



- (51) International Patent Classification:
G06F 17/10 (2006.01)
- (21) International Application Number:
PCT/US2017/057751
- (22) International Filing Date:
21 October 2017 (21.10.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/411,364 21 October 2016 (21.10.2016) US
15/789,887 20 October 2017 (20.10.2017) US
- (71) Applicant: FORWARD NETWORKS, INC. [US/US];
550 California Avenue, Suite 200, Palo Alto, CA 94306 (US).
- (72) Inventors: KAZEMIAN, Peyman; 745 Mediterranean Lane, Redwood City, CA 94065 (US). GANJISAFFAR, Yasser; 1573 Ashcroft Way, Sunnyvale, CA 94087 (US). RADHAKRISHNAN, Sivasankar; 2585 El Camino Real, Apt 211, Santa Clara, CA 95051 (US). HANDIGOL, Nikhil; 1866 Trento Loop, Milpitas, CA 95035 (US).
- (74) Agent: LOHR, Jason et al.; Hogan Lovells US LLP, 3 Embarcadero Center, 15th Floor, San Francisco, CA 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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,

(54) Title: SYSTEMS AND METHODS FOR SCALABLE NETWORK MODELING




FIG. 1

(57) Abstract: Embodiments of the present invention can use a protocol-independent, vendor-independent, efficient and scalable abstraction model for representing the forwarding functionality of networks. Such a model can be used for systematic analysis and verification of networks. Packet header values may be represented as groups of one or more wildcarded bit strings, where unspecified header values are assumed to be fully wildcarded. This representation can describe many combinations of packets in a space-efficient way, enables more efficient tracing and transformation operations, and can even represent traffic from large internet routing tables efficiently. As a result of the scalability benefits of this more effective way to store and operate on packet collections, network modeling can scale to some of the largest, most complicated networks those where the benefits are the greatest.

OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Declarations under Rule 4.17:

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

SYSTEMS AND METHODS FOR SCALABLE NETWORK MODELING

CROSS-REFERENCES TO RELATED APPLICATIONS

[0001] This application claims priority to U.S. Provisional Patent Application No. 5 62/411,364, filed October 21, 2016, titled “SYSTEM AND METHOD FOR SCALABLE NETWORK MODELING,” by Peyman Kazemian, 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 large company may typically 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 state associated with each network device that defines the overall network behavior is distributed throughout the network and is a result of emergent interactions between devices that are configured in vendor- and device-type-dependent ways. Second, the distributed forwarding states are difficult to monitor – often requiring the network 20 administrator to manually login to the device and conduct low-level tests. Third, multiple administrators or users can edit the forwarding states at the same time, resulting in inconsistent configuration, followed by unexpected forwarding states.

[0004] Network models provide a software copy of a network’s behavior, upon which a network administrator can better understand current behavior, troubleshoot 25 problems, analyze whether a network is behaving according to policy, and even try out ways to improve performance and fault tolerance. However, the larger the network, the more difficult it can be to model, because not all models scale easily or accurately. Thus,

techniques and methods of scaling network models are needed to aid in diagnosing problems in large, highly-complex networks, where the advantages of network modeling may be the most beneficial.

BRIEF DESCRIPTION OF THE DRAWINGS

5 [0005] Various embodiments or examples of the invention are disclosed in the following detailed description and the accompanying drawings:

[0006] FIG. 1 is a high-level diagram illustrating a network of devices, in accordance with an embodiment of the present invention.

10 [0007] FIG. 2 illustrates example representations of encoding a collection of packets, in accordance with an embodiment of the present invention.

[0008] FIGs. 3(A)-(C) illustrate example representations of encoding fields of packets using wildcards, in accordance with an embodiment of the present invention.

[0009] FIG. 4 illustrates an example of tracing a packet collection through a rule table, in accordance with an embodiment of the present invention.

15 [00010] FIGs. 5(A)-(C) illustrate examples of rule tables that experience a cross-product problem, in accordance with an embodiment of the present invention.

[00011] FIGs. 6(A)-(C) illustrate examples of IP addresses corresponding to the network model, in accordance with an embodiment of the present invention.

20 [00012] FIG. 7 is a high-level flow diagram illustrating an example method of analyzing networks, in accordance with an embodiment of the present invention.

[00013] FIG. 8 is a high-level block diagram illustrating a computer system, in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[00014] 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
5 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.

[00015] A detailed description of one or more examples is provided below along with accompanying figures. The detailed description is provided in connection with such
10 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
15 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.

[00016] Within a network, data may be transmitted via packets from one networking device to another. Networking devices forward packets based on their header
20 bits, and network behavior can be modeled using a plurality of different types of models. In one embodiment, a data-plane model may be implemented to model the network. In a data-plane model, the network behavior of the data being transmitted is represented by packets and their behavior through the network may be referred to as traffic, flow paths, traffic flow, etc. In some embodiments, state information (e.g., configuration data,
25 forwarding states, IP tables, rules, network topology information, etc.) may be associated with a plurality of devices in a network. State information, network information, and/or any other relevant information about traffic on the network may be received from devices in a network, or obtained from another entity or network-data source. The state

information may be parsed by a network modeling system and used to generate a network model, such as a data-plane model. Generally, network models describe how the network processes packets of data. Using the model, the network modeling system may be able to identify possible flow paths taken by packets of data through the network. The network modeling system may then use the information gathered from generating the network model 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.

[00017] In some embodiments, the network modeling system may generate two different types of models: a data-plane model or a control-plane model. The data plane, the control plane and the management plane are the three basic components of telecommunications architecture. The control plane and management plane serve the data plane, which bears the traffic that the network exists to carry. The data plane is the part of a network that carries user traffic. The data plane may also be referred to as the user plane, forwarding plane, carrier plane or bearer plane. The data plane enables data transfer to and from client computing devices, and may handle multiple data transmissions through multiple protocols. In conventional networking, all three planes may be typically implemented in the firmware of routers and switches. Software-defined networking (SDN) decouples the data and control planes and implements the control plane in software instead, which enables programmatic access to make network administration much more flexible. Data-plane models may be built by combining network state data with device-level forwarding behavior. For example, a data-plane model may consider forwarding lookups and output actions of individual devices and may interconnect the individual devices along known links based on the forwarding lookups and outputs.

[00018] In another embodiment, the network modeling system may generate a control-plane model. Control-plane models may be generated and built by combining network configuration data with protocol behavior. Network models may serve to

represent the individual behavior of packets as they travel through the network. The overall network behavior may be represented by considering the individual behavior of individual network devices using one or more rule tables, in which the rule tables may dictate how the packets travel through the network devices on the network. For example, by dictating how the packets travel through the network via the network devices, network links or flow paths may be known. Then, the network may be modeled by interconnecting the network devices along the known network links, as shown in FIG. 1.

[00019] FIG. 1 is a high-level diagram illustrating a network 100 of networked devices 102, 104, and 106, in accordance with an embodiment of the present invention. The network 100 may include a plurality of networked devices 102, 104, and 106 that are linked by flow paths (e.g., 108). As an example, each networked device 106 may contain a set of rule tables 110, including Table 0 and Table 1. Each rule table in 110 consists of a sequence of rules, for example, transforming some fields in the packet, forwarding them to another table or device port, or dropping the packets. The sequence of rules includes match specifiers 112, 116 and specifies corresponding actions 114, 118 to be applied to the packets when the packets satisfy the match specifiers 112, 116, respectively. As an illustrative example, a packet or a group of packets arriving at network device 106 may be subject to Table 0 and Table 1 in 110, in which the packet may be compared to match specifier 112; in this example, to see if a packet header value matches “11xx..0x”. If match specifier 112 is satisfied, then the corresponding action 114 is executed on the packet. In this example, the packet header value may then be rewritten with “1x01xx..x1” and then sent to port 2. In another illustrative example, the packet may be compared to match specifier 116, and if the match specifier 116 is satisfied, then the corresponding action 118 is executed; in this example, the packet may be dropped.

[00020] According to various embodiments, the network modeling system may generate two different types of models: a data-plane model or a control-plane model. In a control-plane model, the data-plane rules may be inferred from control-plane configuration and behavior, while in a data-plane model, they may be based on state

information that may be collected directly from the devices, obtained from other sources, or obtained from other means. In general, a network model may be built using any combination of control-plane configuration and data-plane state. A network modeling system generates a network model in software to enable a network administrator to trace one or more packets through the network for training, debugging, and analysis purposes. Tracing an individual packet and computing its behavior through the network can be relatively straightforward: existing network analysis systems can identify a packet, match the packet against rules in the first rule table, apply any resulting actions, follow links to other rule tables or other devices, continue tracing through them, and stop when the packet exits the network or is dropped. However, building a complete behavioral model of the network to cover every possible packet and flow path becomes infeasible for large networks. The number of distinct potential packets and paths in a large network leads to long computation times and exorbitant computing resources. Despite the challenges of modeling large or entire networks, modeling the entire network behavior is desirable for several reasons; for example, detecting errors like loops without identifying the packets that might trigger the errors in advance, and providing information about reachability (e.g., finding all packets that that can reach a host in the network). Even tracing a restricted subset of packets, such as all IPv4 packets, can be computationally infeasible when the set of packets to be traced is large.

[00021] A user or network administrator may consult the network model to discover what specific action has happened to a packet and where that action occurred within the network. For example, the user or network administrator may use the network model to determine where a packet is dropped within a network, or to verify that multiple paths are available between two hosts. In some instances, several packets (e.g., group of packets) that have entered a network or have arrived at a specific device in the network are subsequently forwarded to travel through the same flow path. Directing and forwarding the group of packets individually (e.g., directing and forwarding each individual packet one by one) would make behavior analysis infeasible in any large

network. Not only would processing each packet individually be slow and tedious, it would be an inefficient and wasteful use of processing resources. As such, to help make a complete analysis of an entire and/or large network model computationally feasible, when the group of packets enter the network or arrive at a specific device in the network and are forwarded in the same way, the network model may group the packets together into a collection of packets. The network model may determine that the group of packets is to be collectively forwarded in the same manner (e.g., flow path), and continue to process and refer to them as a collection of packets. For example, complex mathematical operations may be applied in each rule table using collections of packets, including matching a packet collection against rule table rules, finding the intersection between a collection of packets and the matching packets of a rule, finding the remaining collection of packets which has not matched a rule, applying actions that transform the collection of packets, and finding the intersection between two collections of packets. All of these operations may be performed on the packet collection as a whole, rather than for each individual potential packet.

[00022] Embodiments of the present invention provide novel techniques of representing a collection of packets that result in an improved model implementation that can be scaled to larger and more complex networks than existing techniques. Determining how to represent a collection of packets (i.e., packet collection) can significantly affect the resulting network model analysis; for example, the representation can affect which mathematical operations to execute, and representation can vary in implementation complexity, memory requirements, and execution efficiency. The network model according to various embodiments provides a means of constructing and analyzing a network model in a reasonable amount of time by adopting an improved packet representation for grouping packets and collections of packets, and performing mathematical operations using the improved representation.

[00023] FIG. 2 illustrates an example 200 of encoding a collection of packets, in accordance with an embodiment of the present invention. There are multiple ways in

which to represent packets such that their behavior within a network can be observed and analyzed. A collection of packets 212 may be routed between an origin and a destination on a network. The packets then travel around the network, locally or globally, wired or wirelessly, and may then be unbundled by another node on the network. During network
5 troubleshooting, administrators and engineers analyzing the networks may need to use a protocol analyzer to open up or look into the packets and to observe their contents. Each packet in the collection 212 includes various fields 202, 204, 206, 208, and 210 that may hold content that can provide important insight to the network administrator as to the packet's behavior on the network. For example, a source MAC address 202 (SRC MAC),
10 a destination MAC address 204 (DST MAC), a source IP address 208 (SRC IP), a destination IP address 210 (DST IP), and other fields 206. As shown in FIG. 2, the packets in the collection 212 differ in their packet header values for field 206, such that the values can be combined using wildcard bits. As shown in FIG. 2, the virtual packet header 214 illustrates a representation of encoding the collection of packets 212 where
15 the 'x' bit may operate as a wildcard bit such that the 'x' bit can represent both 0 and 1 simultaneously, as shown in the field 216.

[00024] FIGs. 3(A)-(C) illustrate example representations of encoding a collection of packet fields using wildcards, in accordance with an embodiment of the present invention. As shown in FIG. 3(A), contiguous power-of-two regions may be represented
20 by a single wildcarded range. A collection of packets may be represented by including wildcarded bit strings in the fields of headers, such that an entire header field can take any value when an 'x' is in all bit positions in the header field. In some embodiments, packet header values may be unspecified, unknown, empty, or irrelevant, and thus, may be implicitly assumed to be fully wildcarded, unless there is a restricted value for that
25 particular packet header field. Representation 300 includes using wildcard bits for header fields that have a contiguous power-of-2-aligned range. For example, an IPv4 subnet like 1.2.3.0/24 can be represented with a single wildcarded bit string that ends in 8 x bits, to represent the 256 possible IP addresses in that subnet, since the value starts at a power of

2 (0) and ends exactly one value before a power of 2 (255). As such, a packet may then be modeled as a point in the $\{0, 1\}^L$ space, where L is the length of the header. To illustrate, in representation 300, a fully wildcarded header 310 of all 'x' bits can represent packet header 302 having the value '00000000', packet header 304 having the value
5 '00000001', packet header 306 having the value '00000010', and all values up to packet header 308 having the values '11111111'.

[00025] FIG. 3(B) illustrates another example representation 320 of a collection of header field values. Wildcard bits can represent a collection of packets that differ in one or more particular bit positions, but share common values elsewhere. Two bit strings that
10 differ in exactly one bit position can be combined into one bit string by wildcarding the differing bit; for example, the values "0000" and "0010" can be represented as "00x0", where the 'x' bit indicates that the corresponding bit is wildcarded and indicates both 0 and 1. This procedure can be performed repeatedly to combine values represented as wildcarded bit strings. When two points are adjacent in this L -dimensional space, they
15 can be represented by a compact representation using wildcards. For example, "0000" and "0010" in a 4-dimensional space are adjacent together and these two points can be represented as "00x0" where the 'x' bit indicates that the corresponding bit is wildcarded and indicates both 0 and 1. To illustrate, in representation 320, consider the range of TCP ports with values from 10 to 12 referenced in packet header 322 having the value '1010',
20 packet header 324 having the value '1011', and packet header 326 having the value '1100', respectively. These three values in 322, 324, and 326 can be represented by two wildcarded bit strings 328 and 330, having the values 101x and 1100 respectively, where the preceding 12 bits of the 16-bit TCP port field are all 0's in both bit strings.

[00026] However, wildcarding may provide no benefit for representing some
25 packet collections. FIG. 3(C) illustrates examples that cannot be represented in fewer values using wildcards. Representation 340 considers two two-bit strings 342 and 344, having the values 01 and 10, respectively. No single wildcarded bit string can represent these two values; two bit strings are required. Representation 340 also considers three

three-bit strings 346, 348, and 350, having the values 001, 010, and 100, respectively. None of them can be combined and three bit strings are required to represent all the values.

[00027] In existing network analysis techniques, wildcarding helps to reduce the number of bit strings needed to represent a packet collection. However, reduction all the way down to a single wildcarded bit string is rare, as individual field values, such as for IP address ranges, Ethernet types, and TCP ports, are often not combinable as illustrated in the representations of FIG. 3(C). Hence, packet collections typically contain a list of multiple bit strings. When packet collections cannot be represented in a more compact way and require many bit strings, tracing them through the network can result in a large number of expensive mathematical operations. Accordingly, the groups of packets may be limited to only those whose headers represent a contiguous region. If representing a group of packets by a single wildcarded bit string is not possible, then the network model is required to use a list of multiple bit strings to describe the packets individually instead of collectively as a group. The individual representation of each packet within a group of packets consumes more storage because it results in a much longer bit string representation. As a result, network models and analysis that refer to individual packets within a group of packets can be difficult to scale without unlimited memory and processing resources.

[00028] Furthermore, in networking, tracing packets through the network involves accessing rule tables that define the traffic flow of packets traveling within a network. Rule tables contain a number of rules where each individual rule may be represented as a condition that defines how packets should be directed; for example, if a packet satisfies a condition, that packet will be directed to a particular flow path defined for that condition. However, tracing using rule tables can be inefficient, particularly when the rule tables represent access control list (ACL) entries, such as those found in a firewall or router. If a large number of individual packet header field values are the result of tracing packets through device A, and then in device B, the packets are treated based on different fields.

As a result, the number of potential packet values to trace can increase rapidly, consuming extra memory and unwanted computation time.

[00029] FIG. 4 illustrates an example of a packet collection entering a rule table, in accordance with an embodiment of the present invention. In example 400, an input packet collection 402 may enter a device 404 in a network model, the device including rule table 408 “Table 0” and rule table 410 “Table 1.” To find out if the first rule “Match0” 412 in rule table 408 “Table 0” matches the packet collection, a mathematical operation, such as an intersection operation, must be performed, as shown at point 440. This operation for “Match 0” 412 may be executed not simply to determine whether there is any overlap, and hence a match, but also to determine the exact overlap. The exact overlap may be used later when applying “Actions0” 414 at point 442, which may potentially modify the packet collection. Subsequently, tracing can continue (1) in the next rule table 410 “Table 1” with a potentially-modified packet collection in point 446, as well as (2) at the next rule (e.g., “Match1” 418 and “Actions1” 420) in the table 408 “Table 0” with the non-overlapping portion of the packet collection, shown at point 448. At point 446, when “Match0” 426 is satisfied and “Actions0 430” is taken, then the output packet collection 406 may exit the device 404, as shown at point 450. As packet collections contain more bit strings, the number of computation cycles to complete the intersection operation increases, such that it can easily dominate the overall computation time. In addition, applying actions, like rewriting header fields at point 442, can become expensive with more bit strings.

[00030] These problems with scaling packet tracing on network models become particularly noticeable when representing a collection of packets denoted by a cross-product of values across header fields. Called the “Cross-Product Problem”, this problem occurs when representing a collection of packets that includes the entire Cartesian product of values across two or more header fields. Enumerating the combinations of header values results in a large number of wildcarded bit strings to represent all the packets in the collection, which adds computation cost to each intersection and rewrite

operation. If the cross-product problem is not addressed through careful packet collection representation and efficient mathematical operations on the collection, the network model may not scale to a practically large size. Either the tracing process will take too long, or may require too many resources, to be useful.

5 **[00031]** FIGs. 5(A)-(C) illustrate examples of computing a network model, in accordance with an embodiment of the present invention. FIGs. 6(A)-(C) illustrate examples of a collection of packets denoted by a cross-product of IP destination and TCP destination port header fields, in accordance with an embodiment of the present invention. For example, the cross-product problem can be illustrated in at least three
10 scenarios while computing a network model: (1) inside a single table in FIG. 5(A); (2) across rule tables in a single device in FIG. 5(B); and (3) across independent rule tables in different devices in FIG. 5(C).

[00032] In FIG. 5(A), the cross-product problem may be seen within a single rule table 502 in network device 500, where individual rules 506, 508, 510, and 512 may be
15 relevant for forwarding a packet. In an example, in corresponding FIG. 6(A), a firewall rule table may match (e.g., 502 in FIG. 5(A)) on two fields, IP destination address 602 “ip_dst” and TCP destination port 610 “tcp_dst_port.” The IP destination address 602 field may include IP addresses 604, 606, and 608, and the TCP destination port 610 field may include ports 612 and 614. The three IP destination subnets (e.g., 604, 606, and 608)
20 and two TCP destination port values (e.g., 612, 614) are not combinable using the wildcard bit string representation. As a result, the set of packets that the rule matches may be represented as six different bit strings to cover each possible combination of IP destination subnet and TCP destination port value, as shown in FIG. 6(B). In FIG. 6(B), each bit string in 630 may include a pair 622 of an IP address and a TCP port, an IP
25 destination address 624 “ip_dst,” a TCP destination port 626 “tcp_dst_port”, and additional wildcard fields 628. Subsequently, each of these six-bit strings 630 produces a different bit string as a result of a tracing operation. Each of these bit strings may yield

even more unique bit strings in the next rule table within the same device, or when the traffic is traced through another device.

[00033] FIG. 5(B) illustrates when the cross-product problem manifests across different rule tables 522 and 524 in a single device 520, when the rules in rule tables match on different header fields. Rule table 522 “table0” may include rules matching field values 528 and 530, while rule table 524 “table1” may include rules matching field values 534 and 536. For example, an Access Control List (ACL) table might contain rules that match on the TCP destination port and permit the traffic, while an IP routing table matches on the IP destination address to forward traffic to an output port. Similar to the previous example illustrated in FIG. 6(A), if there are three IP destination subnets and two TCP destination ports, when a collection of packets is traced through the pipeline of tables 522 and 524 in the device 520 and the match criteria are applied at each rule, the analysis results in 6 different wildcard-bit strings to represent all the packets that may come out of the device, as shown in FIG. 6(B). As with the example shown in FIG. 5(A), tracing through the rule tables yields a significant increase in the number of bit strings, which continues with later devices in the path.

[00034] Similarly, as shown in FIG. 5(C), the cross-product problem also manifests across devices 540 and 542, when a network contains rules that match on different header fields. For example, device 540 might contain rule table 544 with two rules 550 and 552, where the action is to set the output port based on the IP destination address. Another device 542 may contain a firewall rule table 546 with two rules 556 and 558, where the action is to block some traffic based on TCP destination ports. When each packet collection traced through device 540 reaches device 542, the end result is a cross-product of all the rules. As with examples FIG. 5(A) and FIG. 5(B), tracing through the rule tables yields a significant increase in the number of bit strings, which continues with later devices in the path. Regardless of the number of devices or tables involved in a cross-product, the resulting problem is the same.

[00035] To illustrate the Cross-Product Problem, consider the following header field value lists:

my-ip-list = [1.2.3.0/24 10.20.30.0/24 100.200.300.0/24 ...] (*m* elements where *m* = 1, 2, ..., *M*)

5 my-allowed_ports = [80 443 22 8443 ...] (*n* elements where *n* = 1, 2, ..., *N*)

[00036] These lists may be used to create an ACL rule, such as this one:

```
match =
ip_src: my-ip-list;
ip_dst: my-ip-list;
10 ip_proto=TCP/UDP;
tp_dst_port=my-allowed_ports;
tp_src_port=my-allowed_ports;

action =
allow
```

15 **[00037]** The ACL rule above defines the traffic flow of the packet with a match and a corresponding action, such that if the packet matches the parameters defined by the match (e.g., ip_src, ip_dst, ip_proto, tp_dst_port, tp_src_port), then the corresponding action is to allow the packet to pass. *M* elements may be present to match each IP field and *n* elements may be present to match each port. This configuration requires m^2n^2

20 unique bit strings to represent packets that match the ACL rule, which for larger *m* and *n*, can be an immense number of unique bit-string rules. Furthermore, if each rule is considered independently, then any later processing stages will see a much larger number of bit strings, resulting in extra CPU, memory, and disk usage.

[00038] In another example, the situation when packet regions are continuous, representing groups of IP prefixes can also be challenging. As illustrated below, the IP table includes 3 distinct next-hop addresses, A, B, and C:

10.0.0.0/8 --> A

11.0.0.0/8 --> B

12.0.0.0/8 --> C

14.0.0.0/8 --> A

5 15.0.0.0/8 --> B

16.0.0.0/8 --> C

[00039] This configuration may include six unique bit rules to represent the configured behavior, since each group of rules with the same next-hop is non-contiguous. However, a network model according to embodiments of the present application may represent the IP table as three groups, such that additional work in later processing stages may be reduced to a minimum.

10

[00040] Various applications may benefit from the advantages of various embodiments of the present invention. For example, data-plane modeling according to various embodiments may be used to verify that the software code enables the network to function properly. Alternatively, in a production environment, data may be used to validate the current behavior of the network or test the effect of proposed changes.

15

[00041] According to various embodiments, the network modeling system may use a protocol-independent, vendor-independent, efficient and scalable abstraction model for representing the forwarding functionality of networks. Such a model may be used for systematic analysis and verification of networks. The network model may include forwarding rules that are used to determine how packets are processed by devices in the network. In some embodiments, the network model may be periodically updated. In some embodiments, instead of recomputing the entire network model each time the actual network's state changes, the network model may 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

20

25

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.

[00042] According to an embodiment, the network modeling system may address the Cross-Product Problem by implementing an efficient representation for a collection of packets, where each individual header field is represented as a collection of wildcarded bit string values that are possible for that field. The packets included in the collection are packets with header values limited to values specified for the respective header fields in the representation. The wildcarded bit strings for each field are only as wide as the length of that field. For example, for the collection of packets defined in FIG. 6(A), the packet collection may be represented by a single entry containing all the IP destination values 602 and another entry containing all the TCP destination ports 610. Mathematical operations may then be performed on the entire packet collection rather than on individual wildcarded bit strings. While performing these operations, many of them do not require enumerating the cross product to compute the long wildcarded bit strings that cover all the header fields in the packet. The operations may operate directly on the overall packet collection representation or on the set of values for a particular header field. As a result, the network modeling system according to various embodiments may be benefitted by significant improvements in CPU, memory, disk usage costs.

[00043] In an embodiment, as shown in FIG. 6(C), the network modeling system may represent packet header values as groups, for example the representation 660 may include IP destination 662 "ip_dst" to group IP addresses 668, 670 and 672, and TCP destination ports 664 to group ports p1 and p2. Not shown on the diagram, unspecified header values may be assumed to be implicitly wildcarded, with no storage used to represent them. In this way according to various embodiments, a single description can represent many combinations of packets, leading to a network model with improved efficiency and scalability. This representation may then avoid problems with bit-string representations, such as the cross-product rule problem, as well as the inability to

represent traffic from large internet routing tables efficiently. The representation 660 may also be more space-efficient when several header fields are fully wildcarded, both because fewer bit strings are needed, and because bit strings for a particular header field can be stored once and not repeated for each possible combination (cross product) with
5 bit string values of other header fields.

[00044] In some embodiments, fields used in the packet collection representation may not correspond to well-known header fields typically used for packet forwarding. For example, a custom tag may be added within a network to denote an ingress point or a class of traffic. This custom tag may be included in the packet collection. The network
10 model may also introduce its own set of virtual fields to denote some classification of packets that is relevant to the model.

[00045] Although there may be challenges associated with representing packet collections with wildcards, embodiments of the present invention overcome these challenges and provide additional advantages. For example, mathematical operations may
15 be performed on the packet collection, such as analyzing which packets match a rule, or combining two packet collections. According to various embodiments, when these mathematical operations are executed, the resulting output may be represented in a form that satisfies the property that all cross-products of values across the header fields should be included in the result. Alternately, the representation may be able to indicate which
20 packets are excluded from the cross-product. In another example, a network modeling system according to various embodiments may correctly implement each operation, including combination and rewrite operations, despite the large number of ways in which collections of header values may intersect or combine.

[00046] According to various embodiments, the packet collection representation
25 may be implemented in Java, C++, C#, or any suitable programming language, on any computing hardware, such as a general-purpose processor or graphics processor. In some

embodiments, the invention may be implemented directly in hardware, via field-programmable gate array or application-specific integrated circuit.

[00047] FIG. 7 is a block flow diagram illustrating an example method 700 of processing and analyzing packets in a network modeling system according to various
5 embodiments. A networking device, according to various embodiments, may be modeled as a collection of rule tables where each rule table includes a list of rules and each rule includes a match and a collection of actions, such as the rule tables shown in FIG. 1. At step 702, the network modeling system may obtain network information associated with network devices on a network such as the network shown in FIG. 1. The network
10 information may be obtained from a third party entity or collected from the network devices in the network. As shown in FIG. 1, each network device may include one or more rule tables that describe forwarding behavior for packets within the network. Each rule table may include a set of rules, where the rules may be applied in a specific order or in any order (e.g., unordered). In some embodiments, the rules may be implicit or
15 explicit. Each rule is comprised of a match specifier and corresponding actions, such that when a packet or group of packets matches the match specifier, the rule is applied to the packet and the corresponding actions are performed. The forwarding behavior defined in the rule tables of each device in the network may, collectively, cause and define the traffic flow of the network.

[00048] At 704, the network modeling system may identify a group of packets to
20 analyze in a network model. In one embodiment, the group of packets may be identified as entering the network or arriving at a network at particular devices or ports. Identifying the group of packets may be based at least in part on the network information obtained, where the network information includes at least the one or more rule tables and network
25 topology. In some embodiments, the group of packets may be identified based on information obtained about troubleshooting a collection a packets in the network analysis system.

[00049] At 706, the network modeling system may then create a virtual packet header to represent the group or collection of packets identified. The virtual packet header, in order to represent the collection of packets, may include wildcarded bits according to various embodiments. The virtual packet header may include one or more
5 packet header groups, where at least one of the one or more packet header groups represents a union of two or more packet header values.

[00050] The virtual packet header may then be pushed through a network model to simulate network behavior, at 708. When the virtual packet header is pushed through the network, it may enter a network device, for example as shown in FIG. 1. As the virtual
10 packet header enters the network device, the rule tables of the network device may be applied to the virtual packet header. A rule match determines the collection of packets for which this rule applies and may be represented as an HRegion. At 710, a rule match occurs by comparing values in the collection of packets with the values specified in the match specifier. If a collection of packets belongs to this region and matches this rule,
15 then the corresponding actions specified in that rule are applied to that collection of packets. Within a rule table, rules (e.g., match and corresponding action) may be processed in a particular order or in any order. The rules may be explicit in the rule table or implicit. The corresponding actions of rules may be explicit or implicit and include transformations, which can be any operation performed on the virtual packet header that
20 changes at least one packet header value. In some embodiments, the rule tables may be implicit. At 712, the network modeling system may then determine the flow paths of packets traveling in the network based on the behavior of the virtual packet header through the network model.

[00051] In one embodiment, the network modeling system may model packet
25 collections by using a virtual packet header to represent the collection. A virtual packet header specifies the union of header fields over defined protocols. For example, this virtual packet header may include fields that are present in the headers of both IPv4 and IPv6 packets. When representing an IPv4 packet, a virtual packet header may wildcard all

bits that are related to other protocols like IPv6. Each field in the virtual packet header has a corresponding unique identifier, called a field ID. Accordingly, the virtual packet header may identify the packet or class of packets by using wildcards as the packet header values. By doing so, the virtual packet header may have an overall bit size that is smaller than the sum of each individual packet header in the group of packets represented by the virtual packet header, if each had to be processed individually.

[00052] For any given packet, many fields may be fully wildcarded. Instead of representing packet collections using sequences of {0, 1, x} bits, the network modeling system according to various embodiments may use a sparse representation consisting of a collection of field IDs and their corresponding values. Fields not present in this collection may be assumed to be fully wildcarded (i.e., all bits are represented with 'x' values). The value of each field is a collection of possibly-wildcarded values. For example, the virtual packet header may include three fields: f1, f2, and f3, and each of these fields may include four bits. In this example, the example packet collection may be represented using the following compact representation:

$$(f1, \{\text{"0000"}\}), (f3, \{\text{"0001"}, \text{"00x0"}\}) \}$$

[00053] The compact representation above is equivalent to:

$$\{\text{"0000,xxxx,0001"}, \text{"0000,xxxx,00x0"} \}$$

[00054] The first value above ("0000,xxxx,0001") includes four wildcarded bits and hence represents 16 distinct values. Similarly, the second value ("0000,xxxx,00x0") includes 5 wildcarded bits and hence represents 32 distinct values. As a result, the compact representation represents 48 packets in this example.

[00055] Embodiments of the present invention may use header bits ("HBit") to refer to bits that can have one of three values in {0, 1, x}. A sequence of these Hbits may be referred to as a header sequence ("HSeq"). For example, "00x0" in the above example is an HSeq. The value of a header field may be represented as a set of HSeqs which may

be referred to as a header sequence collection (“HSeqCollection”), e.g., “{“0001”, “00x0”}”. A collection of header fields and their corresponding values, as in the above example, may represent a non-contiguous region. Embodiments may use a header region (“HRegion”) to refer to such regions, for example: “{ (f1, {“0000”}), (f3, {“0001”, “00x0”}) }”. In this notation, an Hregion is a map from packet fields to HSeqCollections. If a packet field is fully wildcarded in an Hregion, it will not have an entry in the corresponding map. In other words, the absence of an entry for a packet field in this map indicates that the field is fully wildcarded.

[00056] A collection of packets may be represented, in an embodiment, by input_region (an HRegion) when the collection of packets enters a network device with a rule table, for example, as shown in FIG. 4. The rule table may contain a set of rules, such as [r1, r2, r3], and each rule may have a respective match, such as [match1, match2, match3]. In order to find the fraction of the input_region that matches r1, the network modeling system according to various embodiments may determine the intersection of input_region and match1. The intersection of HRegions is described in greater detail further in the present application. The collection of packets that are in input_region but do not match1 may then be carried forward to the next rule(s).

[00057] To represent packet collections like those carried forward to next rules, which may not be modelable as a single HRegion, according to various embodiments the network modeling system may define a header expression (“HEXpr”). HEXpr may represent an HRegion minus the union of a collection of other HRegions. In other words, the HEXpr represents the difference between a “first packet collection” and one or more “other packet collections”. To illustrate, in the above example, ‘input_region - { match1 }’ (referred to as “hse1”) may be the HEXpr representing the collection of packets that are sent to r2 for processing. In order to find which packets in this collection need to be processed with r2, the networking modeling system may determine the intersection of packets represented by hse1 and the packets represented by match2, which is an HRegion. The intersection of an HEXpr and an HRegion is an operation that may be

invoked when tracing a group of packets that can enter each rule in a rule table. Similarly, 'input_region - { match1, match2 }' (referred to as "hse2") may be the HExpr representing the collection of packets that can match rule r3. In hse2, 'input_region' according to various embodiments may be referred to as a minuend, and 'match1' and 'match2' may be referred to as subtrahends. The intersection of an HExpr and an HRegion is described in greater detail further in the present application.

[00058] In one embodiment, the networking modeling system may perform transformation operations on packet collections. Before a transformation operation, a packet collection (i.e., represented by an HExpr) may match a rule in one of the rule tables of a device on the network, as in the example shown in FIG. 4. Subsequently, the actions specified in that rule may be applied to that packet collection (i.e., a transformation action). These actions may transform the input HExpr to an output HExpr. In addition to HExpr transformations, the actions in the rule table can specify how packet collections are forwarded to device ports or dropped. For example, when a device has the same behavior for all packets in a packet collection arriving at specific ports, the behavior of the device may be modeled using the following format:

$$(\text{inputHExpr}, \text{inputPorts}) \rightarrow (\text{outputHExpr}, \text{outputPorts})$$

[00059] The above representation may indicate that when the collection of packets represented by 'inputHExpr' arrives at any of the ports indicated by 'inputPorts', it is transformed to 'outputHExpr' and forwarded to 'outputPorts'. HExpr transformations can be expressed using rewrite and move actions.

[00060] One example of an HExpr transformation action is a rewrite action, which rewrites the value of a single field in an HExpr to another value. To illustrate, there may be the following HExpr:

$$\{ (f1, \{ "0000" \}), (f3, \{ "0001", "00x0" \}) \}$$

[00061] This HExpr may be rewritten such that “f1” is replaced with {“0001”}, which results in the following HExpr:

$$\{ (f1, \{“0001”\}), (f3, \{“0001”, “00x0”\}) \}$$

[00062] Subsequently, the field “f2” may be rewritten to {“1111”}, resulting in:

5
$$\{ (f1, \{“0000”\}), (f2, \{“1111”\}), (f3, \{“0001”, “00x0”\}) \}$$

[00063] Lastly, the field “f3” may be rewritten to “xxxx”. Because “f3” is written to be a fully wildcarded value, the resulting fully rewritten HExpr is:

$$\{ (f1, \{“0000”\}) \}$$

[00064] Rewriting a field in an HExpr may involve changing at least the value of the minuend of the HExpr, which is described in greater detail further in the present application. In various embodiments, the rewrite operation may be performed in cases that are more complex where HExprs that are being rewritten have subtrahends. Multiple-field rewrites may be considered as a sequence of single-field rewrite actions.

10

[00065] HExpr transformations may also be expressed as a move action. A move action moves the values of a field to another field. To illustrate, there may be the following HExpr:

15

$$\{ (f1, \{“0000”\}), (f3, \{“0001”, “00x0”\}) \}$$

[00066] Moving the value of “f1” to “f4” may result in the following HExpr:

$$\{ (f3, \{“0001”, “00x0”\}), (f4, \{“0000”\}) \}$$

[00067] Moving the value of a field to another HExpr is described in greater detail further in the present application, including how the operation is executed in the presence of subtrahends.

20

[00068] In another embodiment of the present application, another operation that the network modeling system may perform on packets includes intersecting two HExprs. Intersecting HExprs may be a required operation for tracing packet collections through rule tables. To complete such an intersection requires the ability to intersect HSeqs,
 5 which requires the ability to intersect HBits.

[00069] The following table may be used to intersect two HBit values:

| Bit 1 | Bit 2 | Result |
|-------|-------|--------|
| 0 | 0 | 0 |
| 0 | 1 | empty |
| 1 | 0 | empty |
| 1 | 1 | 1 |
| x | 0 | 0 |
| 0 | x | 0 |
| x | 1 | 1 |
| 1 | x | 1 |
| x | x | x |

[00070] According to various embodiments in modeling packet collections, another operation that the network modeling system may perform on packets is intersecting two header sequences (“HSeqs”), where an HSeq is a sequence of HBits.

10 Two HSeqs may be combined if and only if they have the same number of bits. Corresponding bits at each position are intersected and the result is placed in that position in the resulting HSeq. If the intersection of two HBits in the same position is empty, the two HSeqs also have an empty intersection. To illustrate, examples of intersecting two HBits are outlined below:

| HSeq 1 | HSeq 2 | Intersection Result |
|--------|--------|---------------------|
| xxxx | 0000 | 0000 |
| 0x00 | x100 | 0100 |
| 0000 | 1111 | empty |

15 **[00071]** According to another embodiment in modeling packet collections, another operation that the networking modeling system may perform on packets is intersecting

two header sequence collections (“HSeqCollections”), where an HSeqCollection is a collection of HSeqs. For each HSeq pair (s1, s2) where s1 belongs to the first HSeqCollection and s2 belongs to the second HSeqCollection, the intersection of s1 and s2 is computed as disclosed using the same technique for intersecting two HSeqs. If the result of this intersection is nonempty, then it is added to the resulting HSeqCollection.

[00072] To illustrate, there may be two HSeqCollections, as shown below:

$$c1 = \{s11, s12\}$$

$$c2 = \{s21, s22\}$$

[00073] In the above examples of the two HSeqCollections, s11, s12, s21, and s22 may represent different HSeqs. To intersect these two HSeqCollections, the intersections of all possible pairs of HSeqs are considered below:

$$\{s11 \cap s21, s11 \cap s22, s12 \cap s21, s12 \cap s22\}$$

[00074] If the result of each pair’s intersection is nonempty, it may be added to the resulting HSeqCollection.

[00075] Another operation that the network modeling system may perform on packets according for various embodiments includes intersection two header regions (“HRegions”), where an HRegion is a map from packet fields to HSeqCollections. If a packet field is fully wildcarded in an HRegion, it may not have an entry in the corresponding map. Intersecting two HRegions may performed by intersecting the corresponding HSeqCollection values of each field. If a packet field is fully wildcarded in only one of the HRegions, then its value in the intersection result may come from the value in the other HRegion. If it is not fully wildcarded in both the HRegions, its value in the intersection result may be the intersection of the HSeqCollection values in the two HRegions computed using the method explained in the “Intersection of two HSeqCollections” section. If this intersection of the HSeqCollections is empty, the intersection result of the two HRegions may also be empty.

[00076] In another embodiment, another operation that the network modeling system may perform on packets includes intersecting an HExpr and an HRegion. The result of intersecting an HExpr with an HRegion may be another HExpr. This operation may be invoked when tracing the group of packets that can enter each rule in a rule table.

5 To intersect HExpr 'hs' with HRegion 'r', first the minuend of 'hs' is intersected with 'r' using the same technique disclosed to intersect two HRegions. The result of this intersection of 'hs' and 'r', HRegion 'n', may then become the minuend of the result. If 'n' is empty, the result of intersecting 'hs' and 'r' may also be empty; otherwise, each of the subtrahends of 'hs' may be intersected with 'n', again using the technique for
10 intersecting two HRegions.

[00077] For each of these intersections, the following cases may apply. If the result is empty, the subtrahend may be ignored such that it is not included in the resulting HExpr. If the result is equal to 'n', then this subtrahend is subtracting all of 'n' from the result, and the resulting HExpr is empty. In this case, the result of intersecting 'hs' and 'r'
15 is empty, and no other subtrahends need to be processed; otherwise, the result is added as a subtrahend of the resulting HExpr.

[00078] In another embodiment, the network modeling system may perform another operation on packets including intersecting two HExprs. To intersect two HExprs, 'hs1' and 'hs2', the intersection of their minuends is first computed using the
20 technique for intersecting two HRegions, and the result may be referred to as 'm'. If that intersection is empty, the final intersection result may also empty. Otherwise, for each of the subtrahends of 'hs1' and 'hs2', the intersection of that subtrahend with 'm' may be computed. If the result is empty, that subtrahend may be discarded. If the result equals 'm', then that subtrahend is subtracting all the data points in the minuend and the result of
25 the intersection of the two HExpr becomes empty. The result in an HExpr with 'm' as its minuend and subtrahends computed using the above procedure.

[00079] Additionally, according to various embodiments, another operation that the network modeling system may perform on packets includes subtrahend simplification operations with HExprs. In the sections below, the term ‘combination’ may refer to applying a union operation, then attempting to simplify the result. The combination
 5 operation helps to speed up computation, by transforming the set of subtrahends of an HExpr to another set with fewer subtrahends, which represents the exact same collection of packets. Given that the number of subtrahends is reduced, operations that are performed on this HExpr may be computed faster.

[00080] Sometimes a combination of subtrahends results in a single subtrahend,
 10 which is equal to the minuend. In cases like this, every packet represented by the minuend may be excluded by that subtrahend and as a result, this HExpr represents an empty set. Combining subtrahends in this case can help to identify empty packet collections and any tracing operation on that group of packets can stop early. A related operation, subtrahend dropping, also helps to speed up computation and reduce memory
 15 requirements, when an HExpr containing one or more subtrahends can be represented with fewer subtrahends, and possibly none. These two subtrahend simplification operations are applied whenever we create a new HExpr, such as when tracing a collection of packets through a rule table.

[00081] When an HExpr has multiple subtrahends, it is possible that if those
 20 subtrahends are combined, the combination becomes equal to the minuend. In this case, the HExpr is empty and does not match any packet. To illustrate combining subtrahends, there may be the following HExpr:

Minuend = { (f1, {"000x"}) }

Subtrahend 1 = { (f1, {"0000"}) }

25 Subtrahend 2 = { (f1, {"0001"}) }

[00082] In this example, the two subtrahends may be reduced to { (f1, {"000x"}) } which is equal to the minuend and indicates that the HExpr is empty. An empty HExpr

indicates no packet can match it and given that the network model describes how packets are transferred in the network, the network modeling system may detect empty HExprs to determine when to stop. As such, whenever a new HExpr is created, its subtrahends may be combined as much as possible, using a technique for combining two HRegions, described below, and then check if any of the combined subtrahends is equal to the minuend.

[00083] However, the combination does not necessarily result in unique outputs. In an illustrative example, consider the following three subtrahends:

Subtrahend 1 = { (f1, {"0000"}) }
 10 Subtrahend 2 = { (f1, {"0001"}) }
 Subtrahend 3 = { (f1, {"0010"}) }

[00084] Subtrahends 1 and 2 may be combined. Similarly, subtrahends 1 and 3 may be combined. However, subtrahends 2 and 3 cannot be combined. If subtrahends 1 and 2 are combined the following results:

15 Subtrahend a = { (f1, {"000x"}) }
 Subtrahend b = { (f1, {"0010"}) }

[00085] Alternatively, if subtrahends 1 and 3 are combined, the result is:

Subtrahend c = { (f1, {"00x0"}) }
 Subtrahend d = { (f1, {"0001"}) }

20 **[00086]** All of the above collections of subtrahends represent the same collection of packets. According to various embodiments, the network models do not rely on equivalent representations for HExprs and only needs to combine subtrahends to detect empty HExprs. Hence, non-unique results of combinations do not affect correctness, as all results describe an identical collection of packets.

[00087] When a subtrahend differs from its minuend in only a single field and the values in the minuend that are not subtracted by that subtrahend can be represented by an HSeqCollection, subtrahends may be dropped. In dropping subtrahends, that subtrahend is removed and the minuend value is pruned. For example, consider an HExpr where the
 5 minuend is ‘{ (f1, {“000x”}) }’ and the subtrahend is ‘{ (f1, {“0000”}) }’. In this case, we can drop the subtrahend and update minuend to ‘{ (f1, {“0001”}) }’ because “000x” minus “0000” can be expressed as an HSeqCollection, namely “0001”.

[00088] HSeqs may also be combined in another operation that the network modeling system may perform on packets. According to various embodiments, HSeq
 10 values may be combined as much as possible to achieve a more compact and hence more efficient representation of the same packet collection. Two HSeqs may be combined in different cases, such as when (1) the values are equal; (2) one of the values subsumes the other (e.g. ‘00xx’ subsumes ‘0000’); or (3) they differ in only one bit position (e.g. ‘0000’ and ‘0001’ differ in a single bit position and can be combined to ‘000x’). When
 15 combining more than two HSeq values, instead of pairwise combinations, alternative data structures like a trie can be used for a more efficient computation of the final result.

[00089] A combination of two HSeqCollections may also be invoked by first creating the union of all the values in these two collections, putting them into a set. Then, for each pair of HSeqs in this set, the two HSeqs may be combined according to various
 20 embodiments. If the two HSeqs can be combined into one, the two values may be removed from the union and instead added to their combination. Embodiments of the present application include continuing this process until no more combination is possible. To illustrate, consider the following HSeqCollections:

c1 = { “0001”, “00x0” }
 25 c2 = { “0010”, “1000” }

[00090] The union of the values is:

$$\{ \text{"0001"}, \text{"00x0"}, \text{"0010"}, \text{"1000"} \}$$

[00091] However, in this union the 2nd and 3rd values may be combined by combining two HSeqs according to various embodiments. In this specific case, the 2nd value subsumes the 3rd value. As such, the resulting union is:

5
$$\{ \text{"0001"}, \text{"00x0"}, \text{"1000"} \}$$

[00092] As mentioned, according to various embodiments, another operation that the network modeling system may perform on packets includes the combination of two HRegions. Two HRegions may be combined only if they differ in values of at most one field. For example, consider the following two regions:

10
$$r1 = \{ (f1, \{ \text{"0000"} \}), (f3, \{ \text{"0000"} \}) \}$$

$$r2 = \{ (f1, \{ \text{"0000"} \}), (f3, \{ \text{"0001"}, \text{"1000"} \}) \}$$

[00093] These two regions only differ in values of the field f3 and may be combined to the following single region:

$$\text{combined} = \{ (f1, \{ \text{"0000"} \}), (f3, \{ \text{"000x"}, \text{"1000"} \}) \}$$

15 **[00094]** Combining the two HRegions may be performed using the technique described for combining two HSeqCollections to combine the HSeqCollection values corresponding to the single field in which the two HRegions differ, and then copying over the values corresponding to the other fields.

[00095] According to another embodiment, subtracting an HRegion from an HExpr may also be an operation performed on a packet collection by the network modeling system. First, the provided HRegion is intersected with the minuend of the HExpr. The resulting HRegion may be the actual region to be subtracted from the provided HExpr. This region may be added as a new subtrahend. However, in subtrahend simplification operations with HExprs, the HRegion may first be combined with other
 20 subtrahends before adding it as a new subtrahend. If after the combination, any of the
 25

subtrahends becomes equal to the minuend, then the result HExpr may be empty. As an illustrative example, consider the following HExpr:

Minuend = { (f1, {"000x"}) }

Subtrahend 1 = { (f1, {"0000"}) }

- 5 **[00096]** Subtracting { (f1, {"0010"}) } the HRegion from the HExpr results in the following (i.e., the HRegion is simply added as a new subtrahend):

Minuend = { (f1, {"000x"}) }

Subtrahend 1 = { (f1, {"0000"}) }

Subtrahend 2 = { (f1, {"0010"}) }

- 10 **[00097]** However, subtracting { (f1, {"0001"}) } from the HExpr, results in an empty HExpr, as illustrated below:

Minuend = { (f1, {"000x"}) }

Subtrahend 1 = { (f1, {"0000"}) }

Subtrahend 2 = { (f1, {"0001"}) }

- 15 **[00098]** The two subtrahends may be combined to form:

Minuend = { (f1, {"000x"}) }

Subtrahend 1 = { (f1, {"000x"}) }

- [00099]** Subsequently in this case, there is an HExpr where one of the subtrahends is equal to minuend. Hence, all possible packets are excluded from the minuend and
 20 nothing is left.

[000100] In rewriting a field in an HRegion, the value of the field may be simply changed to a new value. For example, rewriting the value of the field f1 to the HSeqCollection {"0000"} in the HRegion

{ (f1, {"0001", "0010"}), (f3, {"0000"}) }

[000101] results in the HRegion

$$\{ (f1, \{ "0000" \}), (f3, \{ "0000" \}) \}.$$

[000102] A field in an HExpr may also be rewritten in another operation that the network modeling system may perform on a packet. For rewriting field 'f' in an HExpr, first, the minuend of the HExpr is rewritten using the technique to rewrite a field of an HRegion. This minuend rewrite step may apply to various embodiments, which vary in the number and form of subtrahends. For example, if the HExpr only contains a minuend, then the rewrite of the field is complete. Alternatively, there may be cases where subtrahends are present.

[000103] The first case is where the value of field f1, the field to be rewritten, is the same between the minuend and each subtrahend. In this case, the value may be rewritten for field f1 in each subtrahend. To illustrate, consider this HExpr:

$$\text{Minuend} = \{ (f1, \{ a, b \}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{ a, b \}), (f2, \{ z \}) \}$$

[000104] The minuend contains multiple values for field f1, a and b. An alternate representation would be to generate one HExpr for each value of field f1 and union them together, as illustrated below:

$$\text{Minuend} = \{ (f1, \{ a \}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{ a, b \}), (f2, \{ z \}) \}$$

U

$$\text{Minuend} = \{ (f1, \{ b \}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{ a, b \}), (f2, \{ z \}) \}$$

[000105] In each HExpr, there is a subtrahend value, which has no intersection with the minuend. In these cases, that non-intersecting subtrahend value may be removed.

Hence, the expression may be reduced to the following:

Minuend = { (f1, {a}) }
 Subtrahend 1 = { (f1, {a}), (f2, {z}) }
 U

Minuend = { (f1, {b}) }
 5 Subtrahend 1 = { (f1, {b}), (f2, {z}) }

[000106] The result of rewriting field f1 to value d may be illustrated by:

Minuend = { (f1, {d}) }
 Subtrahend 1 = { (f1, {d}), (f2, {z}) }
 U

10 Minuend = { (f1, {d}) }
 Subtrahend 1 = { (f1, {d}), (f2, {z}) }

[000107] This result may be simplified to a single HExpr, since both parts of the union are identical, as illustrated below:

Minuend = { (f1, {d}) }
 15 Subtrahend 1 = { (f1, {d}), (f2, {z}) }

[000108] As illustrated above, in this example, each field value of f1 may be replaced with d, even when there are multiple field values. However, in examples where the field values differ, handling these cases requires more care.

[000109] In another embodiment, consider another example where the values for
 20 field f1 (to be rewritten) are different between the minuend and one subtrahend. To illustrate, consider this HExpr example:

Minuend = { (f1, {a, b}) }
 Subtrahend 1 = { (f1, {a}), (f2, {z}) }

[000110] The alternate representation for this HExpr may be illustrated by:

Minuend = { (f1, {a}) }
 Subtrahend 1 = { (f1, {a}), (f2, {z}) }

U

Minuend = { (f1, {b}) }
 5 Subtrahend 1 = { (f1, {a}), (f2, {z}) }

[000111] For the second HExpr, subtrahend 1 may be removed, because it has no intersection with the minuend, yielding the following:

Minuend = { (f1, {a}) }
 Subtrahend 1 = { (f1, {a}), (f2, {z}) }

10 U

Minuend = { (f1, {b}) }

[000112] The result of rewriting field f1 to d becomes:

Minuend = { (f1, {d}) }
 Subtrahend 1 = { (f1, {d}), (f2, {z}) }

15 U

Minuend = { (f1, {d}) }

[000113] In this case, the second HExpr subsumes the first, yielding:

Minuend = { (f1, {d}) }

[000114] In this example, if the value of field f1 (i.e., the field to be rewritten) on a
 20 single subtrahend is different from the value of that field on the minuend, then the
 subtrahend can be discarded when performing the rewrite.

[000115] In a third case, there may be more complexity where an HExpr contains
 two or more subtrahends, and each subtrahend has a value for field f1 that is different
 from the minuend. The challenge is that each subtrahend removes a piece of the minuend,

but the intersection of the non-removed pieces must be accounted for. To illustrate, consider this HExpr:

$$\text{Minuend} = \{ (f1, \{a, b\}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{a\}), (f2, \{z\}) \}$$

$$5 \quad \text{Subtrahend 2} = \{ (f1, \{b\}), (f3, \{w\}) \}$$

[000116] Putting this into the alternate union form yields the following:

$$\text{Minuend} = \{ (f1, \{a\}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{a\}), (f2, \{z\}) \}$$

$$\text{Subtrahend 2} = \{ (f1, \{b\}), (f3, \{w\}) \}$$

10 U

$$\text{Minuend} = \{ (f1, \{b\}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{a\}), (f2, \{z\}) \}$$

$$\text{Subtrahend 2} = \{ (f1, \{b\}), (f3, \{w\}) \}$$

[000117] Where subtrahends have no intersection with the minuend value for field f1, they can be removed. Each HExpr contains one such removable subtrahend. The result is the following:

15

$$\text{Minuend} = \{ (f1, \{a\}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{a\}), (f2, \{z\}) \}$$

U

20

$$\text{Minuend} = \{ (f1, \{b\}) \}$$

$$\text{Subtrahend 2} = \{ (f1, \{b\}), (f3, \{w\}) \}$$

[000118] Now the value of field f1 may be rewritten with d:

$$\text{Minuend} = \{ (f1, \{d\}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{d\}), (f2, \{z\}) \}$$

25

U

Minuend = { (f1, {d}) }

Subtrahend 2 = { (f1, {d}), (f3, {w}) }

[000119] However, these two remaining subtrahends may not be simply directly combined into a single HExpr, even though their minuends are the same, because there is an intersecting piece that remains. As such, the correct single HExpr form of the union-form HExpr above is:

Minuend = { (f1, {d}) }

Subtrahend 1 = { (f1, {d}), (f2, {z}), (f3, {w}) }

[000120] This new subtrahend removes the piece of the minuend, which is shared between the two subtrahends in the union-form HExpr above. According to various embodiments, the network modeling system may automatically compute this singular resulting subtrahend, for HExprs with any number of subtrahends. Each subtrahend that has a different value for field f1 than the minuend may serve as an input in the computation.

[000121] For each pair of such subtrahends, a new subtrahend may be generated. In this new subtrahend, the value for field 'f1' is the combination of the values in the two original subtrahends, where the combination is computed by the technique describe for combining HSeqCollections. The value for each other field is computed from the intersection of the corresponding field values in the two original subtrahends. If the intersection of each non-'f1' field is non-empty, then this subtrahend is to be added, and moves to the next step. However, if the intersection is empty, then no subtrahend is added to the pending HExpr.

[000122] According to various embodiments, with a list of "intermediate" subtrahends, the process may be repeated. For example, merging pairs of remaining subtrahends, until either the newly merged pair has the same value for field f1 as the

minuend, or the list of subtrahends cannot be further reduced. To illustrate, consider the following example HExpr:

Minuend = { (f1, {"0001", "0010", "0100"}) }

Subtrahend 1 = { (f1, {"0001"}), (f2, {"0001", "0010"}) }

5 Subtrahend 2 = { (f1, {"0001", "0010", "0100"}), (f2, {"1000"}) }

Subtrahend 3 = { (f1, {"0100"}), (f2, {"0001", "0100"}) }

Subtrahend 4 = { (f1, {"0010"}), (f2, {"0001", "0010", "0100"}) }

[000123] The value of 'f1' may be rewritten to "1000". First, the minuend may be written as:

10 New Minuend = { (f1, {"1000"}) }

[000124] From the collection of subtrahends, only subtrahend 2 has the same value as the original minuend on 'f1'. Subsequently, that subtrahend may be added after rewriting its 'f1' field. Three subtrahends where the value of field 'f1' differs from the minuend remain. New subtrahends may be added based on each possible pair of values in this set of remaining subtrahends. For each new potential subtrahend, the 'f1' field value comes from the combination of the two input 'f1' values and the value for each non-'f1' field comes from the intersection of the two input non-'f1' values. To illustrate, consider the following example:

Subtrahends (1, 3) = { (f1, {"0001", "0100"}), (f2, {"0001"}) }

20 Subtrahends (1, 4) = { (f1, {"0001", "0010"}), (f2, {"0001", "0010"}) }

Subtrahends (3, 4) = { (f1, {"0010", "0100"}), (f2, {"0001", "0100"}) }

[000125] None of these new potential subtrahends has the same value as the original minuend for f1. However, they can still be further merged. Merging subtrahend pairs (1, 3) and (1, 4) yields the following:

Subtrahends (1, 3) + (1, 4) = { (f1, {"0001", "0010", "0100"}), (f2, {"0001"}) }

Subtrahends (3, 4) = { (f1, {"0010", "0100"}), (f2, {"0001", "0100"}) }

[000126] Merging these two subtrahends yields a single subtrahend below:

Subtrahends (1, 3) + (1, 4) + (3, 4) = { (f1, {"0001", "0010", "0100"}), (f2, {"0001"}) }

5 **[000127]** This simplified subtrahend can now be added. The resulting HExpr contains the following two subtrahends:

Subtrahend 1 = { (f1, {"0001", "0010", "0100"}), (f2, {"1000"}) }

Subtrahend 2 = { (f1, {"0001", "0010", "0100"}), (f2, {"0001"}) }

[000128] These two subtrahends can be combined using the technique described in
10 subtrahend simplification operations with HExprs, yielding a single subtrahend:

Subtrahend = { (f1, {"0001", "0010", "0100"}), (f2, {"0001", "1000"}) }

[000129] Then field f1 may be rewritten, yielding a "final" HExpr of:

Minuend = { (f1, {"1000"}) }

Subtrahend = { (f1, {"1000"}), (f2, {"0001", "1000"}) }

15 **[000130]** Another action that may be performed on packets according to various embodiments may include moving the value of a field to another field in an HRegion. Given a pair of source and target fields, this action moves the value of the source field to the target field. This operation is equivalent to rewriting the source field with the fully wildcarded HSeqCollection and rewriting the target field with the old value of the source
20 field. To illustrate, consider the following HRegion:

{ (f1, {"0001", "0010"}), (f3, {"0000"}) }.

[000131] Moving the value of f1 to f2 results in the following:

{ (f2, {"0001", "0010"}), (f3, {"0000"}) }.

[000132] Moving the value of f1 to f3 results in the following:

$$\{ (f3, \{ "0001", "0010" \}) \}.$$

[000133] Another action that may be performed on packets according to various embodiments may include moving the value of a field to another field in an HExpr. To

5 move value of field 's' to field 't' in HExpr 'hs', first value of field 's' may be moved to field 't' in the minuend of 'hs' using the technique described for moving the value of a field to another field in an HRegion. The resulting HRegion is the minuend of the result HExpr. This shift may be similarly performed on each of the subtrahends of 'hs' to compute the subtrahends of the result HExpr. However, some of the subtrahends should
10 be discarded as part of this operation. To illustrate, consider the following HExpr:

$$\text{Minuend} = \{ (f2, \{ "10xx" \}) \}$$

$$\text{Subtrahend 1} = \{ (f1, \{ "0001" \}) \}$$

$$\text{Subtrahend 2} = \{ (f2, \{ "0000" \}) \}$$

[000134] When moving the value of field f2 to field f1, the first subtrahend should
15 be discarded, because it has been subtracting a region from minuend, which is not being rewritten to a new value. The second subtrahend may not have an entry for the target field (f1) and should be retained in the final result. We just need to move its value. This procedure results in the following HExpr:

$$\text{Minuend} = \{ (f1, \{ "10xx" \}) \}$$

20 **Subtrahend** = $\{ (f1, \{ "0000" \}) \}$

[000135] The representation may be implemented in Java, C++, C#, or any suitable programming language, on any computing hardware, such as a general-purpose processor or graphics processor. In some embodiments, the invention may be implemented directly in hardware, via field-programmable gate array or application-specific integrated circuit.

[000136] 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. Additionally, the network modeling system may be enabled to detect rules within the same or different tables related to a given input rule such that
5 packets matching them will also be matched by the given rule. Rules may also be detected within the same or different tables that are ambiguous due to having the same matching priority, and some or full overlap in match fields.

[000137] Embodiments of the present invention provide a system and method to analyze rules and the network more generally, and provide suggestions for optimization,
10 such as the elimination of rules, modification to rules, or addition of other rules. Furthermore, the networking modeling system according to various embodiment may 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
15 problems.

[000138] FIG. 8 is a high-level block diagram 800 of a computer system, in accordance with an embodiment of the present invention. As shown in FIG. 8, a computer system can include hardware elements connected via a bus, including a network interface 802, that enables the computer system to connect to other computer
20 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 804, 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
25 processors.

[000139] In some embodiments, the computer system can include a graphical user interface (GUI) 806. GUI 806 can connect to a display (LED, LCD, tablet, touch screen,

or other display) to output user viewable data. In some embodiments, GUI 806 can be configured to receive instructions (e.g., through a touch screen or other interactive interface). In some embodiments, I/O interface 808 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 808 can include a wired or wireless interface for connecting to infrared, Bluetooth, or other wireless devices.

[000140] In some embodiments, the computer system may include local or remote data stores 810. Data stores 810 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 810 can include data generated, stored, or otherwise utilized as described herein. For example, data stores 810 can include computed flows 812 and network models 814, generated and stored as described above. Memory 816 can include various memory technologies, including RAM, ROM, EEPROM, flash memory or other memory technology. Memory 816 can include executable code to implement methods as described herein. For example, memory 816 can include a network analyzer module 818 and report generator module 820 that each implement methods described herein.

[000141] 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 above-described invention techniques. The disclosed examples are illustrative and not restrictive.

CLAIMS

What is claimed is:

1. A computer-implemented method for modeling networks, comprising:
 - obtaining network behavior information associated with a plurality of network devices in a network, the network behavior information including at least one rule table for each network device of the plurality of network devices, the at least one rule table including at least a list of rules defining forwarding behavior for each network device in the plurality of network devices within the network, at least one rule in the list of rules including a match specifier and one or more corresponding actions, the match specifier identifying matching packets to which the at least one rule applies, the match specifier and the one or more corresponding actions causing a traffic flow of the matching packets;
 - identifying one or more packets to analyze in a network model, each packet of the one or more packets including packet header values;
 - creating a virtual packet header representing the one or more packets, the virtual packet header including one or more packet header groups, at least one of the one or more packet header groups representing a union of two or more packet header values;
 - pushing the virtual packet header through the network model, the network model specifying the forwarding behavior of the plurality of network devices in the network;
 - applying, to the virtual packet header, a first rule table of at least one network device of the plurality of network devices in the network by comparing the virtual packet header with a first match specifier of a first rule in a first list of rules of the first rule table of the at least one network device in the network; and
 - determining one or more flow paths defining the traffic flow based on the one or more corresponding actions of the at least one rule of the at least one rule table of the at least one network device based on the virtual packet header matching the match specifier of the at least one rule.

2. The computer-implemented method of claim 1, wherein at least one of the one or more packet header groups of the virtual packet header includes at least one packet header value that is unspecified.

3. The computer-implemented method of claim 1, wherein at least one packet header group in the virtual packet header represents a group of two or more packet header values with one or more wildcard bit strings, wherein at least one wildcard bit represents both a 0 and a 1 in a particular bit position in the two or more packet header values.

4. The computer-implemented method of claim 1, wherein the virtual packet header is represented as a difference between a first virtual packet header and one or more other virtual packet headers, the first virtual packet header representing a first group of packets, and the one or more other virtual packet headers representing additional groups of packets.

5. The computer-implemented method of claim 1, wherein determining the one or more flow paths further comprises performing a transformation operation on the virtual packet header, wherein the transformation operation changes at least one packet header value.

6. The computer-implemented method of claim 5, wherein the transformation operation includes one or more of a rewrite operation, a combination operation, a move operation, a shift operation, an intersection operation, and a reduction operation.

7. The computer-implemented method of claim 4, further comprising:
modifying at least one of the first virtual packet header and one of the other virtual packet headers computed as a result of a transformation operation, wherein the transformation operation changes at least one packet header value; and
reducing the result of the transformation operation, wherein the result of the operation generates a different virtual packet header representing an equivalent group of packets.

8. The computer-implemented method of claim 1, wherein the network device behavior is inferred based at least in part on control-plane configuration data.

9. The computer-implemented method of claim 1, wherein the one or more corresponding actions for matching packets includes at least one of: applying transformations to packets, dropping packets, forwarding packets to a subsequent rule, forwarding packets to a subsequent rule table, forwarding packets to a subsequent network device in the network, forwarding packets out of one or more ports, and forwarding packets to another device out of the network.

10. A non-transitory computer readable medium storing code executable by a processor to implement a method, the method comprising:

obtaining network behavior information associated with a plurality of devices in a network, the network behavior information defining forwarding behavior of each device in the network;

creating a virtual packet header representing one or more collections of packets, the virtual packet header including one or more packet header groups, at least one of the one or more packet header groups representing a union of two or more packet header values;

performing analysis operations using the virtual packet header and based at least in part on the network behavior information; and

determining one or more flow paths defining a traffic flow based at least in part on the analysis operations.

11. The non-transitory computer readable medium of claim 10, wherein the network behavior information includes at least one rule table, the at least one rule table including at least one rule defining a flow path of the packet or the collections of packets within the traffic flow of the network, and wherein at least one rule in the at least one rule table is implicit or explicit, and each of the at least one rule tables is implicit or explicit, and at least one rule table contains an ordered or unordered list of rules.

12. The non-transitory computer readable medium of claim 10, wherein the virtual packet header is represented as a difference between a first virtual packet header and one

or more other virtual packet headers, the first virtual packet header representing a first group of packets, and the one or more other virtual packet headers representing additional groups of packets.

13. The non-transitory computer readable medium of claim 10, wherein the analysis operations include a transformation operation on the virtual packet header, wherein the transformation operation changes at least one packet header value of the one or more collections of packets.

14. The non-transitory computer readable medium of claim 10, wherein at least one of the one or more packet header groups of the virtual packet header includes at least one packet header value that is unspecified.

15. The non-transitory computer readable medium of claim 10, wherein at least one packet header group in the virtual packet header represents a group of two or more packet header values with one or more wildcard bit strings, wherein at least one wildcard bit represents both a 0 and a 1 in a particular bit position in the two or more packet header values.

16. 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:

obtain network behavior information associated with a plurality of devices in a network, the network behavior information defining forwarding behavior of each device in the network;

create a virtual packet header representing one or more collections of packets, the virtual packet header including one or more packet header groups, at least one of the packet header groups representing a union of two or more packet header values of the one or more collections of packets;

perform analysis operations using the virtual packet header and based at least in part on the network behavior information; and

determine one or more flow paths defining a traffic flow based at least in part on the analysis operations.

17. The system of claim 16, wherein the virtual packet header is represented as a difference between a first virtual packet header and one or more other virtual packet headers, the first virtual packet header representing a first group of packets, and the one or more other virtual packet headers representing additional groups of packets.

18. The system of claim 16, wherein the analysis operations include a transformation operation on the virtual packet header, wherein the transformation operation changes at least one packet header value of the one or more collections of packets.

19. The system of claim 16, wherein at least one of the one or more packet header groups of the virtual packet header includes at least one packet header value that is unspecified.

20. The system of claim 16, wherein at least one packet header group in the virtual packet header represents a group of two or more packet header values with one or more wildcard bit strings, wherein at least one wildcard bit represents both a 0 and a 1 in a particular bit position in the two or more packet header values.




FIG. 1




FIG. 2

| 300 | | 320 | | 340 | |
|--------------|-------------|----------|------------|--------------|--------------|
| exact | wildcarded | exact | wildcarded | 2-bit values | 3-bit values |
| 302 00000000 | xxxxxxx 310 | 322 1010 | 101x 328 | 342 01 | 001 346 |
| 304 00000001 | | 324 1011 | 1100 330 | 344 10 | 010 348 |
| 306 00000010 | | 326 1100 | | | 100 350 |
| ... | | | | | |
| 308 11111111 | | | | | |

FIG. 3(A)

FIG. 3(B)

FIG. 3(C)




FIG. 4




FIG. 5(A)

FIG. 5(B)

FIG. 5(C)




FIG. 6(A)




FIG. 6(B)

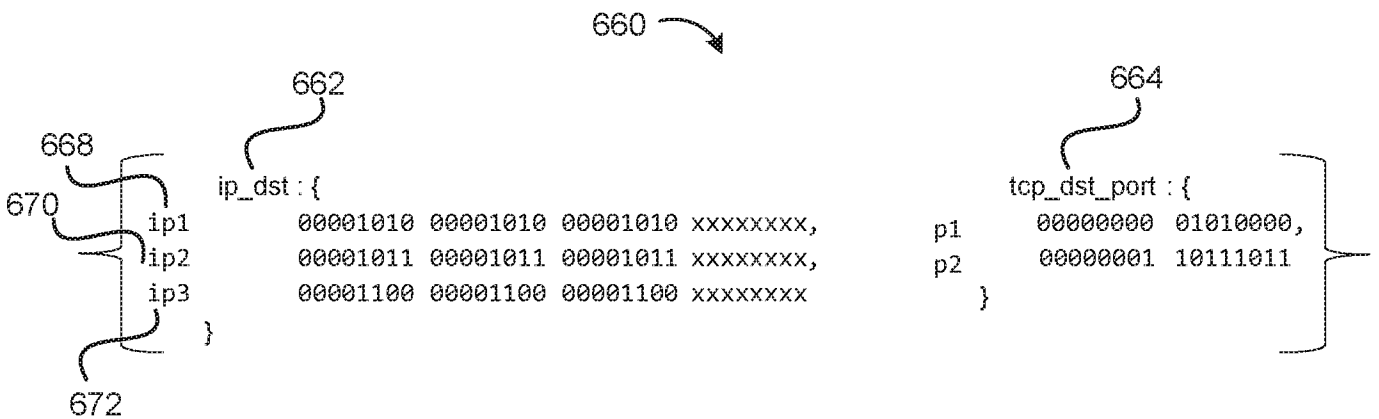


FIG. 6(C)

700

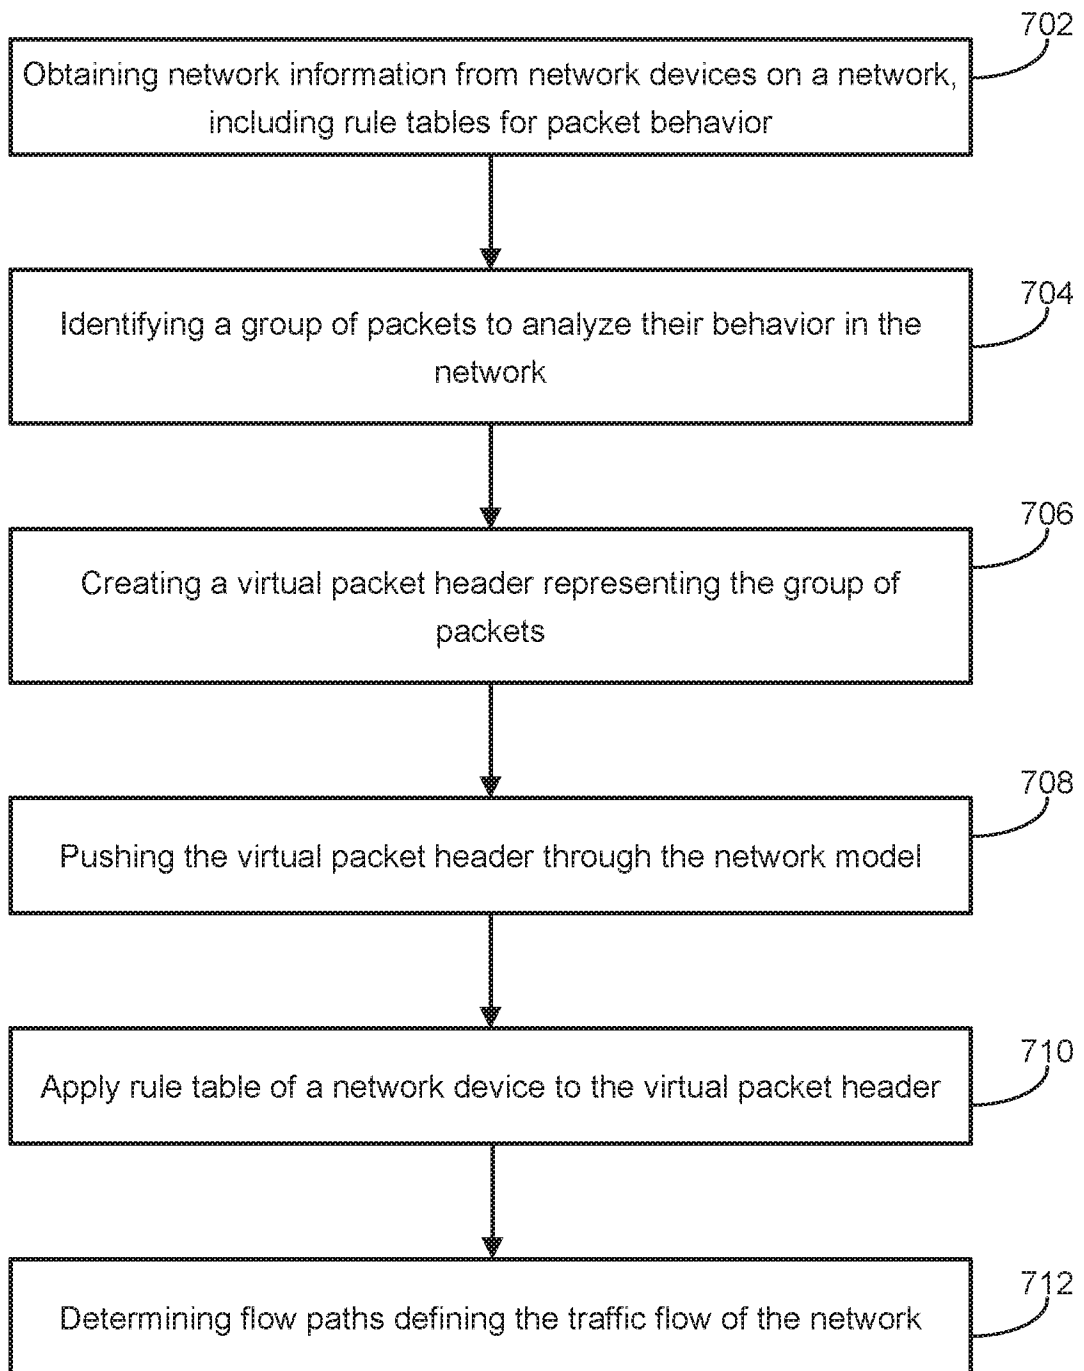


FIG. 7




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 17/57751

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 17/10 (2017.01)
CPC - G06F 17/5018; G06F 17/5036; G06F 17/5009; G05B17/02; G06F 2217/16; G06F 17/5022; H04L 41/145; G06F 11/261; G06F 11/3457; H05K 999/99; G06F 11/3419; G06F 11/3409; G06F 11/3447; G06F 11/3495

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)

See Search History Document

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

See Search History Document

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

See Search History Document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | US 2015/0092778 A1 (Nicira, Inc.), 02 April 2015 (02.04.2015), entire document, especially Abstract; para [0005], [0089], [0119], [0123], [0127], [0172], [0199]-[0200], [0205], [0233]-[0234], [0207], [0287], [0396] | 1-20 |

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

21 December 2017 (21.12.2017)

Date of mailing of the international search report

18 JAN 2018

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

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