCES TO RELATED APPLICATIONS

[0001] This application claims priority to U.S. Provisional Patent Application No. 5 62/031,009, filed July 30, 2014, titled “SYSTEM AND METHOD FOR NETWORK MANAGEMENT,” by David Erickson, et al., which is incorporated herein by reference in its entirety for all purposes.

BACKGROUND

[0002] Networks are growing more and more complex as the backbones of 10 modern information technology systems. A typical large company may employ hundreds or thousands of devices and software components from different vendors to form its network infrastructure. Growth in complexity and size also brings more points of failure, such as forwarding loops, configuration mistakes, reachability issues, or hardware failures.

15 [0003] Diagnosing network failures is difficult for several reasons. First, the forwarding states associated with each network device are distributed throughout the network and defined by their corresponding forwarding tables and other configuration parameters. Second, the distributed forwarding states are difficult to monitor – often requiring the network administrator to manually login to the device and conduct a low- 20 level test. Third, multiple administrators or users can edit the forwarding states at the same time, resulting in inconsistent forwarding states.

[0004] Conventional network diagnosis methods and tools are labor intensive, time consuming and often protocol dependent. For example, network administrators may use rudimentary tools (e.g., ping, traceroute, and Simple Network Management Protocol) 25 to track down network failures. Such methods only cover a tiny fraction of the network state space. The diagnosis process becomes even more difficult with the current trend of increasing network size and complexity.

[0005] Additionally, conventional network diagnosis methods and tools are *ad hoc* in nature as they only solve the manifested issues in the network. For example, simple questions such as “Can host A talk to host B?” or “Can customer X listen to the communication?” are difficult to answer. Thus, conventional tools cannot foresee or prevent problems before they arise.

[0006] On the other hand, large software companies can push out new software products quickly because they have a huge quantity of tests that comprehensively test the behavior of their software products prior to deployment. However, current network management and testing lacks sufficient testing capability to provide confidence that the network will work when deployed, especially when there are changes to the network several times a day.

[0007] Thus, there is a need to develop methods and tools to manage and verify networks in a fast, large-scale, automated, and systematic way.

15

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Various embodiments or examples of the invention are disclosed in the following detailed description and the accompanying drawings:

[0009] FIG. 1 illustrates a high level diagram of a network management system, in accordance with an embodiment of the present invention.

20 [0010] FIG. 2 illustrates a diagram of the network management system, in accordance with an embodiment of the present invention.

[0011] FIG. 3 is a diagram illustrating an example of identifying dead forwarding rules, in accordance with an embodiment of the present invention.

25 [0012] FIG. 4 is a diagram illustrating an example of investigating dead rules, in accordance with an embodiment of the present invention.

[0013] FIG. 5 is diagram illustrating an application of the network management system across multiple devices, in accordance with an embodiment of the present invention.

[0014] FIG. 6 is diagram illustrating a method of network management, in accordance with an embodiment of the present invention.

[0015] FIG. 7 is a high level block diagram of a computer system, in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0016] Various embodiments or examples may be implemented in numerous ways, including as a system, a process, an apparatus, a user interface, or a series of program instructions on a computer readable medium such as a computer readable storage medium or a computer network where the program instructions are sent over optical, electronic, or wireless communication links. In general, operations of disclosed processes may be performed in an arbitrary order, unless otherwise provided in the claims.

[0017] A detailed description of one or more examples is provided below along with accompanying figures. The detailed description is provided in connection with such examples, but is not limited to any particular example. The scope is limited only by the claims and numerous alternatives, modifications, and equivalents are encompassed. Numerous specific details are set forth in the following description in order to provide a thorough understanding. These details are provided for the purpose of example and the described techniques may be practiced according to the claims without some or all of these specific details. For clarity, technical material that is known in the technical fields related to the examples has not been described in detail to avoid unnecessarily obscuring the description.

[0018] In some embodiments, state information (e.g., configuration data, forwarding states, IP tables, rules, network topology information, etc.) can be received from devices in a network. The state information can be parsed and used to generate a network model, which describes how data is processed by the network. Using the model,
5 possible flow paths of data through the network can be identified and used to analyze the network and identify network behavior, such as types of traffic, frequency of rule matches, what kind of transformation occurs as traffic flows through the network, where the traffic gets dropped, etc.

[0019] Using the network model, a network manager or other such entity can
10 create various kinds of tests for the network. For example, a network manager can create tests to verify that “users X, Y, and Z can send information to each other”, “users can send information only across a certain range of IP addresses or transport ports”, or “user A should never be able to talk with user B”, or a wide range of areas.

[0020] There can be various applications that can take advantage of various
15 aspects of the present invention. For example, for a service builder during development, the present invention can be used to make sure the written code enables the network to function properly. Alternatively, in a production environment, the present invention can be used to validate the current behavior of the network or test the effect of proposed changes.

[0021] Furthermore, according to some embodiments, once the network behavior
20 is understood by a network manager, crafted physical packet(s) to match output from the model can be generated and sent into the physical network to make sure that the hardware behaves the way that the model predicts, to detect hardware malfunctions.

[0022] FIG. 1 illustrates a high level diagram 100 of a network management
25 system, in accordance with an embodiment of the present invention. As shown in FIG. 1, network management system 102 can connect to a network 104 that includes a plurality of network devices (e.g., switches, routers, firewalls, load balancers, gateways, bridges, and other network devices). Each network device may be produced by a different vendor

and may implement different operating systems, control interfaces, protocols, etc. Accordingly, network management system 102 can include a plurality of device interfaces 106, which may include vendor-specific, device-specific, and/or protocol-specific interfaces, through which network management system 102 can communicate with each network device.

[0023] In some embodiments, and as described further below with respect to FIG. 2, network management system 102 can include one or more state collectors 108 and one or more parsers 110. The state collectors 108 and parsers 110 can communicate with the network devices using device interfaces 106 to obtain state information for each network device. State information may vary from network device to network device and can include one or more of forwarding states, configuration files, internet protocol (IP) tables, topology information, and rules. State information can be parsed by the parsers to generate a network model 112, which describes how data is processed in the modeled network. A computation engine 114 can use the network model 112 to identify possible path flows that data may travel over in the network model 112. In some embodiments, the path flows may be stored in a data store, such as computed flows data store 116, which may be stored locally with network management system 102 or remotely.

[0024] In some embodiments, network management system 102 can further include a check engine 118 and a query engine 120. As described further below, check engine 118 can analyze flow paths stored in data store 116 to identify various properties, such as path, header, hop counts, quality of service (QoS), queues, ports, physical devices the flow traverses, tables within the device, forwarding type of the packet, packet header modifications, encapsulated packet header properties, or allocated bandwidth, or other properties. This enables the modeled network to be analyzed to verify that the network is functioning according to prescribed standards (e.g., set by a network administrator, developer, or other entity). Query engine 120 can be used to identify particular flow paths that meet queried criteria (e.g., failover flow paths between specified points, flow paths that traverse particular network devices in a particular order, etc.).

[0025] FIG. 2 illustrates an example 200 of a network management system that can be utilized to provide for automated and systematic management of the network via generating and utilizing a virtualized network model of the physical network. According to some embodiments, the present invention includes a system that collects network state information and enables searches for different types of traffic that can exist in the network. As seen in FIG. 2, in some embodiments, a network 104 can comprise numerous network devices including switches and routers from different vendors, such as Cisco and Juniper. Any or all of these network devices can have associated state information, including configuration data and/or forwarding states which may be written by a control plane, and which may be defined by one or more IP tables and other forwarding rules. The forwarding states can be used to implement the network's forwarding policy, which specifies how different classes of traffic are routed in the network as well as how much bandwidth is allocated to those classes of traffic.

[0026] Still referring to FIG. 2, the network management system can use one or more state collectors 108 to collect the state information, such as "snapshots" of the forwarding states, via connecting to the devices through SSH, telnet, SNMP, NETCONF, OpenFlow, or via any other device interface that enables the collectors to collect information from the device. The collected state information can include forwarding states, configuration files, IP tables, OpenFlow rules, topology information, and any other configuration or state made available by the device. Each snapshot may be augmented with information provided by a user, such as topology data. In addition, state collectors can periodically generate the collected state information. In some embodiments, the state collectors 108 can collect the state information only when a change in at least one forwarding state is detected.

[0027] In some embodiments, the network management system can use one or more parsers 110 to translate the collected state information 202 to construct a network model 112. In some embodiments, parsers 110 can be vendor-specific or protocol-specific. For example, the system can use a Cisco parser for parsing Cisco IOS, including reading IOS command output, compressing IP forwarding tables, and generating vendor-

independent forwarding descriptions that model the function of the device. In another example, the system can use a Juniper Junos parser for forwarding states associated with Juniper devices, or an OpenFlow parser for forwarding states associated with OpenFlow devices. In some embodiments, an incremental parser can be used to track the changes in the collected forwarding state and quickly update the network model 112.

[0028] The network model 112 can be a vendor/protocol independent model for data processing (e.g., packet processing) in the network 104. The network model can include forwarding rules 204 that are used to determine how packets are processed by devices in the network. In some embodiments, the network model 112 can be periodically updated. In some embodiments, instead of recomputing the entire network model each time the actual network's state changes, the network model 112 can be incrementally updated by only updating the associated changed forwarding states, which can be realized through rule dependency tracking. In rule dependency tracking, an individual update to a state table, such as a VLAN change at an input port, automatically causes rules that match that input port to update, without requiring the re-generation of other unrelated rules. The list of updated rules may then be used to update the modeling of affected traffic classes and paths in the network, leaving the rest of the model unaffected.

[0029] In some embodiments, the network management system can include Versioned State Dump Files 206 that are associated with the network model 112. The Versioned State Dump Files 206 can support back-tracking to connect specific traffic flows to the configuration and state files and lines of configuration and state files that caused them. In this way, observed connectivity and policy errors can be automatically linked to their root causes.

[0030] As shown in FIG. 2, a computation engine 114 can identify possible flow paths in the network in a port-to-port manner through generating symbolic packets and pushing the symbolic packets through the network model. In addition to the traditional bits (0 or 1), a symbolic packet may have one or more wild cards (*) present. The symbolic packet can replace wild cards with 0 or 1 upon encountering a forwarding rule.

Thus, the symbolic packet can demonstrate how a packet or a class of packets is modified as it travels by composing the transfer functions along the path.

[0031] In some embodiments, the computation engine can generate a minimal set of physical packets so that every forwarding rule is exercised and covered by at least one physical packet. In another embodiment, the physical packets can be created for every possible external outlet, so that the “crawling” of the network is systematic, instead of fractional.

[0032] In some embodiments, the computation engine can save the identified possible flow paths and forwarding rules 206 in a flow path database 116, including the possible flow paths and hops, along with flow path information associated with a given packet traveling on the flow path, such as header value changes, rules that the packet matches, the transformations at each hop, etc.

[0033] In some embodiments, the computation engine 114 can periodically update the flow path database. In another embodiment, the computation engine can incrementally update the flow path database only to reflect the changed forwarding rules and the updated flows.

[0034] A check engine 118 can use the flow path database 116 to check properties of possible flows in the network such as path, header, hop counts (the number of physical or virtual network elements a flow traverses, or the number of tables within those devices it traverses), quality of service properties such as priority, queues, ports, physical devices the flow traverses, tables within the device, physical characteristics such as location(s) or distance, forwarding type of the packet (L2, L3, MPLS, etc.), packet header modifications, encapsulated packet header properties, or allocated bandwidth. In one example, a check can be issued to determine whether a certain flow exists. In another example, a check can be issued to verify that “from Port 1, all the traffic that has IP address IP1 should reach Port 2”. In another example a check could be issued to validate that sufficient failover paths exist between specified locations to ensure traffic and flows will still be able to pass between the locations in the event of a failure on a path. Check

engine 118 can also check for the absence of flows matching specified properties. Meta-checks can also be created that check for higher-level policy such as “user1 endpoints are fully reachable”, or “user1 endpoints are fully isolated”, or “user1 is PCI compliant”, which within the system would be decomposed to one or more basic checks based on
5 properties specified previously. In some embodiments, the system can issue a “can-reach-all” check or a “cannot-reach-any” check. For example, given two subnets, can each IP in the source subnet reach every IP in the destination subnet? This example can be generalized to “header objects”, which are descriptions of packet header values that may contain one or more wildcard values. The corresponding question would be, “can
10 every potential packet value in the source header object reach every potential packet value in the destination header object?”.

[0035] According to some embodiments, certain checks can be defined to run through the network regularly. In some embodiments, the checks can be defined to execute at particular times or particular time intervals. In some embodiments, the checks
15 can be defined to execute upon the detection of particular network events. For example, given it is expected that there is connectivity between two ports using a specific packet, this specific packet can be used to check the connectivity between the two ports whenever there is a change in the network.

[0036] In some embodiments, a query engine 120 can use the flow path database
20 116 to find flows with a certain path, header value, or allocated bandwidth. In some embodiments, a query language can be used to query the flow path database. For example, a query 208 can search for possible flow paths from Port 1 with destination IP subnet 10.1.0.0/16. In another example, a query can search for a range of destination IP addresses that a specific packet can reach. In yet another example, a query can search for
25 all possible flow paths that traverse routers R1, R2, and R3 in that order or go through no more than 3 hops from source to destination. In another example, a query can search for all flows that follow a certain path through the network.

[0037] The present invention can include a query interface to provide query function to the network. For example, a network manager 210 can issue a query 212 for “flows from a particular device”, and the query engine can provide all possible flows from the particular device. A query can also be issued to show what happens to packets going from one device to another device. A query may also find all packets that reach its destination with a particular packet header.

[0038] In addition, embodiments of the present invention can include a check interface to verify whether certain flow(s) should be present or not. For example, the network manager 210 can issue a check for “expect flows from this device to another device”. Embodiments of the present invention can further reveal information as to why a particular check fails, such as showing why no flows exist between two devices, or finding flows when there should be none. For example, when a check “expect flows from device1 to device2” fails, the invention can present the user with information about where the flows from device1 to device2 are being dropped, for instance at a misconfigured firewall device. In another example, when a check “no flows are expected from port1 to port2” fails, the invention can provide details on the one or more flows that go from port1 to port2 thereby causing the check to fail. In another example of a failed can-reach-all check from IP subnetA to IP subnetB, the invention can indicate which IP addresses in subnetA cannot reach IP addresses in subnetB and further provide details on where these respective flows get dropped in the network.

[0039] According to some embodiments, the present invention can generate an analytic report on the network. For example, the analytic report can give suggestions to the network administrators on how to improve the performance of the network, or show whether the network has some security weak points. Furthermore, the present invention can reduce the knowledge burden of the network administrator, as it does not require the administrator to log in to each device, understand device configurations, understand how to retrieve them, or reason about interactions between forwarding tables. Analytic reports may contain suggestions for cross-network decisions such as determining the optimal routing path across networks, or optimal peering or caching locations. Analytic reports

could also contain cross-customer statistics such as most used device types and how the network devices are most frequently used.

[0040] In another example, the present invention can provide a differentiation chart on the behavior history of the network. For example, the network management system can keep taking “snapshots” of the network, and it can reveal why a failure occurs through comparing and differentiating between different snapshots’ network data as well as their behavior.

[0041] According to some embodiments, the present invention can provide a method to compute and query for the differences between the configuration and state of the network at two different points of time or “snapshots”. In some embodiments, this difference can indicate only those changes that actually affected the behavior of the network. For example, suppose there were ten different changes to the border gateway protocol (BGP) configuration on router1, eight of which did not affect the forwarding behavior of the network, whereas two changes affected forwarding behavior through the network. A network manager or user may, in some examples, use this information to determine quickly which changes to a network configuration may be associated with a network failure occurred.

[0042] In some embodiments, the system can identify the behavioral differences of two network snapshots. The behavior differences can be searched such that the returned results only show the flow paths that are added or deleted and match the characteristics specified in the search. For example, a search for flows “from port 1 destined to IP subnet 10.1.0.0/16” may show two added flow paths that can enter the network with VLAN 10 and 11 and one removed flow path which could previously enter the network with VLAN 9.

[0043] The present invention also provides a system to compute and query for any differences in topology as the network evolves. In some examples, this can include the addition or removal of network devices, addition or removal of links between the devices, etc.

[0044] In some embodiments, the present invention can provide an interface to automatically use the information about the history of changes in the network to determine the possible causes of check failures in the network and suggest possible fixes to a network manager or user. In some examples, the user may further add “tags” or notes
5 about past network events that can indicate important events in the evolution of the network, including events that might have caused network behavior to change.

[0045] Embodiments of the present invention can identify when a set of changes to network configuration or to network devices does not cause any behavioral differences to the network. In some embodiments, network checks are created to indicate the
10 intended network behavior or invariants. In other embodiments, the expected network behavior may be inferred automatically via a combination of paths, traffic types and their transformations through the network or the trace of symbolic packets through the network. Further, symbolic packets may also be matched across the different points of time, being compared based on different criteria to not only identify additions and
15 removals from the set of all possible ways the network treats traffic, but also to identify those types of traffic or paths or symbolic packets whose behavior has been modified in some way. For example, traffic from IP address 10.0.0.1 to IP address 10.1.0.1 may have traversed over a path router1 to router2 to router3 in the first network “snapshot”, but traverses over a path through router1 to router4 to router3 in the second network
20 “snapshot”.

[0046] In some embodiments, the system can be used to automatically extract the current forwarding policies of the network and convert them into checks. These checks can be applied to future snapshots to detect any violation of these policies. For example, the current subnet and port connectivity can be extracted and converted to a set of checks.

[0047] According to some embodiments, the network management system can detect and investigate dead forwarding rules. Dead rules are defined as rules that cannot be exercised and thus can be safely removed from the network. FIG. 3 is a diagram
25 illustrating an example of identifying dead forwarding rules, in accordance with an

embodiment of the present invention. As shown in FIG. 3, a forwarding device 302 can include, for example, forwarding table 1 304, forwarding table 2 306, and forwarding table 3 308. Each of forwarding tables 1, 2 and 3 (304, 306, and 308) can include one or more forwarding rules (indicated as hashed boxes), among which some of the forwarding rules are dead rules 312 that cannot be exercised or hit at all. To find these dead rules, the computation engine can generate and push a set of symbolic packets 310 with a wildcard (*) on every header bit to find the dead rules that cannot match any packet. When each possible forwarding path has been identified, those rules that have not been used to forward any packet can be identified as dead rules, such as rule 312.

10 **[0048]** Embodiments of the present invention can determine network test coverage percentage. In one example, given a set of test packets, the system can compute the fraction of network ports, links, rules, queues, etc., covered. In another example, the system can compute the test packets and/or probe locations needed to achieve the desired level of network test coverage.

15 **[0049]** FIG. 4 is a diagram 400 illustrating an example of investigating dead rules, in accordance with an embodiment of the present invention. As shown in FIG. 4, embodiments of the present invention provide a system and method to investigate dead rules. The system and method can investigate why certain rules are not hit; for example, the cause might be a related forwarding table that does not send traffic to the dead rule. For example, as shown in FIG. 4, flow type A 402 directs traffic to dead rule 404 through rule 406. Similarly, flow type B 408 directs traffic to dead rule 404 through rule 410. However, as shown in FIG. 3, rules 406 and 410 are also dead rules, as such there is no path to rule 404.

25 **[0050]** Embodiments of the present invention provide a system and method to detect rules within the same or different tables that have overlapping match fields, and present this to the user.

[0051] Embodiments of the present invention provide a system and method to detect rules within the same or different tables related to a given input rule such that packets matching them will also be matched by the given rule.

5 [0052] Embodiments of the present invention provide a system and method to detect rules within the same or different tables that are ambiguous due to having the same matching priority, and some or full overlap in match fields.

[0053] Embodiments of the present invention provide a system and method to analyze rules and the network more generally, and provide suggestions for optimization, such as the elimination of rules, modification to rules, or addition of other rules.

10 [0054] Embodiments of the present invention provide a system and method to analyze and detect problems within the network, and provide suggestions to the user on how to fix them. These problems include those detected by network checks, queries, or other performance, optimization, or correctness related problems. The invention may also automatically apply changes to fix such problems.

15 [0055] Embodiments of the present invention provide a system and method to modify a snapshot of a network to model and test potential changes to the network before making changes to the actual physical/virtual network. For example, if one or more changes are being made to the network (e.g., configuration information is being updated on one or more network devices in the network), those changes can be tested using the
20 network model prior to pushing those changes to the actual network devices, or prior to adding or removing network devices to or from the network. In this example, changes can include adding or removing network devices to or from the network, updating forwarding tables, or any other configuration changes that may affect forwarding behavior through the network. Checks, as described above, may then be run on the
25 updated network model to identify potential failures associated with the changes made to the network. Once the updated network model has passed the checks, the changes can be pushed to the corresponding network devices in the network. In some embodiments, the network management system can automatically send updated configuration information

to each network device in the network (e.g., using the one or more device interfaces). In some embodiments, the network management system can send updated configuration information only to those devices that are being updated, rather than every device in the network. In some embodiments, the network management system may cause an update interface to be displayed to a user (such as a network administrator or network developer), through which the network management system may receive manual instructions to update the configuration information at one or more devices in the network.

[0056] Embodiments of the present invention can model the effect(s) of configuration change(s) to one or more devices in a network, and how such configuration changes may affect configuration or forwarding state contained in other devices. One example would be modifying the configuration of a distributed protocol such as BGP or OSPF in one device, and modeling its effect on other devices and forwarding state in the network.

[0057] Embodiments of the present invention provide a system and method to take the configuration from a network snapshot and apply it to all devices in the network.

[0058] FIG. 5 is diagram 500 illustrating an application of the network management system across multiple devices, in accordance with an embodiment of the present invention. As shown in FIG. 5, the system and method provided herein can apply to a network including multiple network devices (502, 504, 506) instead of within a single network device. Similar to FIG. 3, symbolic packets 508 with wild card header bits can be sent across the multiple network devices (502, 504, 506) to identify the possible paths that traverse these devices. Although three devices are shown in the example of FIG. 5, embodiments of the present invention can be used with any number of network devices.

[0059] In some embodiments, a request to perform a particular check can be received through a user interface, such as a graphical user interface (GUI) or command line interface (CLI). In some embodiments, check engine 118 can perform a

CanReachAll check of the network. The check can be initiated by receiving an input string, such as: <HO1> CanReachAll <HO2>, where HO1 and HO2 are packet header objects (specification of values for packet header fields). In some embodiments, HO1 can specify values HO1.hValues for a list of header fields HO1.hFields, and HO2 can specify values HO2.hValues for a list of header fields HO2.hFields. In some 5 embodiments, any header fields that are not specified by HO1.hFields can be set as wildcards at network ingress. This check can be used to verify that for all possible values of the header fields HO1.hFields and HO2.hFields as allowed by HO1.hValues and HO2.hValues, respectively, there is complete reachability from HO1 to HO2, i.e., for 10 each possible set of values HO1.hFields can take as allowed by HO1.hValues, there exist packets entering the network with that set of values for HO1.hFields, that will exit the network with all possible set of values HO2.hFields can take as allowed by HO2.hValues. Note that any other header fields unspecified by HO1.hFields and HO2.hValues may not have complete reachability.

15 **[0060]** In some embodiments, a check, such as CanReachAll described above, can further include optional fields. For example, “from” and “to” packet header object (HO) filter specifications can optionally be augmented with a location specification (port/table/device) to narrow the flow path search. Once the flows are retrieved from the database, these location filters are not used for validation of HO reachability.

20 **[0061]** In one example of a CanReachAll check, the string “ipv4_src addr 192.168.2.0/24 canReachAll ipv4_dst addr 192.168.3.0/24” may be received. This string indicates a request to determine whether each source IP address in subnet 192.168.2.0/24 can reach all destination IP addresses in subnet 192.168.3.0/24. The reachability may potentially be valid only for a subset of protocols, port numbers etc. The check condition 25 here does not look for reachability across all protocols. Unless there is source HO based filtering/forwarding in the network, this kind of reachability is likely to always be present as long as the network contains the specified destination IP subnet.

[0062] In another example of a CanReachAll check, the string “router1 port1 ipv4_src addr 192.168.2.0/24 canReachAll ipv4_dst addr 192.168.3.0/24” can be received, which indicates a request to check whether each source IP address in subnet 192.168.2.0/24 at router1 and port1 can reach all destination IP addresses in subnet 192.168.3.0/24. By specifying a port number or location, the search is narrowed to ensure that the source HS at that particular source port number (location) can reach the specified destination IP subnet.

[0063] In another example of a CanReachAll check, the string “ip_proto tcp ipv4_src addr 192.168.2.0/24 canReachAll ipv4_dst addr 192.168.3.0/24” can be received, which indicates a request to check whether each source IP address in subnet 192.168.2.0/24 can reach all destination IP addresses in subnet 192.168.3.0/24 using TCP packets. This check explicitly verifies a specific protocol to address.

[0064] In one example of a CanReachAll check, the string “ip_proto ALL ipv4_src addr 192.168.2.0/24 canReachAll ipv4_dst addr 192.168.3.0/24” can be received which indicates a request to check whether each source IP address in subnet 192.168.2.0/24 can reach all destination IP addresses in subnet 192.168.3.0/24 using all IP protocol types.

[0065] In some embodiments, some assumptions may be made on check string input. For example, in some embodiments it may be assumed that the set of header fields in HO1 and HO2 must be orthogonal. For example, if the same header field (e.g. ipv4_dst) is specified in both HO1 and HO2, then the value for the field in at least one of HO1 and HO2 must contain no wildcard bits. Additionally, for the set of matching flow paths, if any rule in the network can look at the value of a header field hField1 specified in HO1 and dynamically rewrite the value of a header field hField2 specified in HO2 based on the value of hField1 in HO1, at least one of the following conditions must hold: 1) the value of hField1 in HO1 has no wildcard bits and/or 2) the value of hField2 in HO2 has no wildcard bits.

[0066] In some embodiments, a given check can include multiple extensions. For example, the CanReachAll check can include the following extensions.

[0067] 1. <HO1> canReachAll <list of ports P> -- For each possible set of values HO1.hFields can take as allowed by HO1.hValues, there exist packets entering the network that can exit the network at each of the ports in the list P.

[0068] 2. <list of ports P> canReachAll <HO2> -- There exist packets entering the network at each ingress port in the list P that can exit the network with each possible set of values HO2.hFields can take as allowed by HO2.hValues.

[0069] 3. <HO1> canOnlyReach <list of ports P> -- For each possible set of values HO1.hFields can take as allowed by HO1.hValues, there exists no packet that can exit the network at any egress port other than those in the list P.

[0070] 4. <HO1> canReachAll <list of devices D> and <list of devices> canReachAll <HO2> -- Similar to 1 and 2, but with devices instead of ports.

[0071] 5. <list of ports P1> canReachAll <list of ports P2> -- For each source port in the list P1, there exist packets that can reach each of the ports in the list P2.

[0072] 6. <HO1> atEachOf <list of ports P1> canReachAll <HO2> atEachOf <list of ports P2> -- For each source port srcP in the list P1, and each destination port dstP in the list P2, srcP <HO1> canReachAll dstP <HO2> holds. For example, for each possible pair of ports from P1 and P2, there is complete reachability from the source HS HO1 to the destination HS HO2.

[0073] 7. <HO1> onlyAt <list of ports P1> canReachAll <HO2> onlyAt <list of ports P2> -- <HO1> canReachAll <HO2> holds, but we only query for those flow paths that begin at a source port from the list P1 and end at a destination port from the list P2.

[0074] Although the above example CanReachAll check is described with respect to a text-based interface, similar functionality may also be abstracted behind a GUI or other user interface that exposes similar behavior.

[0075] Embodiments of the present invention can provide a multi-domain or “truly distributed” collection and testing system. In this mode, the verification process runs on a mix of local agents (each responsible for a subset of the network) and global agents (responsible for the entire network). Rule changes that can be validated locally are processed immediately by the relevant local agent, and are then synchronized with the global agents. Rule changes that affect multiple local agents must be implemented by the global agents, which may enforce a single global ordering on all rule events. Optionally, the topology and rule knowledge may be presented to each local agent in a way that aggregates or otherwise filters out some of the global knowledge, to enhance scalability.

[0076] Embodiments of the present invention can provide a high availability mode, where multiple collection agents may be run, and they may coordinate together to collect network configuration and state. In the event that one or more fail, the other agents may coordinate to collect the state from which the devices the failed agents were collecting.

[0077] Embodiments of the present invention can identify where different network elements such as IP addresses, IP subnets or VLANs are present in the network. For example, the network model can be used to locate the IP address 192.168.10.2 in the network as being located on router1 port1 and accessible over VLAN 10. In some embodiments, this may be done without sending any real packets through the network. The invention can also be used to locate IP subnets. For example, the network model can determine that the IP subnet 10.1.0.0/24 is spread across 3 different router ports in the network along with specific VLANs on which that the IPs are reachable.

[0078] In some embodiments, a query interface may be provided to locate IP addresses, IP subnets, VLANs, devices, network interfaces or other network elements. The interface may further be augmented with visual diagrams that locate the elements in the network.

[0079] Embodiments of the present invention can provide an interface to trace classes or types of traffic through the network. For example, searching for the traffic

traces from a source IP address 10.0.0.10 to an IP subnet 10.1.2.0/24 can illustrate paths that the queried traffic can take through the network, and provide a manner of refining the search by visually presenting additional filters relevant to the search results.

[0080] In some embodiments, the trace results are classified along various dimensions such as paths, values of header fields at different devices in the network, path length etc., to generate some of the filters that can help refine the trace search results. A network manager or a user of the system can use the suggested filters to refine the trace and rapidly converge to the relevant traffic of interest. For example, the network manager may search for traffic traces from source IP address 10.0.0.10 to IP subnet 10.1.2.0/24 and the system presents the network manager with the list of devices [router1, router2, router3, router4] traversed by the trace results and that some traces relevant to the search are dropped at router2. The network manager then selects router2 and chooses to only view the dropped traffic. The system then presents the network manager with only those trace results that match these additional filters as well.

[0081] Embodiments of the present invention can provide an interface where the user can view virtual packet traces. These are computed by the system without sending actual packets into the network. The trace itself may comprise traffic that traverses different VLANs, or several different IP addresses, rather than describe only a single packet. In some embodiments, the system can also present the user with traces that correspond to individual actual packets that can be seen in the network.

[0082] FIG. 6 is diagram illustrating a method 600 of network management, in accordance with an embodiment of the present invention. At 602, state information can be received from a plurality of network devices in a network. As described above, state information can include information received from a device, or devices, in the network (e.g., forwarding states, configuration files, internet protocol (IP) tables, and rules) and/or information received from a user, such as a developer or administrator, (e.g., network topology information). The state information can be received through one or more device, network, and/or protocol specific interfaces (e.g., SSH, telnet, SNMP,

NETCONF, OpenFlow, or via any other mechanism that enables the collectors to collect information from the device or network).

[0083] At 604, a plurality of network device specific parsers can parse the network information to create parsed network information. In some embodiments, the network device specific parsers can include vendor-specific and/or protocol-specific
5 parsers. For example, a Cisco parser may be used to parse devices using Cisco IOS (e.g., reading IOS command output, compressing IP forwarding tables, and generating forwarding descriptions), a Juniper Junos parser may be used to parse Juniper devices, an OpenFlow parser may be used to parse state information received from OpenFlow
10 devices, etc. In some embodiments, the plurality of network device specific parsers can include one or more incremental parsers used to track changes made to state information received from the plurality of network devices.

[0084] At 606, a network model can be generated based on the parsed network information. The network model can describe how data is processed by the network. As
15 described above, in some embodiments, the network model can include forwarding rules that can be used to determine how data is routed through the network. In some embodiments, the network model can represent a snapshot of the network at a given time and may be updated to reflect changes to the state information received from the plurality of network devices. In some embodiments, the state information used to generate the
20 network model can be backed up.

[0085] At 608, one or more flow paths can be computed using the network model. In some embodiments, the one or more flow paths can include all possible port-to-port paths in the network. Symbolic packets can be pushed through the network model, which include one or more wild card bits. The wild card bits can be updated (e.g., replaced with
25 a '1' or '0') as each forwarding rule is encountered. The changes made to the wild card bits can be used to identify a transfer function associated with a particular flow path. In some embodiments, the one or more flow paths can be stored in a data store. In some embodiments, as changes are detected to the state information (e.g., using the incremental

parser), the one or more flow paths can be updated based on the change, and the updated one or more flow paths can be stored in the data store.

[0086] At 610, analyzing the one or more flow paths to identify network properties. As described above, properties of possible flows in the network can include path, header, hop counts (the number of physical or virtual network elements a flow traverses, or the number of tables within those devices it traverses), quality of service properties such as priority, queues, ports, physical devices the flow traverses, tables within the device, physical characteristics such as location(s) or distance, forwarding type of the packet (L2, L3, MPLS, etc.), packet header modifications, encapsulated packet header properties, or allocated bandwidth.

[0087] In some embodiments, the flow paths can be analyzed by querying the data store based on a particular network property, receiving flow paths that match that query, and then comparing the received flow paths to a rule associated with the particular network property. For example, a query can be submitted to identify a number of failover paths between particular locations (i.e., a particular network property) and flow paths matching the failover paths may be returned. The number of flow paths returned may then be compared to a predetermined number (i.e., a rule associated with the particular property). Similar checks may also be performed to validate other network properties, as described above. In some embodiments, a report may be generated based on the comparison of the subset of the one or more flow paths to the at least one rule, the report including one or more suggested network configuration changes. For example, if the number of failover paths is less than the predetermined number, the report may include suggested changes to increase the number of available paths.

[0088] In some embodiments, the one or more flow paths can be analyzed to identify one or more violations of at least one network property. For example, as described above, a network property may be associated with one or more rules. When the network state information does not meet the conditions defined in the one or more rules, a violation of the network property may be indicated. Network properties and

associated rules may be received from a user, such as a network administrator, developer, or other user, inferred by the network management system based on current network state information, or extracted from other third party sources. In some embodiments, when a violation of a network property is identified, a report identifying a configuration of a network device associated with the one or more identified violations can be generated. The generated report can be stored and/or a user can be sent an alert including the generated report.

[0089] FIG. 7 is a high-level block diagram 700 of a computer system, in accordance with an embodiment of the present invention. As shown in FIG. 7, a computer system can include hardware elements connected via a bus, including a network interface 702, that enables the computer system to connect to other computer systems over a local area network (LAN), wide area network (WAN), mobile network (e.g., EDGE, 3G, 4G, or other mobile network), or other network. The computer system can further include one or more processors 704, such as a central processing unit (CPU), field programmable gate array (FPGA), application-specific integrated circuit (ASIC), network processor, or other processor. Processers may include single or multi-core processors.

[0090] In some embodiments, the computer system can include a graphical user interface (GUI) 706. GUI 706 can connect to a display (LED, LCD, tablet, touch screen, or other display) to output user viewable data. In some embodiments, GUI 706 can be configured to receive instructions (e.g., through a touch screen or other interactive interface). In some embodiments, I/O interface 708 can be used to connect to one or more input and/or output devices such as mice, keyboards, touch-sensitive input devices, and other input or output devices. I/O interface 708 can include a wired or wireless interface for connecting to infrared, Bluetooth, or other wireless devices.

[0091] In some embodiments, the computer system may include local or remote data stores 710. Data stores 710 can include various computer readable storage media, storage systems, and storage services, as are known in the art (e.g., disk drives, CD-

ROM, digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, relational databases, object storage systems, local or cloud-based storage services, or any other storage medium, system, or service). Data stores 710 can include data generated, stored, or otherwise
5 utilized as described herein. For example, data stores 710 can include computed flows 712 and network models 714, generated and stored as described above. Memory 716 can include various memory technologies, including RAM, ROM, EEPROM, flash memory or other memory technology. Memory 716 can include executable code to implement methods as described herein. For example, memory 716 can include a network analyzer
10 module 718 and report generator module 720 that each implement methods described herein.

[0092] Although the foregoing examples have been described in some detail for purposes of clarity of understanding, the above-described inventive techniques are not limited to the details provided. There are many alternative ways of implementing the
15 above-described invention techniques. The disclosed examples are illustrative and not restrictive.

CLAIMS

What is claimed is:

1. A computer-implemented method, comprising:
receiving state information from a plurality of network devices in a network;
parsing, by a plurality of network device-specific parsers, the network information to create parsed network information;
generating a network model based on the parsed network information, wherein the network model describes how data is processed by the network;
computing one or more flow paths using the network model; and
analyzing the one or more flow paths to identify network properties.
2. The computer-implemented method of claim 1, wherein the state information includes one or more of forwarding states, configuration files, internet protocol (IP) tables, and rules received from the plurality of network devices.
3. The computer-implemented method of claim 1, wherein the state information further includes network topology data received from a user.
4. The computer-implemented method of claim 1, wherein analyzing the one or more flow paths to identify network properties further comprises:
verifying network correctness invariants.
5. The computer-implemented method of claim 1, further comprising:
storing the one or more flow paths in a data store.
detecting, using at least one of the plurality of parsers, a change to the state information associated with at least one network device in the network;
updating the one or more flow paths based on the change; and
storing the updated one or more flow paths in the data store.
6. The computer-implemented method of claim 1, further comprising:
storing the one or more flow paths in a data store; and

wherein analyzing the one or more flow paths to identify network properties further comprises:

querying the data store based on at least one network property;
receiving a subset of the one or more flow paths from the data store in response to the query; and
comparing the subset of the one or more flow paths to at least one rule associated with the at least one network property.

7. The computer-implemented method of claim 6, further comprising:
generating a report based on the comparison of the subset of the one or more flow paths to the at least one rule, the report including one or more suggested network configuration changes.

8. The computer-implemented method of claim 1, further comprising:
analyzing the one or more flow paths to identify one or more violations of at least one network property.

9. The computer-implemented method of claim 8, further comprising:
generating a report identifying a configuration of a network device associated with the one or more identified violations.

10. A system, comprising:
one or more processors; and
one or more memory devices including instructions that, when executed by the one or more processors, cause the system to:
receive state information from a plurality of network devices in a network;
parse, by a plurality of network device-specific parsers, the network information to create parsed network information;
generate a network model based on the parsed network information, wherein the network model describes how data is processed by the network;
compute one or more flow paths using the network model; and

analyze the one or more flow paths to identify network characteristics.

11. The system of claim 10, wherein the state information includes one or more of forwarding states, configuration files, internet protocol (IP) tables, and rules received from the plurality of network devices.

12. The system of claim 10, wherein the state information further includes network topology data received from a user.

13. The system of claim 10, further comprising:
storing the one or more flow paths in a data store;
detecting, using at least one of the plurality of parsers, a change to the state information associated with at least one network device in the network;
updating the one or more flow paths based on the change; and
storing the updated one or more flow paths in the data store.

14. The system of claim 10, further comprising:
storing the one or more flow paths in a data store; and
wherein analyzing the one or more flow paths to identify network characteristics further comprises:

querying the data store based on at least one network property;
receiving a subset of the one or more flow paths from the data store in response to the query; and
comparing the subset of the one or more flow paths to at least one rule associated with at least one network property.

15. The system of claim 14, further comprising:
generating a report based on the comparison of the subset of the one or more flow paths to the at least one rule, the report including one or more suggested network configuration changes.

16. A non-transitory computer readable storage medium including instructions that, when executed by one or more processors, cause the system to:

receive state information from a plurality of network devices in a network;
parse, by a plurality of network device-specific parsers, the network information to create parsed network information;

generate a network model based on the parsed network information, wherein the network model describes how data is processed by the network;

compute one or more flow paths using the network model; and

analyze the one or more flow paths to identify network characteristics.

17. The non-transitory computer readable storage medium of claim 16, wherein the state information includes:

one or more of forwarding states, configuration files, internet protocol (IP) tables, rules received from the plurality of network devices; and

network topology data received from a user.

18. The non-transitory computer readable storage medium of claim 17, further comprising:

storing the one or more flow paths in a data store;

detecting, using at least one of the plurality of parsers, a change to the state information associated with at least one network device in the network;

updating the one or more flow paths based on the change; and

storing the updated one or more flow paths in the data store.

19. The non-transitory computer readable storage medium of claim 17, further comprising:

storing the one or more flow paths in a data store; and

wherein analyzing the one or more flow paths to identify network characteristics further comprises:

querying the data store based on at least one network property;

receiving a subset of the one or more flow paths from the data store in response to the query; and

comparing the subset of the one or more flow paths to at least one rule associated with at least one network property.

20. The non-transitory computer readable storage medium of claim 13, further comprising:

generating a report based on the comparison of the subset of the one or more flow paths to the at least one rule, the report including one or more suggested network configuration changes.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

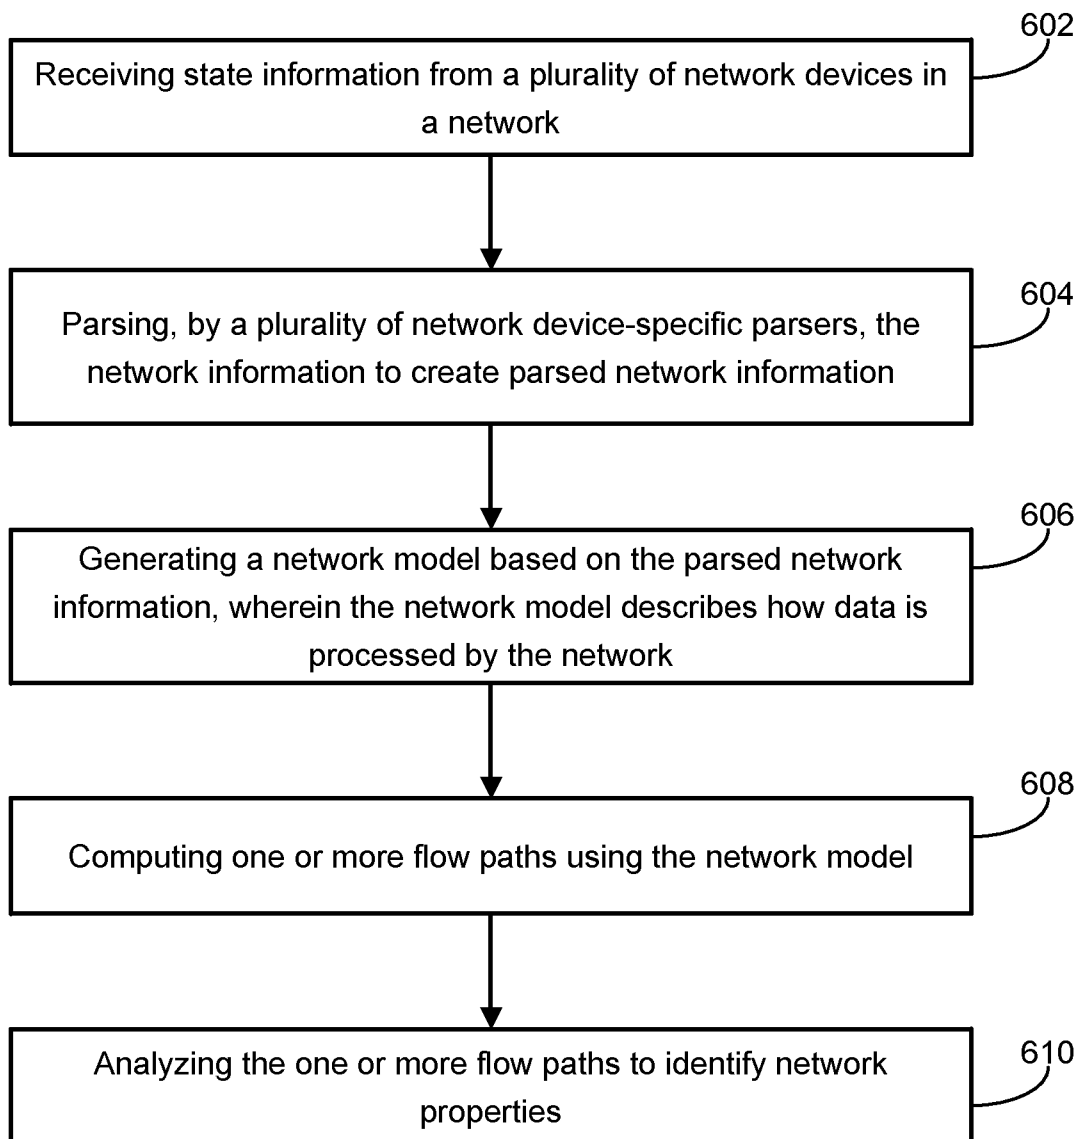

600 

FIG. 6




FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US15/42966

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - H04L 12/00, H04L 12/02, H04 L1/24 (2015.01)

CPC - H04L 41/14, H04L 41/145, H04L 41/08

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)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

IPC(8): H04L 12/00, H04L 12/02, H04 L1/24 (2015.01)

CPC: H04L 41/14, H04L 41/145, H04L 41/08

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

PatSeer (US, EP, WO, JP, DE, GB, CN, FR, KR, ES, AU, IN, CA, INPADOC Data)

Keywords: Receive, Network, Parameters, Parse, Generate, Model, Testing

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/0165207 A1 (ENGEL, G et. al.) 12 June 2014, Figure 2A, paragraphs [0044]-[0047], [0049], [0050], [0073], [0088], [0099].	1, 2, 10, 11, 16, 17
---		-----
Y		3-9, 12-15, 18-20
Y	US 2005/0076104 A1 (LISKOV, B et. al.) 07 April 2005, paragraph [0010].	3, 12
Y	US 2004/0199808 A1 (FREIMUTH, D et. al.) 07 October 2004, paragraphs [0009], [0010].	4
Y	US 6512824 B1 (HSIEH, J et. al.) 28 January 2003, Column 7, Lines 5-15.	5, 13, 18
Y	US 2011/0283016 A1 (UCHIDA, S) 17 November 2011, Figure 2, Claim 8, paragraphs [0102], [0103].	6, 7, 14, 15, 19, 20
Y	US 7937755 B1 (GURUSWAMY, K) 03 May 2011, Abstract, Column 8, Lines 55-60.	8, 9

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

28 September 2015 (28.09.2015)

Date of mailing of the international search report

23 OCT 2015

Name and mailing address of the ISA/

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774