

# Practical Network-wide Packet Behavior Identification by AP Classifier

Huazhe Wang  
University of Kentucky  
huazhe.wang@uky.edu

Chen Qian  
University of Kentucky  
qian@cs.uky.edu

Ye Yu  
University of Kentucky  
ye.yu@uky.edu

Hongkun Yang  
University of Texas, Austin  
yanghk@cs.utexas.edu

Simon S. Lam  
University of Texas, Austin  
lam@cs.utexas.edu

## ABSTRACT

Identifying the network-wide forwarding behaviors of a packet is essential for many network management applications, including rule verification, policy enforcement, attack detection, traffic engineering, and fault localization. Current tools that can perform packet behavior identification either incur large time and memory costs or do not support real-time updates. In this paper we present AP Classifier, a control plane tool for packet behavior identification. AP Classifier is developed based on the concept of atomic predicates which can be used to characterize the forwarding behaviors of packets. Experiments using the data plane network state of two real networks show that the processing speed of AP Classifier is faster than existing tools by at least an order of magnitude. Furthermore, AP Classifier uses very small memory and is able to support real-time updates.

## CCS Concepts

• **Networks** → **Packet classification; Network management; Network reliability;**

## Keywords

Packet behavior identification; Network verification; Software defined networking

## 1. INTRODUCTION

Managing packet forwarding in a large network is a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [Permissions@acm.org](mailto:Permissions@acm.org).

CoNEXT '15 December 01-04, 2015, Heidelberg, Germany

© 2015 ACM. ISBN 978-1-4503-3412-9/15/12...\$15.00

DOI: <http://dx.doi.org/10.1145/2716281.2836095>

complex problem. Software defined networking (SDN) simplifies network management by decoupling the control plane from devices that forward packets, to be referred to as *boxes*.<sup>1</sup> More specifically, control plane applications, including routing [5, 9], traffic engineering [4], access control [26], measurement [32], and policy enforcement [27], are implemented as software in a logically centralized controller. The controller specifies forwarding actions of packets by writing directly into flow tables in each box in the form of rules, through a standard API such as OpenFlow [25].

Let a *flow* be an equivalence class of packets defined on a subset of fields in the packet header, e.g., the 5-tuple consisting of source address, destination address, source port, destination port, and protocol type. All packets of a flow have the same forwarding behaviors in a network (also referred to as the flow's behaviors) when there is no data plane update. Network-wide *packet behavior identification* is a control plane function that discovers the actual forwarding behaviors of the packets in a flow (or a set of flows) including their forwarding paths, where they stop or are dropped, and which boxes they traverse, by analyzing network state in the data plane [13]. Packet behavior identification is necessary for SDN management in the following situations.

**Verification of flow properties.** For network flows, the control plane may specify pre-defined flow behaviors that satisfy application requirements or network policies, called *flow properties*. We highlight several typical flow properties.

- **Forwarding correctness:** The control plane must ensure that flow packets can be forwarded to the destination (e.g., a host or an egress router), or dropped if they are not allowed to reach the destination.

<sup>1</sup>We use “box” to refer to any network device that forwards packets, including routers, switches, and functional middle boxes such as firewalls, NATs and intrusion detection systems (IDSes).

- **Policy enforcement:** Network policies may require packets of the flow to go through various middle boxes. For example, HTTP traffic should be forwarded through a sequence of middle boxes: firewall, IDS, and web proxy [27]. Other types of traffic may be required to traverse different sequences of middle boxes.
- **Quality of service:** Some applications require guaranteed flow quality. For example, a multi-tenant cloud should provide certain levels of bandwidth or latency for its users based on service level agreements [6, 21].
- **VLAN isolation:** A cloud provider guarantees that packets in a virtual network (VLAN) cannot travel to another VLAN.

Any data plane update could change the behaviors of a number of flows. Prior to data plane updates, the controller needs to verify that the data plane, with the new updates, can forward the packets correctly and comply with the flow properties. Such verification requires packet behavior identification for the flows that will be affected by the new rules.

**Attack detection.** Data plane attacks to an SDN may change the correct packet behaviors or send packets with abnormal behaviors, such as data plane DDoS attacks. An efficient data plane attack detection method should verify data plane forwarding behaviors and be aware of the behaviors that violate network policies. For example, a recent work SPHINX [11] uses flow graphs to represent actual network operations and detect abnormal changes.

**Traffic engineering.** Centralized traffic engineering [5] [8] [22] determines the forwarding paths for flows to maximize network throughput. When the controller is notified about a new flow, it needs to identify its packet behaviors in the current data plane and check whether they can meet application requirements such as bandwidth or latency. If not, the controller needs to modify the data plane to install a desired forward path for the flow.

**Localization of network faults.** When the control plane finds a flow property violation at any time, it should identify the actual flow behaviors in the network and compare them with the expected behaviors. In this way the controller can find the part of the data plane that contains faults, called fault localization [33].

A practical packet behavior identification method must satisfy three requirements. First, it provides a high throughput in responding to packet behavior queries. According to recent measurement results [15] [7], a large data center network may see hundreds of thousands of new flows per second. SDNs should support hundreds of data plane updates per second [14] and each update may need to query multiple flows to verify correctness.

Hence a desired throughput should exceed one million packet queries per second (1 Mqps). Second, the query structure should fit into a small and fast memory such as cache. Third, the query structure can be updated in real time under data plane changes to ensure that query results reflect the current network state.

Unfortunately none of the existing solutions can meet all of the requirements stated above. A straightforward approach is to maintain copies of flow tables of all boxes in the controller. However even for a medium-scale network used in [17], tens of GBs are required to store all rules [13]. Due to slow search speed among flow tables and disk I/Os, the query throughput is very low. Very recently, Inoue *et al.* [13] propose to use a multi-valued decision diagram (MDD) to classify flows to different sets of network-wide behaviors. However, an MDD cannot be updated in real time. It has to be reconstructed for updates and the reconstruction time needs 0.1 second to a few seconds.<sup>2</sup> Given the network update frequency of current networks, many updates could happen during the reconstruction time. Hence the controller may never get an MDD that correctly reflects the *current* network state.

In this work, we propose a network-wide packet behavior identification method called *AP Classifier*, where AP stands for Atomic Predicates, a concept developed in [30]. Each atomic predicate specifies a set of packets that have the same forwarding behavior in the network. We develop a novel data structure, called AP Tree, which can rapidly classify a packet to an atomic predicate. The packet behavior can then be easily computed using the atomic predicate. We evaluated the performance of AP Classifier using the data plane network state, including forwarding tables and access control lists (ACLs), from two real networks, namely: Internet2 [2] and a Stanford campus network [17]. Our results show that AP Classifier, running on a general-purpose desktop computer, only uses a few MBs memory and supports more than two millions of queries per second. In addition it can be updated in real time (< 4 ms for 95% updates in Internet2 and < 1 ms for 95% updates in Stanford).

The balance of this paper is organized as follows. Section 2 presents related work. We discuss the network model and background knowledge in Section 3. We introduce the framework of AP Classifier in Section 4. The algorithms to construct an AP Tree are presented in Section 5 and the update and reconstruction methods of an AP Tree for dynamic networks are presented in

<sup>2</sup>The paper [13] claims that if a data plane update does not change the existing packet behaviors, MDD update can be finished in tens of milliseconds. However from examining update traces of the Route Views Project [3], it is unlikely that a data plane update does not change the existing packet behaviors.

Section 6. We present experimental results in Section 7 and a discussion in Section 8. Finally we conclude this work in Section 9.

## 2. RELATED WORK

Network-wide packet behavior identification is equivalent to reachability computation for a specific packet. This problem is related to, but different from, network reachability analysis which has been studied for over a decade. Xie *et al.* [28] present a model for static reachability analysis of data plane network state. Quarnet [18] represents ACLs as firewall decision diagrams to compute network reachability. Ant eater [24] and Header Space Analysis (HSA) [17] are both general-purpose tools to check network invariants but not in real time.

For real-time applications, NetPlumber [16] makes use of HSA to detect network invariant violations. Veriflow [19] stores all data plane rules in a multi-dimensional prefix tree (trie) and determines the Equivalence Classes (ECs) of packets. An EC is defined to be a set of packets that have identical forwarding actions in all boxes. Veriflow then checks network invariants by analyzing reachability graphs of ECs. AP Verifier [30] uses predicates to represent packet filters that guard input and output ports of boxes. It enables very fast computation of reachability trees by first computing atomic predicates for the set of predicates. Instead of computing the conjunction of two predicates during reachability computation, AP Verifier computes the intersection of two sets of integers representing atomic predicates. AP Verifier is more time and space efficient than prior tools.

All of the above methods focus on analyzing network-wide invariants (e.g., reachability, loop-freedom) but were not designed to identify the reachability of a specific packet. For example, they can determine whether it is possible to reach box  $B$  from box  $A$  but cannot tell whether a given packet can reach  $B$ . AP Verifier [30] can check whether all packets entering a port in the network pass through a waypoint (e.g., a firewall) but cannot tell whether a specific packet traverses a given waypoint.

One possible solution to the packet behavior identification problem is to obtain all related data plane rules of the packet by searching the trie created in Veriflow and then compute the forwarding path based on the rules. However storing all rules requires non-trivial memory cost (tens of GBs for the Stanford network) which could cause disk I/Os during query processing. As a result, using the Veriflow trie for packet behavior identification was shown to be very slow by Inoue *et al.* [13] who proposed a tool that can quickly classify a packet to an EC. Its main drawback is that their MDD structure cannot correctly represent the current network state because it does not support real-time updates, especially for SDNs where data plane updates are frequent [20].

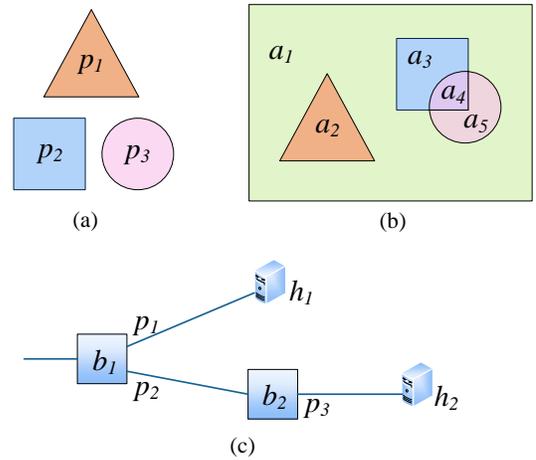


Figure 1: (a) Three predicates. (b) The packet header space and five atomic predicates. (c) A sample network including the three predicates.

Recently, Network Optimized Datalog is proposed as a general specification language to model high-level abstraction of network beliefs and dynamism [23]. A new approach to derive data plane from network configurations is presented in [12].

## 3. MODEL AND BACKGROUND

We model a network as a directed graph of boxes, each of which has a forwarding table as well as input and output ports guarded by access control lists (ACLs). Each packet has a fixed-size header, including all fields that are evaluated by forwarding tables and ACLs in the network. A flow is then a sequence of packets that have the same values in the evaluated header fields.

Following the concepts in [30], forwarding tables and ACLs are all packet filters. Each ACL can be specified by a *predicate*. The set of packets that are allowed by the ACL are evaluated to true by the predicate. Similarly, by analyzing a forwarding table, each output port can be specified by a forwarding predicate. The set of packets that can be forwarded to the port are evaluated to true by the predicate.<sup>3</sup> Forwarding tables and ACLs can be converted to predicates using the algorithms in [30]. A predicate  $P$  specifies the set of packets for which  $P$  evaluates to true. Hence if a packet can travel through a sequence of packet filters, it is evaluated to true by the conjunction of predicates corresponding to the packet filters.

Given a set of predicates, we can compute a set of *atomic predicates*. Due to space limitation, we do not repeat the formal definition of atomic predicates, which can be found in [30]. A proved property of the set of atomic predicates is that they specify the minimum set of equivalence classes in the set of all packets. *The pack-*

<sup>3</sup>All predicates are represented by binary decision diagrams (BDDs) in our implementation of AP Classifier.

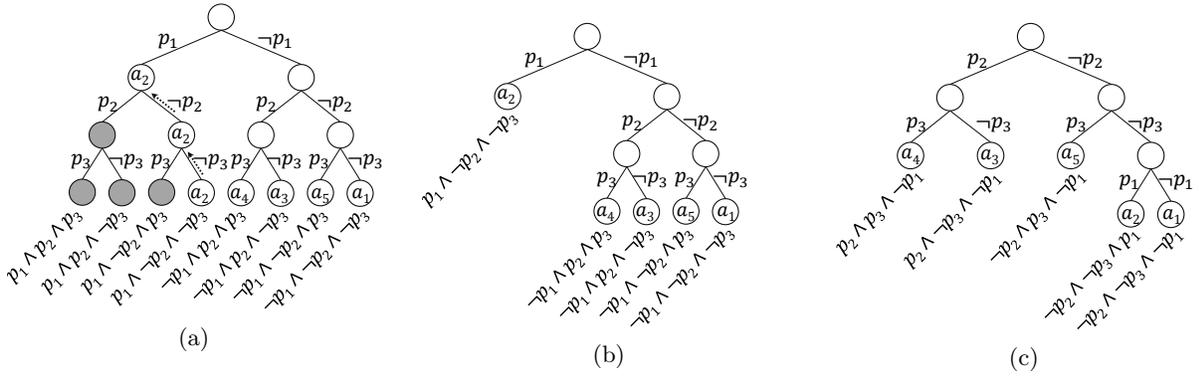


Figure 2: The AP Tree of predicates in Figure 1(b).

ets that are evaluated to true by the same atomic predicate have identical behaviors at all boxes. For a set of predicates  $P = \{p_1, p_2, \dots, p_k\}$ , each atomic predicate  $a_i$  is in the form  $a_i = q_1 \wedge q_2 \wedge \dots \wedge q_k$ , where  $q_j \in \{p_j, \neg p_j\}$ . (Note that  $a_i$  in the previous sentence is an atomic predicate only if it is not false.) Every predicate is equal to the disjunction of a subset of atomic predicates. Every packet is evaluated to true by one and only one atomic predicate.

As an illustration, Fig. 1(a) shows three predicates  $p_1$  (triangle),  $p_2$  (square), and  $p_3$  (circle), each of which represents a set of packets that are evaluated to true by a predicate. Each predicate specifies a set of packets that can pass the corresponding packet filter. Fig. 1(b) shows the three predicates in the packet header space. All packets in this example can be classified into five equivalence classes specified by five atomic predicates,  $a_1$  to  $a_5$ . Each predicate is equal to the disjunction of a subset of atomic predicates. For example,  $p_2 = a_3 \vee a_4$ . Also,  $a_4 = \neg p_1 \wedge p_2 \wedge p_3$ . All packets evaluated to true by  $a_4$  have identical behaviors: they can pass the filters of  $p_2$  and  $p_3$  but cannot pass  $p_1$ .

In the network shown in Fig. 1(c), Let  $p_1$  specify the set of packets that can be forwarded at box  $b_1$  to its output port to host  $h_1$ ,  $p_2$  specify the set of packets that can be forwarded at box  $b_1$  to its output port to box  $b_2$ , and  $p_3$  specify the set of packets that can be forwarded at box  $b_2$  to its output port to host  $h_2$ . A packet specified by  $a_4 = \neg p_1 \wedge p_2 \wedge p_3$  is forwarded at  $b_1$  by the path  $b_1 \rightarrow b_2 \rightarrow h_2$ . A packet specified by  $a_5 = \neg p_1 \wedge \neg p_2 \wedge p_3$  is forwarded to  $h_2$  if it is at  $b_2$ , but will be dropped if it is at  $b_1$ . An atomic predicate characterizes the behaviors of all packets it evaluates to true.

## 4. DESIGN FRAMEWORK

AP Classifier is a program designed for a SDN controller. It computes the network-wide behaviors for an input packet (or flow). AP Classifier performs two-stage processing for a packet. First, using the AP Tree, it classifies the packet to the atomic predicate that evaluates

to true for the packet. Second, AP Classifier determines all forwarding paths for the packet by using the atomic predicate, network information, and ingress box of the packet.

### 4.1 AP Tree

Using the algorithms presented in [30], the controller first converts each ACL to a predicate and the forwarding table of each box to  $m$  predicates, where  $m$  is the number of output ports of the box. Let  $P = \{p_1, p_2, \dots, p_k\}$  be the set of predicates of all boxes in the network. The controller constructs an AP Tree which is a binary tree. The root is labeled by  $p_1$ . At level  $i$ , the  $2^i$  internal nodes are each labeled by  $p_i$ . Starting from the root, at each internal node, the input packet is evaluated by the predicate in the label. If the result is true, the packet continues to be evaluated in the left sub-tree. Otherwise it goes to the right sub-tree. An AP Tree with  $(k + 1)$  levels can be constructed from evaluating each of the  $k$  predicates at each level of internal nodes. A leaf node is then labeled by  $q_1 \wedge q_2 \wedge \dots \wedge q_k$ ,  $q_i \in \{p_i, \neg p_i\}$ , which specifies the set of packets reaching the leaf. From Theorem 2 in [30], leaf labels (that are not false) represent the atomic predicates of  $P$ . Fig. 2(a) shows the AP Tree of the three predicates in Fig. 1(b). Shaded circles indicate leaf labels that are false. We will show that two sub-trees in an AP tree do not necessarily have a same predicate order in Section 5.3.

To classify a packet to an atomic predicate, AP Classifier simply searches the AP Tree by evaluating the packet until the leaf labeled by the atomic predicate is found. At each node, the packet is evaluated by checking the BDD of the predicate. Since predicates on sibling nodes are disjoint, for a given packet, the path from the root to the leaf is exclusive and determinate.

In the worst case, there could be  $2^k$  atomic predicates and finding a leaf needs to evaluate all  $k$  predicates. However, it is found that the number of atomic predicates is surprisingly small for real networks [30]. Hence many leaves specify empty sets of packets. For

example, in Fig. 2(a),  $p_1 \wedge p_2 \wedge p_3$ ,  $p_1 \wedge p_2 \wedge \neg p_3$ , and  $p_1 \wedge \neg p_2 \wedge p_3$  are all false according to the relationships in Fig. 1(b). Hence no packet can reach any of these three leaves. We use the following rule to “prune” the AP Tree: If no packet reaches a sub-tree, i.e., all leaves in the subtree are labeled by false predicates, the subtree is removed from the AP Tree. If an internal node has only one child, it is removed from the AP Tree as there is no need to check the predicate. We define the depth of a leaf to be the number of predicates evaluated to reach the leaf. After pruning, the average depth of all leaves in the AP Tree can be reduced and each node has either 0 or 2 children. Fig. 2(b) shows the pruned AP Tree whose average depth is  $(1 + 3 + 3 + 3 + 3)/5 = 2.6$ .

An important observation is the following: If predicates are placed at the levels in a different order, the average depth of the AP Tree may be different. In Fig. 2(c), the predicates are placed at three levels in the order of  $p_2, p_3, p_1$ . The average depth of all leaves in the pruned AP Tree is 2.4. *An important contribution of this work is an algorithm to find an order of predicates that substantially reduces the average depth of an AP Tree.*

For examples, each of the Internet2 and Stanford networks includes hundreds of thousands of forwarding rules, which can be converted to 161 (Internet2) or 507 (Stanford) predicates. Using our AP Tree construction algorithm, the average depth of the AP Tree is only 10.6 (Internet2) or 16.8 (Stanford). In an unpruned AP Tree, a packet needs to be evaluated by 161 or 507 predicates. AP Classifier only requires it to be evaluated by 10.6 or 16.8 predicates, on average, thus improving the query throughput by more than an order of magnitude. The detailed algorithm design of AP Tree construction is presented in Section 5.

## 4.2 Computing packet behaviors

The second stage of AP Classifier determines the network-wide behaviors of the queried packet from the network information, the ingress box, and the atomic predicate determined in the first stage.

Since the atomic predicate is in the form  $q_1 \wedge q_2 \wedge \dots \wedge q_k$ ,  $q_i \in \{p_i, \neg p_i\}$ , for any predicate  $p_j$ , AP Classifier can easily check whether the predicate evaluates to true or false for the packet. Recall that  $p_j$  represents a packet filter of an ACL or output port. Hence AP Classifier can determine at any box whether the packet is dropped and which port it is forwarded to. Starting from the ingress box, i.e., the box that sees the packet first in the network, AP Classifier finds the output port to which the packet is forwarded and then determines the next-hop box. If the packet is a multicast packet, it may be forwarded to multiple ports. AP Classifier continues to find the forwarding ports on the next-visited boxes until the packet reaches the destination or is dropped. The

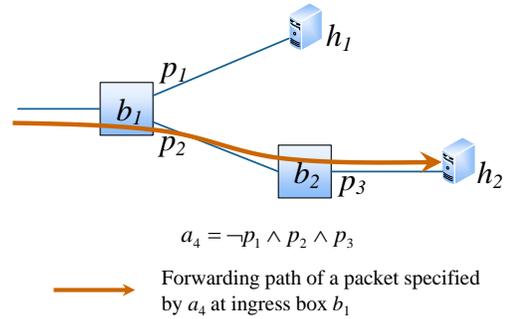


Figure 3: Computing forwarding path for a packet

packet behaviors are thus obtained.

Fig. 3 shows an example to illustrate how to compute network-wide forwarding paths for a given packet. Consider a packet which arrives at the ingress box  $b_1$  and it is classified to atomic predicate  $a_4$  by searching the AP Tree. The representation,  $\neg p_1 \wedge p_2 \wedge p_3$ , of  $a_4$  shows that the packet is forwarded to  $b_2$  because  $p_1$  is false and  $p_2$  is true for the packet. Similarly at  $b_2$ , the packet is forwarded to  $h_2$  because  $p_3$  is true for the packet.

We ran experiments to evaluate the speed of the above approach on a general-purpose desktop computer. We found that, for the Internet2 and Stanford datasets, the throughput is greater than 15M and 10M packets per second, respectively. Note that this throughput is much higher than the throughput in the first stage. Therefore, the main effort of this work is to optimize the construction, search, and update of the AP Tree.

## 5. AP TREE OPTIMIZATION

The most challenging problem in designing AP Classifier is to construct an AP Tree with minimized average depth, which can support dynamic updates.

### 5.1 Query throughput versus average depth

To reduce the query time and improve the query throughput, the optimization goal of AP Tree construction is to reduce the average depth of leaves. We conduct a set of experiments to justify the correlation of reducing the average depth and improving the throughput. We use the Internet2 network containing 161 predicates and the Stanford network containing 507 predicates. In each experiment, we randomly order the  $k$  predicates for placement at levels of the AP Tree. Then we query the generated tree using sample packets and measure the query throughput. In Fig. 4, we show the relationship between query throughput and average depth for 100 random generated AP Trees for each network. After pruning, the average depth of the AP Tree of Internet2 varies from 15.9 to 44.2, and the average depth of the AP Tree of Stanford varies from 39.1 to 92.5. From the two sub-figures in Fig. 4, it is obvious that an AP Tree with smaller average depth provides higher query

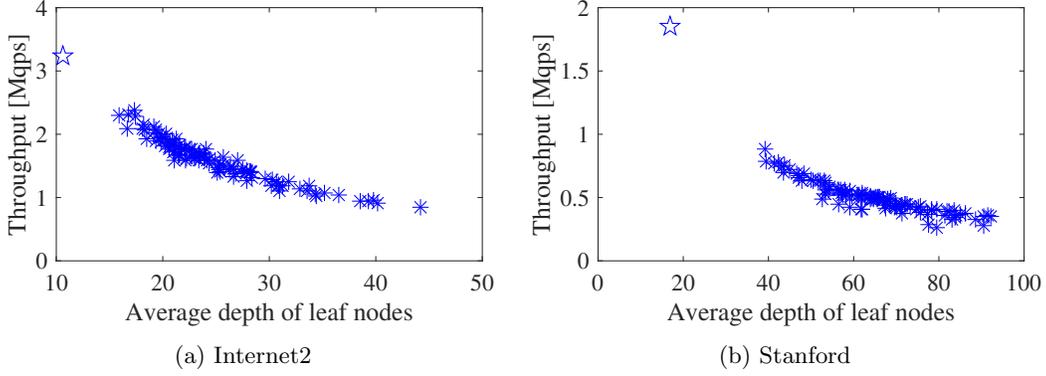


Figure 4: Query throughput versus average depth of leaves

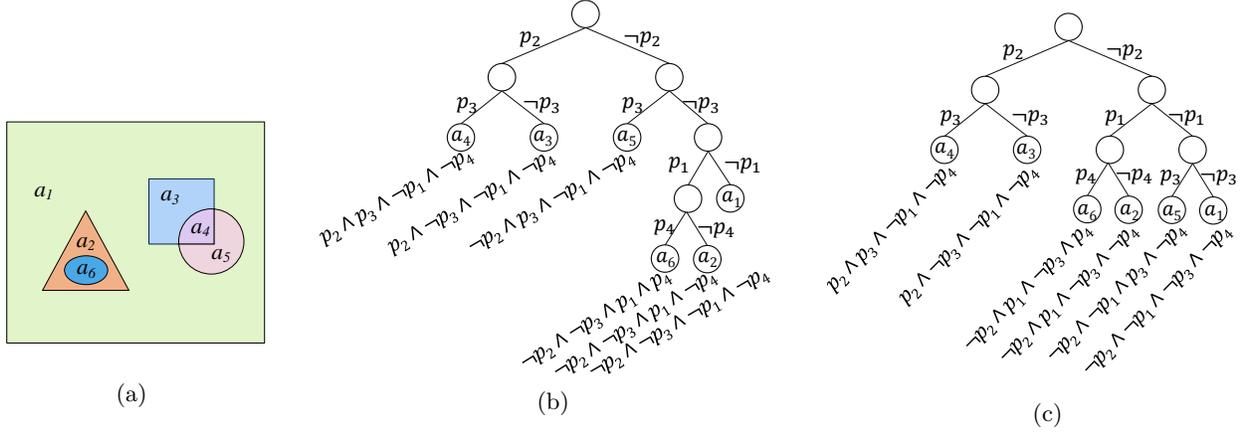


Figure 5: Additional example

throughput. The star in each figure represents the performance of the AP Tree constructed by AP Classifier. The query throughput of AP Classifier is 3.35 Mqps (Internet2) and 1.82 Mqps (Stanford), substantially higher than any random construction.

## 5.2 Quick-Ordering algorithm

Given a fixed number of leaves, a balanced binary tree minimizes the average leaf depth. Compare the two AP Trees in Fig. 2(b) and (c) whose average depths are 2.6 and 2.4, respectively. The one in Fig. 2(c) is more balanced and hence has less average depth. The reason for the imbalance in Fig. 2(b) is that  $p_1$  is placed at a higher level of the tree. According to properties of atomic predicates, every predicate is equal to the disjunction of a subset of atomic predicates. The number varies from one to the number of all atomic predicates. In this example,  $p_1$  is a predicate that is equal to a single atomic predicate. Hence the left child of the node labeled as  $p_1$  must be a leaf node representing the atomic predicate. However, the right sub-tree may include more levels, causing the imbalance.

In fact, an analysis of the two real network data planes shows that many predicates are equal to a single

atomic predicate. One fast yet effective ordering of predicates is to place those predicates at lower levels. For example, in Fig. 2(c),  $p_1$  is placed at the lowest level.

**Notation.** Let  $R(p)$  denote the subset of atomic predicates whose disjunction is  $p$ .  $|R(p)|$  denotes the cardinality of  $R(p)$ .

In the *Quick-Ordering algorithm*,  $|R(p_i)|$  is counted for each predicate  $p_i$ . Then the AP Tree is constructed by placing all predicates onto the tree in descending order of  $|R(p_i)|$ .

## 5.3 Optimized AP Tree construction

Developing on Quick-Ordering, we propose a more sophisticated ordering method that further reduces the average depth of the AP Tree.

One important observation is that, for two sub-trees whose roots are siblings, their predicate orders can be different. In the example of Fig. 5(a), we now have four predicates  $p_1$  (triangle),  $p_2$  (square),  $p_3$  (circle), and  $p_4$  (ellipse), which determine six atomic predicates,  $a_1$  to  $a_6$ . If the predicates are added in the order  $p_2, p_3, p_1, p_4$ , the pruned AP Tree is shown in Fig. 5(b). However, for the sub-tree rooted at the right child of the root, its subtree is more balanced if the predicate order is

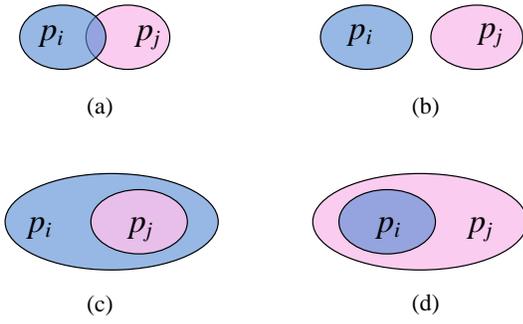


Figure 6: Relationships of two predicates

$p_1, p_3, p_4$ , as shown in Fig. 5(c).

For a given set of predicates  $P = \{p_1, p_2, \dots, p_k\}$ , the atomic predicates  $A = \{a_1, a_2, \dots, a_n\}$  is determined. The number of leaf nodes of the AP Tree is  $n$ , because each leaf node corresponds to an atomic predicate. We define  $F(Q, S)$  as the *minimal sum of leaf depths* of the subtree (which is a part of the AP Tree) whose nodes include the set of predicates  $Q$  and leaves are the set of atomic predicates  $S$ . In the example of Fig. 5(c), let  $Q = \{p_1, p_3, p_4\}$  and  $S = \{a_1, a_2, a_5, a_6\}$ ,  $F(Q, S) = 8$ .  $F(Q, S)$  can be calculated recursively using the following equations. Let  $H(Q, S, p)$  be the minimal sum of leaf depths if the root of the sub-tree is  $p$ . If  $S \cap R(p_i) \neq \emptyset$  and  $S \cap R(\neg p_i) \neq \emptyset$ ,  $H(Q, S, p)$  is the sum of three components:  $F(Q \setminus \{p\}, S \cap R(p))$  and  $F(Q \setminus \{p\}, S \cap R(\neg p))$  are recursive computing for the left and right sub-trees and extra  $|S|$  needs to be added because the depth of every leaf increments by 1. We have

$$H(Q, S, p) = F(Q \setminus \{p\}, S \cap R(p)) + F(Q \setminus \{p\}, S \cap R(\neg p)) + |S|$$

If  $S \cap R(p) = \emptyset$ , the left sub-tree will be pruned and the leaf depths do not increase. Hence,

$$H(Q, S, p) = F(Q \setminus \{p\}, S \cap R(\neg p))$$

Similarly, if  $S \cap R(\neg p) = \emptyset$ , we have,

$$H(Q, S, p) = F(Q \setminus \{p\}, S \cap R(p))$$

In addition, we have the following recursive equation.

$$F(Q, S) = \begin{cases} 0 & \text{if } |S| = 1 \\ \min_{p_i \in Q} H(Q, S, p_i) & \text{otherwise} \end{cases} \quad (1)$$

When  $|S| = 1$ , it is easy to see that the sub-tree contains only one leaf, hence  $F(Q, S) = 0$ . Otherwise, the predicate  $p_i \in Q$  is selected as the root of the sub-tree such that  $p_i$  minimizes  $H(Q, S, p_i)$ .

Using the above formula, it is possible to compute  $F(P, A)$ . By recording the selection of  $p_i$  at each recursion, the optimized AP Tree can also be constructed.

However, the time complexity of solving this recursion is as high as  $O((2^k) * k!)$ , where  $k$  is the cardinality of  $P$ . We need to propose an efficient heuristic algorithm to simplify the recursion. At a level of recursion, we need to find the predicate  $p_i$  that minimizes

$H(Q, S, p_i)$ . Instead of trying all predicates, we propose an easier way to decide which predicate to select.

We define a pair-wise relation between two predicates that implies which one is better to select. If  $H(Q, S, p_i) < H(Q, S, p_j)$ , we say that  $p_i$  is *superior* to  $p_j$  and  $p_j$  is *inferior* to  $p_i$ , denoted as  $p_i \xrightarrow{S} p_j$ . If  $H(Q, S, p_i) = H(Q, S, p_j)$ , we say  $p_i$  and  $p_j$  are in the same order, denoted as  $p_i \overset{S}{\sim} p_j$ .

We compare two predicates in four cases based on their logical relationships, as shown in Fig. 6. Here,  $p_i$  and  $p_j$  refer to predicates which are equal to union of atomic predicates in  $S \cap R(p_i)$  and  $S \cap R(p_j)$  respectively.

1) Packets specified by  $p_i$  intersect with those of  $p_j$  (Fig. 6(a)). If we place  $p_i$  to the root and  $p_j$  to the children of the root, we have

$$\begin{aligned} H(Q, S, p_i) &= |S| + F(Q - \{p_i\}, S \cap R(p_i)) \\ &\quad + F(Q - \{p_i\}, S \cap R(\neg p_i)) \\ &= |S| + F(Q - \{p_i, p_j\}, S \cap R(p_i) \cap R(p_j)) \\ &\quad + F(Q - \{p_i, p_j\}, S \cap R(p_i) \cap R(\neg p_j)) \\ &\quad + |S \cap R(p_i)| \\ &\quad + F(Q - \{p_i, p_j\}, S \cap R(\neg p_i) \cap R(p_j)) \\ &\quad + F(Q - \{p_i, p_j\}, S \cap R(\neg p_i) \cap R(\neg p_j)) \\ &\quad + |S \cap R(\neg p_i)| \end{aligned}$$

If we place  $p_j$  to the root and  $p_i$  to the children, we can get  $H(Q, S, p_j)$  similarly. Since  $|S \cap R(p_i)| + |S \cap R(\neg p_i)| = |S \cap R(p_j)| + |S \cap R(\neg p_j)| = |S|$ ,  $H(Q, S, p_i) = H(Q, S, p_j)$ . We have  $p_i \overset{S}{\sim} p_j$ .

2) Packets specified by  $p_i$  disjoint with those of  $p_j$  (Fig. 6(b)).  $p_i \wedge p_j$  is false. If we place  $p_i$  to the root and  $p_j$  to the children of the root, the sub-tree representing  $p_i \wedge p_j$  will be pruned. Hence

$$\begin{aligned} H(Q, S, p_i) &= |S| + F(Q - \{p_i, p_j\}, S \cap R(p_i) \cap R(\neg p_j)) \\ &\quad + F(Q - \{p_i, p_j\}, S \cap R(\neg p_i) \cap R(p_j)) \\ &\quad + F(Q - \{p_i, p_j\}, S \cap R(\neg p_i) \cap R(\neg p_j)) \\ &\quad + |S \cap R(\neg p_i)| \end{aligned}$$

Similarly, if we place  $p_j$  to the root and  $p_i$  to the children,

$$\begin{aligned} H(Q, S, p_j) &= |S| + F(Q - \{p_i, p_j\}, S \cap R(p_j) \cap R(\neg p_i)) \\ &\quad + F(Q - \{p_i, p_j\}, S \cap R(\neg p_j) \cap R(p_i)) \\ &\quad + F(Q - \{p_i, p_j\}, S \cap R(\neg p_j) \cap R(\neg p_i)) \\ &\quad + |S \cap R(\neg p_j)| \end{aligned}$$

Despite of the same terms, if  $|S \cap R(\neg p_i)| < |S \cap R(\neg p_j)|$ ,  $p_i \xrightarrow{S} p_j$ . If  $|S \cap R(\neg p_i)| = |S \cap R(\neg p_j)|$ ,  $p_i \overset{S}{\sim} p_j$ . Otherwise  $p_j \xrightarrow{S} p_i$ .

3) Packets specified by  $p_j$  are a subset of those of  $p_i$ . If we place  $p_i$  to the root and  $p_j$  to the children of the root, the sub-tree representing  $\neg p_i \wedge p_j$  will be pruned.

Hence

$$\begin{aligned} H(Q, S, p_i) = & |S| + F(Q - \{p_i, p_j\}, S \cap R(p_i) \cap R(p_j)) \\ & + F(Q - \{p_i, p_j\}, S \cap R(p_i) \cap R(\neg p_j)) \\ & + F(Q - \{p_i, p_j\}, S \cap R(\neg p_i) \cap R(\neg p_j)) \\ & + |S \cap R(p_i)| \end{aligned}$$

If we place  $p_i$  to the root and  $p_j$  to the children of the root, the sub-tree representing  $p_j \wedge \neg p_i$  will be pruned.

$$\begin{aligned} H(Q, S, p_j) = & |S| + F(Q - \{p_i, p_j\}, S \cap R(p_j) \cap R(p_i)) \\ & + F(Q - \{p_i, p_j\}, S \cap R(\neg p_j) \cap R(p_i)) \\ & + F(Q - \{p_i, p_j\}, S \cap R(\neg p_j) \cap R(\neg p_i)) \\ & + |S \cap R(\neg p_j)| \end{aligned}$$

Therefore if  $|S \cap R(p_i)| < |S \cap R(\neg p_j)|$ ,  $p_i \xrightarrow{S} p_j$ . If  $|S \cap R(p_i)| = |S \cap R(\neg p_j)|$ ,  $p_i \xrightarrow{S} p_j$ . Otherwise  $p_j \xrightarrow{S} p_i$ .

4) Packets specified by  $p_i$  are a subset of those of  $p_j$ . Similar to the above cases, we can get if  $|S \cap R(\neg p_i)| < |S \cap R(p_j)|$ ,  $p_i \xrightarrow{S} p_j$ . If  $|S \cap R(\neg p_i)| = |S \cap R(p_j)|$ ,  $p_i \xrightarrow{S} p_j$ . Otherwise  $p_j \xrightarrow{S} p_i$ .

We then design the key criterion of predicate selection for each level of recursion, namely: *We select a predicate that is not inferior to any other predicate.* The **algorithm** is presented as follows: For each level of recursion, a predicate  $p_s$  is maintained, initially being  $p_1$ . A linear scan is performed from  $p_2$  to  $p_k$ . For a predicate  $p_i$ , if  $p_i \xrightarrow{S} p_s$ , then  $p_s$  is set to  $p_i$ . At the end,  $p_s$  is selected as the root node of the subtree for this level of recursion.

To prove the correctness of the above algorithm, we need to show that  $p_s$  is indeed not inferior to any other predicate. A sufficient condition is that the superior/inferior relation is acyclic, i.e., there are no three predicates  $p_a, p_b, p_c$  such that  $p_a \xrightarrow{S} p_b$ ,  $p_b \xrightarrow{S} p_c$ , and  $p_c \xrightarrow{S} p_a$ . We have proved the acyclic property by exhaustion. Our proof is not shown herein due to space limitation.

**Time efficiency of AP Tree construction.** In the AP Tree construction algorithm presented above, we avoid the time-intensive operation of computing the conjunction of two predicates represented as BDDs. Instead, our algorithm computes the intersection of two sets of integers that are identifiers of atomic predicates, as suggested in [30]. Intersections of integer sets can be computed much more quickly than conjunctions of BDDs.

The time complexity of determining relationship between two predicates is  $O(n)$ , where  $n$  is the number of atomic predicates. For each level of recursion, a linear scan needs  $O(k'n)$  time, where  $k'$  is the number of predicates in the current level. The overall complexity of building an AP Tree depends on the number of levels as well as the balance of the tree. Here we only provide the complexity analysis for a balanced AP Tree. For a balanced AP Tree, there are  $2^l$  nodes at level  $l$ . For each node,  $k' \leq (k-l)$ . Hence at level  $l$ , the time complexity

is at most  $2^l(k-l)n$ . Since  $l \leq \log_2 n$ ,  $2^l(k-l)n < kn^2$ . Since there are  $\lceil \log_2 n \rceil$  levels, the overall time complexity is upper-bounded by  $O(kn^2 \log n)$ .

## 6. AP TREE UPDATE AND RECONSTRUCTION

An important requirement of practical packet behavior identification is to support dynamic network changes, including link and rule changes, both of which require addition and deletion of predicates. We design fast AP Tree update methods for adding a predicate and deleting a predicate while maintaining tree correctness. However, after a large number of updates, an AP Tree will experience performance degradation. Hence we also design an AP Tree reconstruction method that periodically rebuilds the tree to optimize its performance while performing packet query processing at the same time.

### 6.1 Real-time update of an AP Tree

The SDN data plane of a network is frequently updated by rule installation and deletion. When a rule is inserted into or removed from a forwarding table or an ACL, it may change one or more predicates. The set of atomic predicates may change as well. We use the method presented in [29] to convert a rule insertion or deletion to predicate change. If there is no predicate change after a rule update, AP Classifier does not need to update the AP Tree. Otherwise, AP Classifier performs the methods presented below to remove the old predicate and add the updated predicate in the AP Tree. These methods are also used after addition/deletion of a network link which requires addition/deletion of predicates.

**Add a predicate.** When a new predicate  $p$  is added, for each leaf node representing an atomic predicate  $a$  in the current AP Tree, AP Classifier computes  $a \wedge p$  and  $a \wedge \neg p$ . If none of them is false, two children are added to the leaf node, representing  $a \wedge p$  and  $a \wedge \neg p$  respectively. If one and only one of the two conjunctions is false, the label of the leaf node is replaced by the other conjunction. If both conjunctions are false, AP Classifier does nothing to this leaf node.

**Delete a predicate.** To delete an existing predicate  $p$  from the AP Tree, AP Classifier does not remove all internal nodes labeled by  $p$ . This is because after the removal of a node, merging the two sub-trees rooted at its children is very difficult. Instead, we still keep  $p$  in the AP Tree, but mark it as “deleted” in the list of all predicates. A query packet is still processed by the AP Tree to find its leaf node representing its atomic predicate. It is still evaluated by the deleted predicates to determine which sub-tree to visit next. However, in the second stage of AP Classifier, i.e., computing packet behaviors, AP Classifier just ignores all predicates that have been deleted.

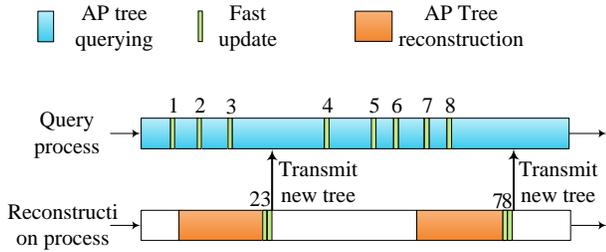


Figure 7: Real-time update and query processing

## 6.2 Parallel reconstruction of an AP Tree

Although, the AP Tree updates in AP Classifier are fast and maintain correctness of packet behavior identification, the AP Tree is no longer optimized and the query throughput will degrade over time. Hence AP Classifier also reconstructs the AP Tree to optimize it from time to time. To enable query processing at the same time as tree reconstruction, AP Classifier runs two processes in parallel, called the *query process* and *reconstruction process*, executing on two different cores. The start of a reconstruction is triggered by an event, e.g., query throughput is lower than a threshold or the number of updates on the current AP Tree is higher than a threshold. During reconstruction, the query process still maintains the old AP Tree by performing updates, and responds to queries. After the reconstruction process has built a new tree, the new tree needs to be updated for data plane changes that have occurred during the reconstruction period, if any. The updated new tree is then transmitted to the query process to replace the old tree.

Fig. 7 shows an example of the parallel reconstruction of an AP Tree. The query process performs AP Tree search to respond to queries as well as updates when data plane changes happen. In this example, the first reconstruction starts shortly after the change that requires update 1, which is included in the construction of a new tree. However, when the new tree is finished, two changes that require updates 2 and 3 have occurred during the reconstruction period. The new tree does not reflect these two updates. Thus the reconstruction process also applies these two updates to the new tree. Then the updated new tree is sent to the query process to replace the old AP Tree. Similarly the second reconstruction begins after changes that require updates 4, 5, and 6. The new tree constructed needs to be updated for changes (that require updates 7 and 8) which occur during the reconstruction period, before it can be sent to the query process. Note that if there is no data plane change during a reconstruction period, the new AP Tree is optimized.

## 7. EXPERIMENTAL EVALUATION

We have implemented and evaluated AP Classifier on

Table 1: Statistics of the two real networks

	Stanford		Internet2
No. of rules	Forwarding	ACL	Forwarding
	757170	1584	126017
No. of predicates	507	71	161
No. of atomic pred.	494	21	216

a general purpose desktop computer with quadcore@3.2G and 16GB memory. Our implementation and evaluation include all functional components for packet behavior identification from scratch, including computing atomic predicates, classifying packets using the AP Tree, and computing packet behaviors. (In comparison, prior work on this problem only implements and evaluates a single function, namely: classifying packets to equivalence classes [13].) For our experimental evaluation, we use forwarding tables and ACLs from two real networks: Internet2 [2] and Stanford network [17]. Internet2 includes 126,017 forwarding rules and the Stanford network includes 757,170 forwarding rules and 1,584 ACL rules. The predicates and atomic predicates are computed using the method in [30]. We compare AP Classifier with possible solutions by utilizing two state-of-art tools, namely Header Space Analysis (HSA) [17] and AP Verifier [30]. We do not compare AP Classifier with MDD [13] because it relies on a special method for MDD construction and the source code is not publicly available. Furthermore, its method does not support dynamic updates.

### 7.1 Depths of leaf nodes

In this set of experiments, we show the depths of leaf nodes in an AP Tree, which can demonstrate effectiveness of the proposed tree construction algorithms. We evaluate and compare three methods, Best from Random, Quick-Ordering, and Optimized AP Tree construction (OAPT), for both Internet2 and Stanford networks. The Best from Random method generates a random order of predicates for placement on levels of an AP tree and performs pruning. It constructs 100 AP trees and chooses the tree with the minimal average depth of leaf nodes. Quick-Ordering is presented in Section 5.2 and OAPT is presented in Section 5.3.

Fig. 8 shows the average depth of leaf nodes in an AP tree. For Internet2, the average depth of Best from Random is 16.0, worse than those of Quick-Ordering (13.0) and OAPT (10.6). OAPT reduces the average depth by 34% compared to Best from Random and 19% compared to Quick-Ordering. For the Stanford network, Best from Random also has the highest average depth (39.0), followed by Quick-Ordering (24.2) and OAPT (16.9). OAPT shows significant improvement: It reduces the average depth by 57% compared to Best from Random and by 30% compared to Quick-Ordering.

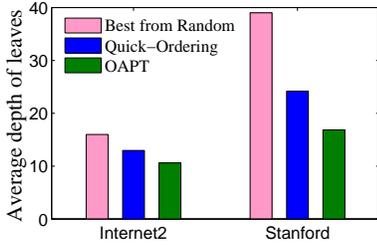
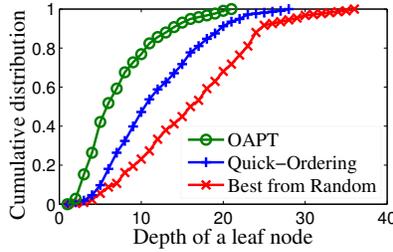
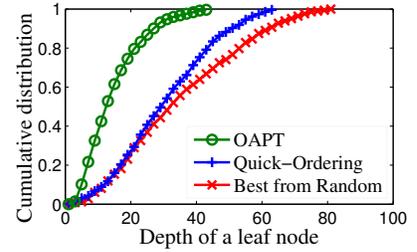


Figure 8: Average depth of leaves



(a) Internet2



(b) Stanford

Figure 9: Cumulative distribution of the depths of leaf nodes in AP Trees

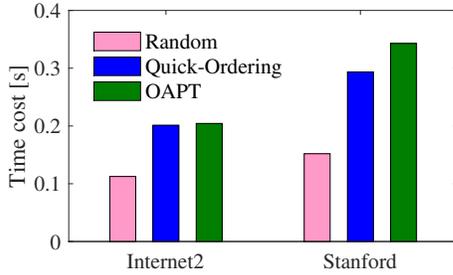


Figure 10: Overall construction time of AP Classifier

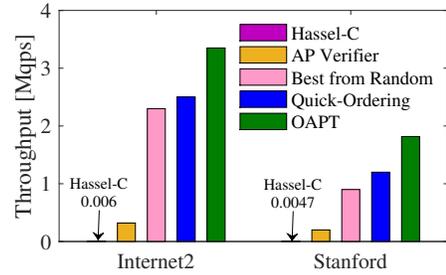


Figure 11: Query throughput for static networks

Fig. 9 shows the cumulative distribution of depths of leaf nodes in an AP Tree. For Internet2, the leaf depths of Quick-Ordering are clearly smaller than Best from Random. However for the Stanford network such improvement is not very significant. OAPT has clearly smaller depths for all percentiles compared to the other two methods. For Internet2 80% of the leaf nodes in the OAPT tree have a depth less than 11 and for Stanford this number is 21. The maximum depths are 24 and 46 for Internet2 and Stanford, respectively.

## 7.2 Memory Usage

AP Classifier stores all predicates and atomic predicates as BDDs and also, for each predicate, a set of integer identifiers of atomic predicates. In the AP Tree a node only stores a pointer to the labeled predicate or atomic predicate. Since pointers use very little memory, the memory costs of different methods are very close. Hence we only show the memory cost of AP Classifier using OAPT.

The total memory cost of AP Classifier for Internet2 is 4.79 MB and that for Stanford is 2.15 MB. Although Internet2 has fewer predicates than Stanford, it requires more memory because BDDs of the Internet2 predicates are more complex than those of Stanford. Unlike the results of [13] that only show memory cost of the search structure, our memory costs account for all components for packet behavior identification, including the network topology, predicates, atomic predicates, and AP Tree. We found that AP Classifier uses very small memory

and can be stored in cache.

## 7.3 AP Tree construction time

Fig. 10 shows times to construct AP Trees using the three methods for the two networks. Note that the time cost is the overall construction time that includes the times for computing atomic predicates as well as for AP Tree construction. The Random method costs the least time but it is only for one random construction. To find the best AP Tree from a large number of random constructions takes substantially longer time. Quick-Ordering and OAPT have similar time costs, 201.36 ms and 204.39 ms, for Internet2. For the Stanford network, OAPT requires 342.77 ms for Stanford, a little longer compared to Quick-Ordering (293.36 ms).

## 7.4 Query throughput for static networks

In this set of experiments, we measure the throughput of AP Classifier to process packet queries, in number of queries per second (qps). Packet headers used for queries in the experiments are generated randomly with respect to the atomic predicates. The throughput results for static networks are shown in Fig. 11. For Internet2, AP Classifier using OAPT can achieve 3.4 Mqps, higher than Best from Random by 102% and Quick-Ordering by 52%. For Stanford network, AP Classifier using OAPT can achieve 1.8 Mqps, higher than Best from Random by 46% and Quick-Ordering by 34%. For both networks, the throughput of AP Classifier is much higher than 1 Mqps, which is enough to satisfy most

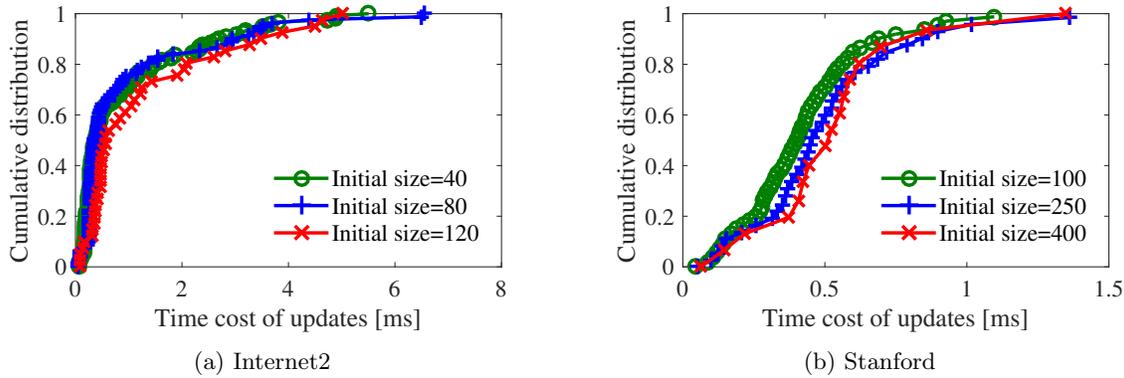


Figure 12: Cumulative distributions of time cost for adding a predicate.

application requirements in SDN.

For static networks, we can use the open-source tool Hassel-C [1] that implements HSA [17] to perform packet behavior identification for a specific packet. By providing the input port and a specific query packet, Hassel-C computes the reachability tree of the query packet. (For a unicast packet, the reachability tree is a forward path to the packet’s destination.) The query throughputs of using Hassel-C to perform packet behavior identification are 6 Kqps and 4.7 Kqps for Internet2 and Stanford, respectively, which are about 1000 times slower than the query throughputs of AP Classifier. They are also plotted in Fig. 11 but they are very small and barely visible. We also compare AP Classifier with AP Verifier [30]. We first use AP Verifier to compute all atomic predicates, and perform a linear search of all atomic predicates for the query packet until the packet matches an atomic predicate. Results in Fig. 11 show that AP Verifier is also much slower, though its throughput is improved a lot compared to Hassel-C.

## 7.5 Dynamic Networks

In this set of experiments, we first construct the AP Tree using a number of predicates and then keep adding new predicates. We measure the time cost to add each new predicate and update the AP Tree. Fig. 12 (a) shows the cumulative distribution of time cost for adding a predicate in the Internet2 network. The initial number of predicates is set to 40, 80, and 120 for three different experiments. From the figure we find that about 80% of the predicate additions are finished in 2 ms. It may take 5-6 ms in worst cases. We do not observe obvious differences when the initial numbers of predicates are different. Fig. 12 (b) shows the results of similar experiments for Stanford. The initial number of predicates is set to 100, 250, and 400 for three different experiments. Over 90% of the predicate additions are finished in 1 ms. Deleting a predicate does not require extra computation, hence there is no result for deletions.

**Query throughput for dynamic networks.** We also evaluate the throughput of AP Classifier in practi-

cal environments where additions and deletions of rules and predicates happen over time. At the beginning of each experiment, a number of predicates are chosen randomly from the set of predicates of a network to construct the initial AP Tree. Starting from time 0, the arrivals of change events requiring the addition or deletion of predicates are modeled by a Poisson process. Each update operation can be adding a new predicate or deleting an existing predicate. In all experiments, equal numbers of additions and deletions are inserted to the event queue. A reconstruction is triggered every 0.4 s. During every reconstruction, AP Classifier answers queries and performs updates as explained in Section 6.2. We compare AP Classifier with two possible methods, APLinear and PScan, APLinear utilizes AP Verifier [30] to compute atomic predicates and performs a linear search for the query packet until the packet matches an atomic predicate. Note that BDDs of atomic predicates are more complex than those of predicates. Hence APLinear is not efficient. PScan performs a scan on all predicates using the query packet and decides whether the packet is filtered by the predicate. Both methods can be used to identify packet behaviors.

Fig. 13 shows the throughputs of AP Classifier, APLinear, and PScan in dynamic networks. The  $x$ -axis is time and the  $y$ -axis is throughput measured in Mqps. We conduct two sets of experiments whose update rates are 100 updates/s and 200 updates/s. From all subfigures in Fig. 13, we find that AP Classifier is faster than the other two methods by an order of magnitude.

Note that starting from time 0, the throughput of AP Classifier slowly decreases as an increasing number of updates make the AP Tree less optimized. The first reconstruction starts at time 0.4 s and finishes at about 0.6 s in Fig. 13(a) and (c), and 0.7s in Fig. 13(b) and (d). When a reconstruction finishes, the throughput immediately goes back to a high value (4 Mqps in (a) and (c), and 2 Mqps in (b) and (d)). Furthermore, the throughput does not degrade in the long-term view. Comparing results of the two different update rates, we

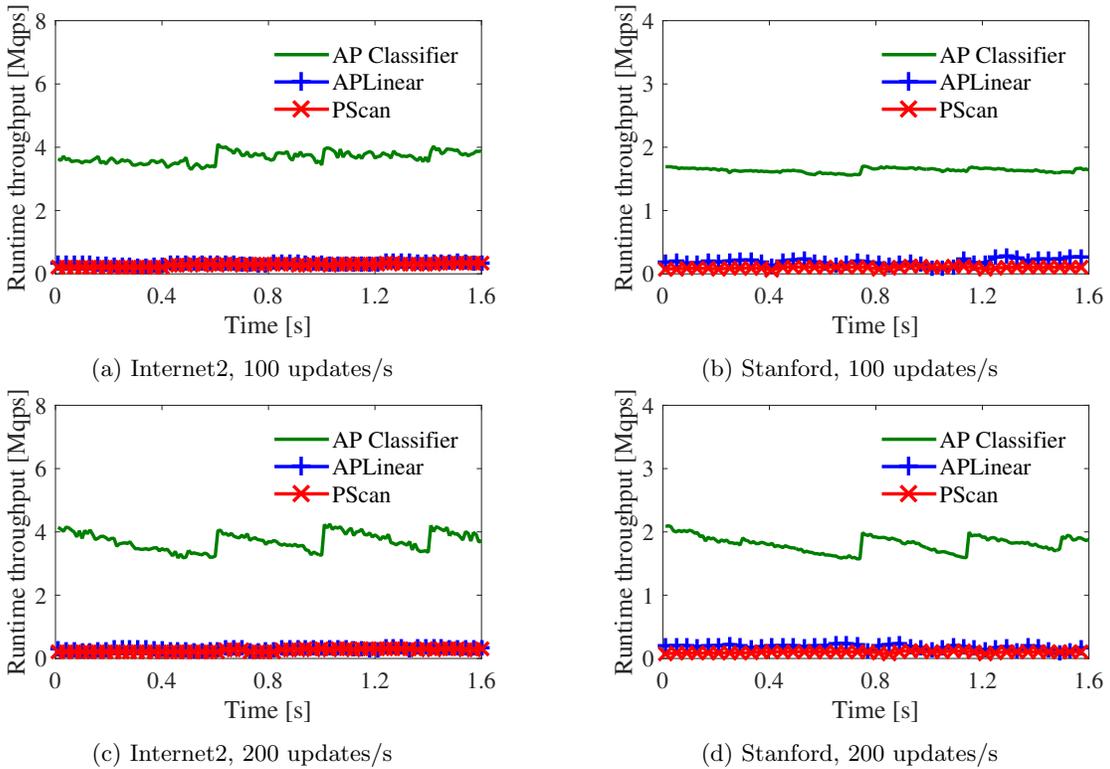


Figure 13: Query throughput for dynamic networks.

find that the average throughput of AP Classifier does not drop much even after the update rate is doubled. Hence AP Classifier is fast and robust for practical dynamic networks.

## 8. DISCUSSION

**Optimization for packet distribution.** In the proposed algorithms of AP Classifier, we assume that, for a packet query, leaf nodes (atomic predicates) have equal probability to be visited. Therefore minimizing the average depth of leaf nodes maximizes the query throughput. However, practical network flows may not be distributed uniformly with respect to the set of atomic predicates. For example, if many queried packets may eventually visit a leaf in a very deep position and leaves close to the root are rarely visited, the throughput decreases. To improve the query throughput for uneven packet distribution, we assign weights to atomic predicates such that leaf nodes that are visited frequently will be placed relatively close to the root. The experimental results also show that our method is very robust to different packet distributions. Detailed design and experimental results are skipped due to page limit.

**Dealing with Packet Header Changes.** In modern networks, some middleboxes may make changes to packet headers, such as Network Address Translation (NAT). In these cases, when the AP Classifier searches the AP Tree using the particular packet header, it will

tell that the packet stops at the middlebox. If AP Classifier realizes that the packet stops at a middlebox, it needs to check whether the packet is dropped, e.g., by a firewall, or the packet header has been changed based on the middlebox policies. If it is the later case, the AP Classifier will consider the packet to be a new one and compute the forwarding behaviors for the new packet header by searching the AP Tree again. Such process may repeat multiple times until the computed path ends at the packet destination. It is hard for AP Classifier to handle the middleboxes that determine packet behaviors based on high-level payloads or real time policies, as such load balancers and adaptive routing methods [10, 31] in data center networks. We leave a full design to future work.

## 9. CONCLUSION

We propose AP Classifier for network-wide packet behavior identification that can be utilized by many important network management applications. We design algorithms to construct the AP Tree for a network, which can be used to quickly classify a packet to an atomic predicate. Each atomic predicate represents the network-wide forwarding behaviors of a set of packets. Experimental results using the datasets of two real networks show that the proposed AP Tree construction algorithm can optimize the average depth of leaf nodes. AP Classifier can process millions of packet queries per

second. The speed is faster than existing tools by at least an order of magnitude. Furthermore, it uses only a few MBs memory. It can be updated in real time and is robust under dynamic data plane changes.

## 10. ACKNOWLEDGEMENT

The authors thank the CoNEXT reviewers for their constructive comments and suggestions. Huazhe Wang, Chen Qian, and Ye Yu were sponsored by National Science Foundation grant CNS-1464335 and University of Kentucky College of Engineering Startup Grant. Hongkun Yang and Simon S. Lam were supported by National Science Foundation grant CNS-1214239.

## 11. REFERENCES

- [1] Hassel-C. <http://bitbucket.org/peymank/hassel-public/>.
- [2] The internet2 observatory data collections. <http://www.internet2.edu/observatory/archive/data-collections.html>.
- [3] University of oregon route views project. <http://www.routeviews.org>.
- [4] S. Agarwal, M. Kodialam, and T. Lakshman. Traffic engineering in software defined networks. In *Proc. of IEEE INFOCOM*, 2013.
- [5] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat. Hedera: dynamic flow scheduling for data center networks. In *Proc. of USENIX NSDI*, 2010.
- [6] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron. Towards predictable datacenter networks. In *Proc. of ACM SIGCOMM*, 2011.
- [7] T. Benson, A. Akella, and D. A. Maltz. Network traffic characteristics of data centers in the wild. In *Proc. of ACM IMC*, 2010.
- [8] T. Benson, A. Anand, A. Akella, and M. Zhang. Microte: Fine grained traffic engineering for data centers. In *Proc. of ACM CoNEXT*, 2011.
- [9] Q. Chen, C. Qian, and S. Zhong. Privacy-preserving cross-domain routing optimization – a cryptographic approach. In *Proc. of IEEE ICNP*, 2015.
- [10] W. Cui and C. Qian. Difs: Distributed flow scheduling for adaptive routing in hierarchical data center networks. In *Proc. of ACM/IEEE ANCS*, 2014.
- [11] M. Dhawan, R. Poddar, K. Mahajan, and V. Mann. Dynamic scheduling of network updates. In *Proc. of ACM SIGCOMM*, 2014.
- [12] A. Fogel, S. Fung, L. Pedrosa, M. Walraed-Sullivan, R. Govindan, R. Mahajan, and T. Millstein. A general approach to network configuration analysis. In *Proc. of USENIX NSDI*, 2015.
- [13] T. Inoue, T. Mano, K. Mizutani, S. Minato, and O. Akashi. Rethinking packet classification for global network view of software-defined networking. In *Proc. of IEEE ICNP*, 2014.
- [14] S. Jain et al. B4: Experience with a Globally-Deployed Software Defined WAN. In *Proceedings of ACM Sigcomm*, 2013.
- [15] S. Kandula, S. Sengupta, A. Greenberg, P. Patel, and R. Chaiken. The nature of data center traffic: measurements & analysis. In *Proc. of ACM IMC*, 2009.
- [16] P. Kazemian, M. Chang, H. Zeng, G. Varghese, N. McKeown, and S. Whyte. Real time network policy checking using header space analysis. In *Proc. of USENIX NSDI*, 2013.
- [17] P. Kazemian, G. Varghese, and N. McKeown. Header space analysis: Static checking for networks. In *Proc. of USENIX NSDI*, 2012.
- [18] A. R. Khakpour and A. X. Liu. Quantifying and querying network reachability. In *Proc. of IEEE ICDCS*, 2010.
- [19] A. Khurshid, X. Zou, W. Zhou, M. Caesar, and P. B. Godfrey. Veriflow: Verifying network-wide invariants in real time. In *Proc. of USENIX NSDI*, 2013.
- [20] M. Kuzniar, P. Peresini, and D. Kostic. What you need to know about SDN flow tables. In *Proc. of PAM*, 2015.
- [21] X. Li and C. Qian. Traffic and failure aware vm placement for multi-tenant cloud computing. In *Proceedings of IEEE IWQoS*, 2015.
- [22] H. Liu, S. Kandula, R. Mahajan, M. Zhang, and D. Gelernter. Traffic engineering with forward fault correction. In *Proc. of ACM SIGCOMM*, 2014.
- [23] N. P. Lopes, N. Bjørner, P. Godefroid, K. Jayaraman, and G. Varghese. Checking beliefs in dynamic networks. In *Proc. of USENIX NSDI*, 2015.
- [24] H. Mai, A. Khurshid, R. Agarwal, M. Caesar, P. B. Godfrey, and S. T. King. Debugging the data plane with Anteater. In *Proc. of ACM SIGCOMM*, 2011.
- [25] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. Openflow: enabling innovation in campus networks. *ACM SIGCOMM CCR*, 38(2):69–74, 2008.
- [26] A. Nayak, A. Reimers, N. Feamster, and R. Clark. Resonance: Dynamic access control for enterprise networks. In *Proc. of ACM WREN*, 2009.
- [27] Z. A. Qazi, C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu. Simple-fying middlebox policy enforcement using SDN. In *Proc. of ACM SIGCOMM*, 2013.
- [28] G. Xie, J. Zhan, D. A. Maltz, H. Zhang, A. Greenberg, G. Hjalmtysson, and J. Rexford. On static reachability analysis of IP networks. In *Proc. of IEEE INFOCOM*, 2005.
- [29] H. Yang and S. S. Lam. Real-time verification of network properties using atomic predicates. Technical Report TR-13-15, The Univ. of Texas at Austin, Dept. of Computer Science, Aug. 2013.
- [30] H. Yang and S. S. Lam. Real-time verification of network properties using atomic predicates. In *Proc. of IEEE ICNP*, 2013, extended version in *IEEE/ACM Transactions on Networking*.
- [31] Y. Yu and C. Qian. Space shuffle: A scalable, flexible, and high-bandwidth data center network. In *Proc. of IEEE ICNP*, 2014.
- [32] Y. Yu, C. Qian, and X. Li. Distributed collaborative monitoring in software defined networks. In *Proc. of ACM HotSDN*, 2014.
- [33] H. Zeng, P. Kazemian, G. Varghese, and N. McKeown. Automatic test packet generation. In *Proc. of ACM CoNEXT*, 2012.