

# Alpaca: Compact Network Policies with Attribute-Encoded Addresses

Nanxi Kang, Ori Rottenstreich, Sanjay G. Rao and Jennifer Rexford

**Abstract**—In enterprise networks, policies (e.g., QoS or security) are often defined based on the categorization of hosts along dimensions such as the organizational role of the host (faculty vs. student), and department (engineering vs. sales). While current best practices (VLANs) help when hosts are categorized along a single dimension, policy may often need to be expressed along multiple orthogonal dimensions. In this paper, we make three contributions. First, we argue for *Attribute-enCoded IPs (ACIPs)*, where the IP address allocation process in enterprises considers attributes of a host along all policy dimensions. ACIPs enable flexible policy specification in a manner that may not otherwise be feasible owing to the limited size of switch rule-tables. Second, we present Alpaca, algorithms for realizing ACIPs under practical constraints of limited-length IP addresses. Our algorithms can be applied to different switch architectures, and we provide bounds on their performance. Third, we demonstrate the importance and viability of ACIPs on data collected from real campus networks.

## I. INTRODUCTION

Managing large enterprise networks is challenging. A typical enterprise has many users who belong to different departments (e.g., sales and engineering, or computer science and history), and play different roles (e.g., faculty, staff, administrators, and students). In addition, the network supports diverse end-hosts running different operating systems and offering different services. In response, network administrators want to enforce policies—such as access control and quality of service—that group hosts along multiple different *dimensions*. For instance, one policy may restrict access to a database to all employees in the sales department, while another may offer a higher bandwidth limit to senior managers across all departments, and yet another may restrict access for old hosts running a less secure operating system.

### A. Enforcing Policies in Today’s Enterprises

To enforce policies in today’s enterprises, network administrators typically rely on virtual local area networks (VLANs) [2]. A host joining the network is assigned to a VLAN based on its MAC address or the physical port of the access switch. Hosts in the same VLAN are assigned an IP address in the same IP prefix, even if they are not located near each other. Traffic flows freely between hosts in the same VLAN, while traffic between different VLANs traverses an IP router that can enforce policy. Since a host can only belong to a

single VLAN, administrators typically assign hosts to VLANs based on a single dimension (e.g., department or role), which has several major limitations:

- The routers interconnecting VLANs need long lists of data-plane “rules” to classify traffic along all relevant “dimensions” of the source and destination hosts.
- No security or QoS policies can be imposed on intra-VLAN traffic, forcing administrators to use a VLAN only to group hosts that should “trust” each other.
- Since VLAN tags are removed from traffic destined to the Internet, the border router must classify all return traffic from the Internet to assign VLAN tags.

For instance, suppose a security policy depends on both user role and department. If VLANs were created based on the user’s department, then expressing a policy based on role in a concise fashion is challenging. And, if two users in the different department are close to each other, the traffic follows inefficient paths through intermediate routers to go between VLANs.

The rise of more flexible network switches, with open interfaces to separate control software, enables an attractive alternative. In recent years, switches built with commodity chipsets expose a pipeline of rule-tables that perform match-action processing on packet headers. While the rule tables are relatively small (with small thousands of rules per stage, to limit power and cost), the switches support programmatic interfaces (e.g., OpenFlow, OF-Config, and OVSDB [3]) that enable new ways of controlling the network by installing the rules in support of higher-level policies. One natural approach, adopted in the early Ethane system [4], directs the first packet of each flow to a central controller, which consults a database—containing all the relevant host attributes and high-level policies—and reactively install rules for forwarding the remaining packets of the flow. However, reactive, fine-grained solutions like Ethane have high overhead and do not scale to large enterprises. FlowTags [5] tags packets and uses the tags to enforce network policies. Yet, this approach, when applied to enterprise networks, requires installing many extra tagging rules at the edge switches to classify hosts along all dimensions. Moreover, these tags can confound middleboxes that don’t understand the tag. Instead, we need a proactive design that can aggregate hosts along many dimensions, while keeping switch rule tables small.

### B. Attribute-encoded IP Addresses

In the current enterprise networks, although IP addresses are assigned based on attributes (e.g., a separate IP prefix

A preliminary version of the paper appeared in ACM *CoNEXT* 2015 [1]. N. Kang is with Databricks (email: knx1029@gmail.com). O. Rottenstreich and J. Rexford are with the Department of Computer Science, Princeton University, NJ, USA (e-mails: {orir, jrex}@cs.princeton.edu). S. Rao is with the School of Electrical and Computer Engineering, Purdue University, IN, USA (email: sanjay@ecn.purdue.edu).

per VLAN), it is just on a *single* dimension. We argue that a complete IP address management scheme should consider *all* dimensions and enable compact representation of policies. Our key idea is to assign each host an IP address based on *all* dimensions of the policy—that is, an *attribute-encoded IP address* (ACIP). Our solution, Alpaca (algorithms for policies on attribute-encoded addresses), computes an *efficient address assignment*, based on the policy dimensions and the attributes of individual hosts, and a *compact list of switch rules* that realize a specific network policy.

Alpaca proactively generates a small number of coarse-grained rules in the switches, without using VLANs. Alpaca greatly simplifies enterprise management by (i) enabling administrators to specify policies on many orthogonal dimensions while achieving an order of magnitude reduction in rules; (ii) allowing policy to be correctly maintained even when a host connects to the network at a new location; and (iii) simplifying federated management, where different teams manage different parts of the enterprise network.

Alpaca is unique in gaining from the flexibility in the address allocation and the possibility to consider multiple dimensions of all hosts in the allocation process. Through extensive experiments we demonstrate that even while using efficient existing rule compression schemes not doing so can result in a potentially large memory increase.

While Alpaca has many potential advantages, several practical issues must be considered:

**Limited IP address space:** Since most enterprises have limited IP address space, a naive ACIP assignment can easily exhaust all the available bits.

**Heterogenous group sizes:** Some combinations of policy attributes are much more common than others. As such, an efficient address assignment cannot simply devote a portion of the bits to representing group size.

**Minimizing churn:** It is important to support changes in the policy attributes associated with a host (*e.g.*, due to a user moving to a different role or department), or changes in the set of attributes themselves (*e.g.*, due to the creation of a new department).

**Multi-stage switch pipelines:** Rather than assuming switches have a single rule table (as in OpenFlow 1.0 [6]), a practical solution should capitalize on the multi-stage pipelines in modern switch chipsets [7]–[9].

We propose a family of algorithms to generate ACIPs and the associated rule tables, starting with a simple strawman that devotes a separate set of address bits to each policy dimension. This solution minimizes the number of rules but consumes too much IP address space, making it infeasible in most practical settings. We then propose two other algorithms that can keep the number of rules small, while respecting constraints on the number of bits in the IP address space. Our algorithms optimize based on the characteristics of modern switch chipsets. Conventionally, switch chip-sets have a single TCAM rule-table that supports a few thousands of wildcard rules matching on multiple header fields [6], [10], [11]. More recently, we have seen the emergence of switch chip-sets with a pipeline of multiple tables, where each table could be a TCAM or a larger SRAM that supports prefix matching

on source IP or destination IP [7]–[9]. As such, our first algorithm generates ACIPs that enable policies to be expressed by solely IP *prefixes*, useful for rule tables that support IP prefix matching. Our second algorithm targets generating rules that perform arbitrary *wildcard* matching on IP addresses, in exchange for a reduction in the number of rules. Together, these algorithms can capitalize on the unique capabilities of a variety of commodity switch architectures.

While assigning IP addresses to hosts, we make sure any pair of hosts are always allocated distinct addresses. This property is maintained even for two hosts that share the same attributes along all dimensions.

Section II presents related work. Next, in Section III, we present a case study of multiple campus networks, to underscore the need for policies along multiple dimensions. Section IV introduces ACIPs and formulates the optimization problem Alpaca must solve, followed by Sections V, VI and VII that present our algorithms. Using access control data from two large campuses, the experiments in Section VIII show that Alpaca can reduce the number of ACL rules on existing networks by 60% – 68% for switches with multiple tables, while requiring only 1 more bit of the IP address space than needed to represent the number of hosts in the network. Further, Alpaca can support futuristic scenarios with policies based on multiple dimensions, while requiring an order of magnitude fewer rules than VLAN-based configurations optimized for a single dimension. We also compare Alpaca with known rule minimization techniques such as ORTC, TCAM Razor and BitWeaving [12]–[14]. In particular, ORTC finds for a given address allocation the optimal representation of a dimension with a minimal number of prefix rules [12]. We demonstrate that taking into account multiple dimensions is crucial to achieve efficient representations. Our results also show the advantages of a flexibility in the address allocation process. Section IX concludes the paper.

## II. RELATED WORK

Various approaches have been described to deal with the limited memory size of rule matching tables. Alpaca is unique in representing an arbitrary number of dimensions with a complete freedom in the address allocation process. It relies on new developed algebra for the performance analysis.

**Address permutation:** Wei et al. propose to swap addresses between two blocks of users to reduce the number of rules [15], but the algorithm can only handle up to two dimensions. Meiners et al. use permutation of the bits in addresses to create prefix patterns so compression algorithms can apply [14]. But it only discovers the optimization potential within the original address allocation. Liu et al. [16] suggested a novel approach of topological transformation to reduce the representation memory size of classifier. Earlier to the lookup that determines the classification value, a transformation is applied to the header fields. This transformation, operated in small lookup tables, enables to minimize the total memory for all tables by taking advantage of the semantics of the specific represented function. While this approach assumes a single classifier, Alpaca considers the representation of an

arbitrary number of dimensions. Another important difference is that in topological transformation, the representation of the transformation is implemented through matching rules and requires memory whose amount is affected by the identity of the selected transformation. In Alpaca addresses are assigned by DHCP servers and it is always better to select an assignment that enables minimal representation of the dimensions.

**Rule optimization:** Minimizing prefix rules matching one header field is easy [12], but minimizing prefix rules or wildcard rules in general cases is NP-hard [13], [17]. Optimal solutions are developed to minimize prefix rules in special cases [12], [18]–[21]. Additional approaches [13], [22], [23] are presented to compress rules in general cases. In particular, [24] suggests to decompose a single rule list into a pipeline of rule lists to minimize the total number of rules. All of these works take the rule-set as an input and explore the potential for minimization, which, in fact, is limited by the original (unoptimized) address allocation. In comparison, Alpaca *generates* the rule list as an output through a smart address allocation process to minimize the number of rules.

**Information encoding:** Huffman coding encodes attributes of a single dimension using prefixes, but its goal is to minimize the weighted sum of the prefix lengths of all attributes. Hence the prefixes do not match the group sizes. SoftCell [25] embeds two dimensional information, *i.e.*, location and middlebox service chain, in the NAT-ed IP addresses. Its encoding mechanism is a special case of BitSegmentation. Concise representation of small sets has been shown useful for the encoding of attributes in SDN [26].

**Attribute-based policy enforcement:** Ethane [4] proposes to implement access control at the network edge by directing the first packet of every flow to a controller, which consults the attributes of hosts and install microflow rules on the switch. FlowTags [5] tag packets based on host attributes and match tags to enforce network policies. Another approach, NetAssay [27], supports network traffic monitoring policies by pushing specific switch-rules for each host given their current IPs. All these work do not optimize IP allocation and install many host-specific rules.

### III. CASE STUDY: DIVERSE ENTERPRISE POLICIES

In this section, we present a case study of 25 enterprise networks, to identify the challenges in representing sophisticated policies, and the implications for Alpaca. Specifically:

- We present a qualitative analysis of the security and quality-of-service policies employed by 22 universities, plus one individual department that runs its network separately from the campus IT group. The analysis indicates that networks must often apply policy along several logical *dimensions*, with multiple *attributes* as possible categories in each dimension. However, the analysis also points to policies that are desirable but difficult to realize in practice.
- We analyze router configuration data from two other large campuses. The analysis provides further confirmation that there is significant commonality in policy across hosts, but also points to how an inefficient assignment of IP addresses can lead to an unnecessary “blow-up” in rule-table size.

- We study host-registration data for one department-level network, to understand the dimensions of network policies and the number and size of host “groups” with these attributes. The analysis has important implications for the design of Alpaca.

#### A. Policies on Multiple Dimensions

Many universities make descriptions of their high-level network policies available online (see <http://tinyurl.com/pwvlygx> for a summary). Most schools classify hosts by the owner’s *role* (*e.g.*, faculty, student, staff, visitors), *department*, *residence* (*e.g.*, a particular dormitory), and *usage* (*e.g.*, research vs. education). In addition, many schools associate each host with a *security level* (with around ten different integer values) and whether the host is currently viewed as *compromised* (with a “yes” or “no” value). Some schools also classify hosts by *bandwidth quota* and *past usage*, to inform rate-limiting policies, and by whether they offer *core services* (*e.g.*, email and web servers). Based on these documents, and our discussions with the administrators of the computer-science department’s network of one university (University A), we learned about the following example policies.

**Security:** Schools use the security level to limit which external users can access a given host (and in what way). For example, hosts at the lowest security level might be blocked from receiving unsolicited traffic from external hosts; that is, these hosts cannot run public services. Other security levels correspond to different restrictions on which transport port numbers are allowed (*e.g.*, port 80 for HTTP, but not port 22 for SSH or 109, 110, and 195 for POP3). Some schools allow individual departments to state their own access-control lists, applicable only to hosts with IP addresses in that department’s address block. When administrators identify an internal host as compromised, they change the *compromised* attribute and significantly restrict the host’s access to network services. In addition, users in the *visitor* category typically have access to a limited set of services on the campus (*e.g.*, no access to the printers or campus email servers and compute clusters). One school restricts access to compute clusters in dormitories to the students residing in that particular dorm.

**Quality of Service:** Some universities impose a different default bandwidth quota based on the host’s *role*, but allow students and postdocs to purchase a higher quota. Some universities employ rate-limiting policies that depend on the user’s bandwidth usage on previous days (*e.g.*, users whose bandwidth usage exceeded a certain level were rate-limited to a lower level). Hosts offering core services are excluded from bandwidth usage calculations for both the users responsible for the service machines and the owners of the access machines, to avoid that traffic counting against their usage caps. Also, some schools offer higher quality-of-service for hosts assigned for educational use (*e.g.*, for streaming high-quality media in a classroom). The administrators of University A also expressed a desire to perform server load balancing for internal Web services based on role, to prevent heavy load from one group of users from compromising the performance of other users.

**Administrator “wish-lists”:** Our discussions with the administrators of University A also indicated that there were

many additional policies that they would ideally like to implement in the network, but did not do so since they were hard to realize in practice. University A assigns hosts to VLANs based on role (*e.g.*, faculty, staff, and student), for traffic isolation, to prevent packet sniffing and excessive broadcast traffic. The administrators would like to apply access-control policies based on device usage, device ownership, and OS, but do not do so today, since this would require exhaustive enumeration of IP addresses in the switch configuration. Likewise, the administrators expressed a desire to apply flexible QoS policies based on (i) the way a device is used (*e.g.*, research vs. infrastructure machines) and (ii) whether the host is owned by the department (unlike a “Bring Your Own Device” host).

Our discussions also revealed additional challenges with federated network management. The campus network assigns IP addresses in blocks based on location (*e.g.*, building). This raises challenges in applying security policies that restrict access to users affiliated with computer science department. Currently, the policy works correctly for hosts that are physically in the CS building, since these hosts are assigned a prefix from the CS subnet. However, when a CS user works in another building (common for faculty with dual appointments in other departments), the host receives a different IP address outside the CS subnet and the user is no longer able to access the CS resources. While the administrators could conceivably update network configurations dynamically to reflect the IP addresses that should have access, the management complexity is a deterrent. More generally, federated management would be much easier if network administrators had concise ways to represent security and QoS policies based on host attributes.

### B. Potential for Concise Rules with ACIPs

Existing techniques for assigning IP addresses to hosts can lead to a large numbers of rules in the switches. To quantify this problem, we analyzed the access-control policies in router configuration files for two university networks (University B and University C). Prior work shows that hosts in a network may be partitioned into a small number of policy units [28]—*i.e.*, a set of hosts that have identical reachability policies in terms of their communication with the rest of the network. Though the number of policy units is small, the number of ACL rules to express policy could still be as large as the square of the number of policy units. Thus, we go beyond [28], and not only identify policy units, but also calculate the number of rules required if ACLs were written in terms of policy units rather than the existing IP assignment.

Specifically, we consider two hosts as belonging to the same source policy unit (SPU) if and only if packets sent by these hosts to all destinations are treated identically in every ACL across all routers. Likewise, we consider two hosts as belonging to the same destination policy unit (DPU) if and only if packets from all sources to these hosts are treated identically in every ACL across all routers. We then compute the total number of rules needed to represent each ACL if it were more compactly expressed in terms of its source and destination policy units. Our results show that the number of ACL rules required is much smaller than the product of the

Dimensions	#Attributes	Example Attributes
Role	8	Faculty, Student
Security Level	16	1, 2, ..., 16
Status	6	In service, In testing
Location	7	–
Usage	3	Research, Infrastructure, ...
CS_owned	2	Yes, No
OS	5	MacOS, Windows, Linux, ...

TABLE I  
HOST DATA FOR CS DEPARTMENT (UNIVERSITY A)

source and destination policy units, and indicates that a smart ACIP allocation, which classifies hosts into their SPUs and DPUs efficiently, can potentially offer significant reduction for ACLs of the two universities (Section VIII).

### C. Diverse Attributes and Group Sizes

To better understand the attributes of hosts, we collected data about the 1491 registered hosts of the CS department of one university (University A). Each host is associated with seven dimensions of information, as summarized in Table I. In this network, (i) hosts are assigned to separate VLANs based on role and (ii) role, security level, and status are considered in access-control policy.

Given the number of attributes in each dimension, hosts could theoretically have 161,280 (*i.e.*,  $8 \times 16 \times 6 \times 7 \times 3 \times 2 \times 5$ ) combinations of attributes. In practice, only 287 unique combinations exist; for example, no visitor has a CS\_owned host. In addition, some combinations are much more popular than others. One group of hosts—belonging to one Linux-based compute cluster—has 109 members (more than 7% of all hosts). The large number of attributes and the diversity of group sizes have important implications for address assignment in Alpaca, which encodes host attributes in the IP address to enable more compact representations of policies.

Consider a naive address allocation scheme that performs *BitSegmentation*, by (1) concatenating a binary encoding of the host attributes along each dimension, where dimension  $i$  with a set of attributes  $D_i$  requires  $\lceil \log \|D_i\| \rceil$  bits, and (2) using the remaining bits to distinguish hosts with the same attributes along all dimensions, requiring  $\lceil \log X \rceil$  bits, where  $X$  is the size of the largest group.

The resulting encoding would enable very compact rules in the switches, using wildcard patterns to match on any attribute. However, this solution is impractical, even for this small network. Representing the seven dimensions would require 19 bits, and representing the largest group (with 109 members) would require 7 bits, for a total of 26 bits—a highly inefficient allocation of IP address space.

## IV. ALPACA OVERVIEW

### A. ACIP allocation with Alpaca

We present Alpaca, a system that embeds host attributes in IP addresses to enable compact policies. Figure 1 shows an overview of the system. Alpaca takes as inputs from the network operator the set of policy dimensions, and a database that lists the attributes associated with each host. It instructs

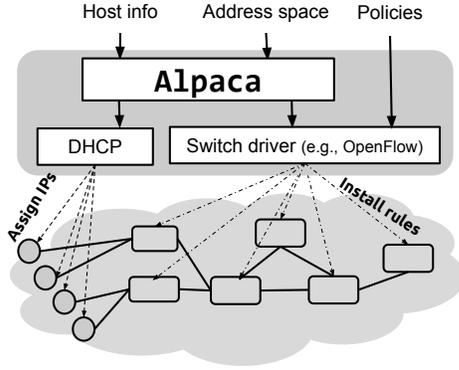


Fig. 1. Use Alpaca in a network.

DHCP servers how to assign IP addresses to hosts based on the results of Alpaca algorithm. If a host needs to be assigned a new address (e.g., it moves to a new location), the original attributes along with the new location are used in determining the new IP address. Meanwhile, the switch driver instantiates the network policies, which are defined on attributes, by installing match-action processing rules on switches. Alpaca coordinates with the switch driver such that rules can correctly classify the IP addresses of hosts to the corresponding attributes along different dimensions.

While ACIP allocation enables operators to express policies defined on multiple dimensions with switch rules, Alpaca must adapt to the following constraints:

**Switch rule-table sizes:** Switches impose hard constraints on the maximum number of rules to be installed. Therefore, Alpaca should optimize the classification of hosts for different switch architectures so that the attribute-based policies can be compactly expressed with switch rules that stay within the rule-table sizes.

**Address space:** IPv4 is widely deployed and likely to remain for the foreseeable future. Many enterprises have 16 or fewer bits for public IPv4 address space, or up to 24 bits for private IPv4 address space (i.e., 10.0.0.0/8). The limited address space calls for an efficient ACIP allocation that encodes attributes without wasting IPs. Though the address space constraint may be relaxed if IPv6 is fully deployed in a network, the allocation still needs to take the size of address space into account to correctly represent attributes.

**Dynamics:** Attributes of a host along any dimension may change. Alpaca must be able to handle changes in the attributes of a host while ensuring that only the IP address of that host changes, and that IP addresses of other hosts are not impacted. Alpaca should also handle (1) addition or removal of new attributes in existing dimensions without impacting the existing IP address allocation, and (2) addition or removal of dimensions, which is however relatively rare and may require significant changes.

## B. Problem Formulation

Given an IP address space of  $W$  bits, a set  $U$  of  $N$  hosts ( $N \leq 2^W$ ), a set of  $M$  dimensions and the attributes for the hosts along each dimension, an Alpaca algorithm computes an assignment of IPs to individual hosts and  $M$  sets of classification rules. Consider a multi-table switch architecture

where each rule-set corresponds to a dimension and appears in a dedicated table. A classification rule consists of an address pattern  $p$  and an attribute  $a$ ; rule  $(p, a)$  means that any host with IP matching  $p$  has attribute  $a$ .

Our optimization metric is the total size of  $M$  rule-sets, i.e., the number of classification rules that decide the attributes of all hosts along all dimensions in a multi-table switch architecture. We would like to find an IP address assignment that minimizes this total number of required rules.

An example is shown in Table II for the case of  $N = 16$  hosts with  $M = 2$  dimensions of Department and Role. Their values for each of the hosts are described in (a). Two IP address assignments  $T_1, T_2$  for each of the hosts are shown in (b). In (c), we can see encodings of the  $M$  dimensions based on the address assignments. With assignment  $T_1$ , 7 rules are required for the Department and 2 for the Role and the total cost is 9 rules. Alternatively, with assignment  $T_2$ , 5 and 8 rules are required for the two dimensions with a larger cost of 13.

Given an assignment, the number of required rules depends on the encoding technique. We mainly focus (Section V) on a simple technique named binary prefix [29], where prefix rules with disjoint matching sets of addresses are used to map addresses to attributes (as illustrated in Table II(c)). Due to its simplicity, this encoding technique can be implemented in simple memories or in simplified TCAMs without the expensive priority encoder component.

We generalize our solution for a more general encoding technique with wildcard matching rules that can be implemented in conventional TCAMs (Section VI).

## C. Overview of Alpaca algorithms

Alpaca consists of a series of algorithms targeted at different scenarios. The algorithms are described in Sections V-VII.

**The Prefix** algorithm computes prefix classification rules with a proven approximation ratio to the optimal case. It uses address space efficiently, requiring exactly  $\lceil \log_2 N \rceil$  bits. It is specially designed for multi-table switches with SRAMs that support a large number of prefix rules.

**The Wildcard** algorithm computes wildcard classification rules. It uses a small address space and can be applied on multi-table switch architectures.

Both prefix and wildcard algorithms by themselves do not handle dynamics and host attribute changes.

**The Slack** algorithm refines the prefix and wildcard algorithms by taking advantage of one more bit in the address space for an allocation that works well under dynamics in host attributes.

## V. ALPACA PREFIX ALGORITHM

In the following we describe Alpaca algorithms for assigning IPs to individual hosts. The first algorithm is designed for switch chipsets that allow prefix rules.

This section presents an address allocation algorithm that optimizes the number of prefix rules to represent attributes along multiple dimensions. It targets at multi-table switch architectures with IP prefix matching tables. We first introduce

(a) Input dimensions Dept, Role (b) Assignment Examples

Hosts	Dept	Role	address	$T_1$	$T_2$
$h_1 - h_5$	CS	Faculty	0000	$h_1$	$h_1$
$h_6 - h_7$	CS	Student	0001	$h_2$	$h_2$
$h_8 - h_{10}$	EE	Faculty	0010	$h_3$	$h_3$
$h_{11} - h_{16}$	EE	Student	0011	$h_4$	$h_4$
			0100	$h_8$	$h_5$
			0101	$h_9$	$h_6$
			0110	$h_{10}$	$h_7$
			0111	$h_5$	$h_8$
			1000	$h_{11}$	$h_9$
			1001	$h_{12}$	$h_{10}$
			1010	$h_6$	$h_{11}$
			1011	$h_7$	$h_{12}$
			1100	$h_{13}$	$h_{13}$
			1101	$h_{14}$	$h_{14}$
			1110	$h_{15}$	$h_{15}$
			1111	$h_{16}$	$h_{16}$

(c) Dimension representation based on the assignments

$T_1$		$T_2$	
Dept	Role	Dept	Role
00** → CS	0*** → Faculty	00** → CS	00** → Faculty
010* → EE	1*** → Student	010* → CS	0100 → Faculty
0110 → EE		0110 → CS	0101 → Student
0111 → CS		0111 → EE	0110 → Student
100* → EE		1*** → EE	0111 → Faculty
101* → CS			100* → Faculty
11** → EE			101* → Student
			11** → Student
7 + 2 = 9 rules		5 + 8 = 13 rules	

TABLE II

ALPACA ILLUSTRATION. A COMMON IP ASSIGNMENT IS USED TO REPRESENT ALL DIMENSIONS. THE COST OF AN ASSIGNMENT IS GIVEN BY THE SUM OF ITS COSTS FOR THE DIMENSIONS, I.E., THEIR TOTAL NUMBER OF REQUIRED RULES.

the notation, then discuss the optimal solution for a single dimension and the generalization to multiple dimensions.

We use the following notations when illustrating the algorithms. Let  $\alpha$  be a dimension and  $A = \{a_1, a_2, \dots\}$  be the set of associated attributes. For instance in the example from Table II the dimension Dept has the set of attributes {CS, EE} and the dimension Role has the set {Faculty, Student}. We view  $\alpha$  as a function that maps every host to an attribute, i.e.,  $\alpha(x) \in A$  is the attribute of host  $x$ . Let  $T$  be an ACIP allocation. We use  $C_\alpha(T)$  to denote the minimum number of rules to represent dimension  $\alpha$ . Likewise, for a set of dimensions  $\mathcal{D} = \{\alpha, \beta, \dots\}$ ,  $C_{\mathcal{D}}(T)$  represents the total number of rules to present all the dimensions in  $\mathcal{D}$  using the allocation  $T$ , i.e.,  $C_{\mathcal{D}}(T) = \sum_{\phi \in \mathcal{D}} C_\phi(T)$ . We define  $opt_\alpha = \min_T C_\alpha(T)$  and  $opt_{\mathcal{D}} = \min_T C_{\mathcal{D}}(T)$  to be the minimum number of rules to represent dimension  $\alpha$  and the set of dimensions, respectively.

#### A. A single dimension

We start with the simplest case: assigning addresses to represent exactly *one dimension*. Consider the dimension *Role*: each attribute of Role, such as Faculty or Student, should have its own set of rules for the hosts. As a prefix pattern matches a power-of-two number of hosts (e.g., 0\*\*\* stands for 8 hosts and 111\* stands for 2 hosts), one attribute might need several rules. The rules of different attributes do not overlap, i.e., matches are disjoint. Below, we describe a simple algorithm that finds an optimal address allocation to represent

#### Algorithm 1 Optimal algorithm for a single dimension $\alpha$

---

```

for attribute  $a_i \in A$  do
   $U_i = \{x \in U, \text{ where } \alpha(x) = a_i\}$ 
   $n_i = |U_i|$ 
Start from address  $0 \dots 0$ ,  $R = \emptyset$ 
for  $j = W$  to 0 do
  for attribute  $a_i \in A$  do
    if  $n_i \geq 2^j$  then
      Allocate in  $T$ ,  $2^j$  next addresses to hosts from  $U_i$ 
      Remove these  $2^j$  hosts from  $U_i$ 
      Add to  $R$  a corresponding rule:
        with  $j$  *s and attribute  $a_i$ 
       $n_i = n_i - 2^j$ 
return address allocation of hosts  $T$ , set of rules  $R$ 

```

---

one dimension.

Given the dimension function  $\alpha : U \rightarrow A$ , the algorithm returns the address allocation function  $T : U \rightarrow \{0, 1\}^W$ . The core idea is to treat the number of hosts for each attribute as the sum of power-of-twos and use a prefix rule for each power-of-two. Specifically, we first partition hosts into  $|A|$  sets based on their attributes. Let  $n_i$  be the number of hosts with attribute  $a_i$  ( $i = 1, \dots, |A|$ ) and  $bin(n_i)$  be the binary representation of  $n_i$ . We represent  $n_i$  as the sum of *distinct* power-of-twos based on  $bin(n_i)$ . For example,  $bin(14) = b1110$  and 14 is represented as  $8 + 4 + 2$ . For each attribute  $a_i$ , we further partition the set of hosts into subsets according to the representation of  $n_i$ . For example, if  $n_i = 14$ , we partition the set of hosts into 3 subsets with size 2, 4 and 8 respectively. Subsets for different attributes are considered in a descending size. Hosts are allocated addresses based on the index of the subset they belong to among the order subsets. The pseudo-code of the algorithm is shown in Algorithm 1.

Let  $\|bin(n_i)\|$  be the number of 1s in the binary representation of  $n_i$ , e.g.,  $\|bin(14)\| = 3$ . The above algorithm constructs  $\sum_{i=1}^{|A|} \|bin(n_i)\|$  subsets. We show that the resulting address allocation needs exactly  $\sum_{i=1}^{|A|} \|bin(n_i)\|$  rules for  $\alpha$ . In other words, each subset takes a single rule to represent. To prove it, we consider a subset of size  $2^j$ . Since subsets are sorted in a descending size, any previous subset must have a size of  $2^i$  for some  $i \geq j$ . Hence, the sum of the sizes of all previous subsets are multiples of  $2^j$ . This guarantees that all  $2^i$  hosts in the current set can be all represented by a single prefix rule with exactly  $j$  wildcards.

**Example 1.** We demonstrate the algorithm for the first dimension of Department for the hosts from Table II. We have seven hosts ( $h_1 - h_7$ ) with the attribute CS and nine hosts ( $h_8 - h_{16}$ ) with the attributes EE. They satisfy  $n_1 = 7 = 4 + 2 + 1 = 2^2 + 2^1 + 2^0$  and  $n_2 = 9 = 8 + 1 = 2^3 + 2^0$ . We consider the powers of two  $2^3, 2^2, 2^1, 2^0$  in this descending order. We allocate the first  $8 = 2^3$  addresses 0000 – 0111 to the first hosts with EE  $h_8 - h_{15}$ , adding the rule  $0*** \rightarrow EE$ . Then, we allocate the next  $4 = 2^2$  addresses 1000 – 1011 to the first hosts with attribute CS  $h_1 - h_4$ , adding the rule  $10** \rightarrow CS$ . Next, we allocate the following  $2 = 2^1$  addresses 1100 – 1101

to the hosts  $h_5, h_6$  with attribute CS. The corresponding rule is  $110^* \rightarrow CS$ . Next we allocate a single address  $1110$  to  $h_7$  and  $1111$  to  $h_{16}$ , one for each of the attributes with corresponding rules  $1110 \rightarrow CS$ ,  $1111 \rightarrow EE$ . There are a total of  $\|bin(7)\| + \|bin(9)\| = 3 + 2 = 5$  rules.

To prove the algorithm optimality, we show the following.

**Property 1.** For a dimension  $\alpha$ , let  $n_i = |\{x \in U \mid \alpha(x) = a_i\}|$  be the number of hosts that have to be mapped to an attribute  $a_i$ ,  $i = 1, 2, \dots, |A|$ . The minimal number of rules that can represent  $\alpha$  in any ACIP allocation satisfies  $opt_\alpha = \min_T C_\alpha(T) = \sum_{i=1}^{|A|} \|bin(n_i)\|$ .

*Proof.* An integer  $n_i$  cannot be described as a sum of less than  $\|bin(n_i)\|$  powers of two. Due to their disjointness, the matching rules for an attribute match a total number of hosts  $n_i$  that equals the sum of their matching sizes. Accordingly, an attribute with  $n_i$  hosts requires at least  $\|bin(n_i)\|$  rules.  $\square$

### B. Two dimensions

We now consider the case of two dimensions. Let  $\alpha : U \rightarrow A, \beta : U \rightarrow B$  be the dimensions under consideration, where  $B = \{b_1, \dots, b_{|B|}\}$  is the set of attributes in the second dimension. We observe a clear tradeoff between shortening the representation of these two dimensions. While we could choose the address allocation to be for instance the optimal for  $\alpha$  and use the minimal number of rules for  $\alpha$ , we may have to use many more rules to represent  $\beta$ . Since the address allocation is shared by both dimensions, in most cases we cannot find an allocation that favors both dimensions. This is summarized in the following property describing the relationship between the optimal allocation for two dimensions and the optimal allocation for each dimension. We recall that the optimal allocation minimizes the sum of the number of rules to represent each dimension.

**Property 2.** The optimal allocation for the dimensions  $\alpha, \beta$  satisfies  $opt_{\alpha,\beta} \geq opt_\alpha + opt_\beta$ . An equality  $opt_{\alpha,\beta} = opt_\alpha + opt_\beta$  is achieved if there exists an allocation that is optimal for the dimension  $\alpha$  as well as for the dimension  $\beta$ .

*Proof.* Consider an optimal representation of  $\alpha, \beta$  with  $opt_{\alpha,\beta}$  rules. It is composed of legal representations for each of the dimensions with the same allocation. Thus it has at least  $opt_\alpha$  rules for the first dimension and at least  $opt_\beta$  for the second and a total of at least  $opt_\alpha + opt_\beta$  rules.  $\square$

To obtain an upper bound on the optimal number of rules, we construct a special allocation below. Let  $\gamma$  be a new (dummy) dimension, which is the product of  $\alpha$  and  $\beta$ . The corresponding set of attributes is  $C = \{c_1, \dots, c_{|A| \cdot |B|}\}$ . For a host  $x \in U$ , if  $\alpha(x) = a_i, \beta(x) = b_j$  then  $\gamma(x) = c_{(i-1) \cdot |B| + j}$ . The dimension  $\gamma$ , denoted by  $\gamma = \alpha \times \beta$ , has the property that  $\gamma(x)$  determines the attributes  $\alpha(x), \beta(x)$  for the same host  $x$ . Considering the optimal allocation for  $\gamma$ , we deduce:

**Property 3.** The cost of an optimal mapping for the dimensions  $\alpha, \beta$  and for  $\gamma$  satisfy  $opt_{\alpha,\beta} \leq 2 \cdot opt_\gamma$ .

*Proof.* Consider the representation of  $\gamma$  under some allocation  $T$  with  $C_\gamma(T)$  rules. We can obtain a representation of the

---

### Algorithm 2 Algorithm for two dimensions $\alpha, \beta$

---

Define a dimension  $\gamma$  with attributes  $C = \{c_1, \dots, c_{|A| \cdot |B|}\}$ :

$\gamma(x) = c_{(i-1) \cdot |B| + j}$  if  $\alpha(x) = a_i, \beta(x) = b_j$

Find an optimal address allocation  $T$  for  $\gamma$  by Algorithm 1

Let  $R$  be the corresponding set of rules

For  $\alpha$ , define  $R_\alpha$  based on  $R$ :

replace an attribute  $c_{(i-1) \cdot |B| + j}$  of a rule by  $a_i$

For  $\beta$ , define  $R_\beta$  based on  $R$ :

replace an attribute  $c_{(i-1) \cdot |B| + j}$  of a rule by  $b_j$

**return** address allocation of hosts  $T$ , sets of rules  $R_\alpha, R_\beta$

---

dimension  $\alpha$  (or of  $\beta$ ) with the same number of rules by only modifying the attribute of each rule, i.e., replacing attribute in  $C$  with the corresponding attribute in  $A$  (respectively in  $B$ ). Therefore,  $C_\alpha(T) \leq C_\gamma(T)$ ,  $C_\beta(T) \leq C_\gamma(T)$ . Finally,

$$C_\alpha(T) + C_\beta(T) = C_{\alpha,\beta}(T) \leq 2C_\gamma(T). \quad (1)$$

$\square$

**Example 2.** We demonstrate the algorithm for the two dimensions *Department, Role* from Table II. Denote the attributes of the two dimensions as  $\{CS, EE\} = \{a_1, a_2\}$  and  $\{Faculty, Student\} = \{b_1, b_2\}$ . Accordingly we define the dimension  $\gamma$  with attributes  $C = \{c_1, \dots, c_4\}$  as  $\gamma(h_i) = c_1$  for  $i \in [1, 5]$ ,  $\gamma(h_i) = c_2$  for  $i \in [6, 7]$ ,  $\gamma(h_i) = c_3$  for  $i \in [8, 10]$  and  $\gamma(h_i) = c_4$  for  $i \in [11, 16]$ . With regards to  $\gamma$ ,  $n_1 = 5, n_2 = 2, n_3 = 3, n_4 = 6$ . We apply Algorithm 1 for  $\gamma$ . This results in a set of rules  $R$  of  $2+1+2+2 = 7$  rules ( $00^{**} \rightarrow c_1$ ,  $01^{**} \rightarrow c_4$ ,  $100^* \rightarrow c_2$ ,  $101^* \rightarrow c_3$ ,  $110^* \rightarrow c_4$ ,  $1110 \rightarrow c_1$ ,  $1111 \rightarrow c_3$ ) based on some address allocation  $T$ . For this allocation, we can derive a representation for *Department* with seven rules by replacing  $c_1, c_2$  by CS and  $c_3, c_4$  by EE. We can also build a representation for *Role* with the same number of rules by replacing  $c_1, c_3$  by Faculty and  $c_2, c_4$  by Student. This results in a representation for the two dimensions with 14 rules.

We remark here that the representations obtained from that of  $\gamma$ , are not necessarily the minimum representations of  $\alpha$  or  $\beta$  with the address allocation  $T$ : we can sometimes further compress rules for each dimension. For instance, for *Role* we derived the seven rules ( $00^{**} \rightarrow Faculty$ ,  $01^{**} \rightarrow Student$ ,  $100^* \rightarrow Student$ ,  $101^* \rightarrow Faculty$ ,  $110^* \rightarrow Student$ ,  $1110 \rightarrow Faculty$ ,  $1111 \rightarrow Faculty$ ). Here, it is possible merge the last two of them to a single rule  $111^* \rightarrow Faculty$  to achieve a representation for this dimension with only six rules and with  $7 + 6 = 13$  rules for the two dimensions.

Next we show that an optimal allocation for  $\gamma$  is a 2-approximation of the optimal allocation for  $\alpha, \beta$ , i.e., the number of rules it generates is at most twice of the optimum.

**Property 4.** The costs of optimal mappings for the dimensions  $\alpha, \beta$  and for  $\gamma$  satisfy  $opt_\gamma \leq opt_{\alpha,\beta}$ .

*Proof.* Consider some allocation  $T$ , it must satisfy

$$C_\gamma(T) \leq C_\alpha(T) + C_\beta(T) = C_{\alpha,\beta}(T). \quad (2)$$

This is because a group of hosts that cannot be represented using a single rule in  $\gamma$  must have different attributes in at

least one of  $\alpha$  and  $\beta$ , thus requiring a minimum of two rules to represent these hosts in that dimension. In particular, considering the optimal allocation  $T_{\alpha,\beta}$  for  $\alpha, \beta$  we have  $C_\gamma(T_{\alpha,\beta}) \leq C_{\alpha,\beta}(T_{\alpha,\beta}) = \text{opt}_{\alpha,\beta}$  and the claim follows.  $\square$

Finally, we conclude the following property:

**Property 5.** *Let  $T_\gamma$  be an optimal allocation for  $\gamma = \alpha \times \beta$ . Then,  $C_{\alpha,\beta}(T_\gamma) \leq 2 \cdot \text{opt}_{\alpha,\beta}$ , i.e., this allocation is a 2-approximation for the two dimensions  $\alpha, \beta$ .*

*Proof.* Let  $T_\gamma$  be the optimal solution of  $\gamma$  and  $T_{\alpha,\beta}$  be the optimal solution for  $\alpha, \beta$  together. Substituting  $T$  with  $T_{\alpha,\beta}, T_\gamma$  in Equation 1 and 2, we obtain  $C_{\alpha,\beta}(T_\gamma) \leq 2C_\gamma(T_\gamma)$  and  $2C_\gamma(T_{\alpha,\beta}) \leq 2C_{\alpha,\beta}(T_{\alpha,\beta})$ . Since  $C_\gamma(T_\gamma) \leq C_\gamma(T_{\alpha,\beta})$  and  $C_{\alpha,\beta}(T_{\alpha,\beta}) = \text{opt}_{\alpha,\beta}$  we deduce the result.  $\square$

To summarize, for two given dimensions  $\alpha, \beta$ , we calculate  $\gamma = \alpha \times \beta$  and find its optimal allocation  $T_\gamma$  by the algorithm for a single dimension. We then use this allocation to represent each of the dimensions  $\alpha, \beta$ . A pseudo-code of the algorithm is given in Algorithm 2.

### C. General number of dimensions

We can generalize the solution for two general dimensions from Section V-B to handle a set  $\mathcal{D}$  of an arbitrary number of dimensions  $M = |\mathcal{D}|$ . Similar to computing  $\gamma$  for the two dimensions, we introduce a dimension  $\Pi_{\mathcal{D}}$ , whose attributes for a host  $x \in U$  is a vector of length  $|\mathcal{D}|$  with the attributes of all dimensions in  $\mathcal{D}$  for that host. The number of attributes in  $\Pi_{\mathcal{D}}$  equals the product of the attributes number in the dimensions of  $\mathcal{D}$ . We find an optimal assignment for this new dimension  $\Pi_{\mathcal{D}}$  based on Algorithm 1.

We derive that the optimal allocation for  $\Pi_{\mathcal{D}}$  is an  $M$ -approximation to the optimal allocation for  $\mathcal{D}$ . As in the case of two dimensions, we can obtain a representation for each of the  $M$  dimensions from the representation of  $\Pi_{\mathcal{D}}$ . To do so, we use the same set of rules while replacing an attribute of  $\Pi_{\mathcal{D}}$  by the single attribute it is associated with in each of the  $|\mathcal{D}|$  dimensions. The guaranteed approximation ratio on the memory size holds since the cost of an optimal allocation for  $\mathcal{D}$  is at least the cost of an optimal allocation for  $\Pi_{\mathcal{D}}$ .

### D. Exploiting Redundant Addresses

The analysis so far assumed that the number of available addresses equals the number of hosts. We can further reduce the representation cost when the number of available addresses is larger than the number of hosts. In practice, we expect some slack in the number of addresses, to handle dynamics in the number of hosts, and changes in host attributes. We demonstrate that for a single dimension. A similar approach can lead to cost reduction also for multiple dimensions.

Consider a dimension  $\alpha$  (with a set of attributes  $A$ ) defined on a set of hosts  $U$ . Addresses of  $W$  bits should be allocated to the set of hosts satisfying  $|U| < 2^W$ . Let again  $n_i$  be the number of hosts with an attribute  $a_i$ ,  $i \in [1, |A|]$ . In an address allocation, there are  $2^W - |U|$  addresses that are not allocated to any of the hosts. For these addresses, the matching rule can be associated with an arbitrary attribute. Furthermore, there is

no requirement for the existence of a matching rule for these addresses. Of course, for the correctness of a representation of a dimension, the matching rule for the address assigned to a host must have its required attribute.

Consider a representation of  $\alpha$ . Let  $h_i$  be the number of addresses that match rules with a value  $a_i$ . For the correctness of the representation necessarily  $h_i \geq n_i$ . Furthermore, since the total number of available addresses is  $2^W$  then necessarily  $\sum_{i \in [1, |A|]} h_i \leq 2^W$ . There exists such a representation with  $\sum_{i \in [1, |A|]} |h_i|_b$  rules.

Notice that  $\forall n_i \in \mathbb{N}$  the value  $h_i = 2^j$  for  $j = \lceil \log_2(n_i) \rceil$  satisfies  $|h_i|_b = 1$  and  $h_i < 2 \cdot n_i$ . This means that by allocating less than twice the number of addresses for some attribute, we can represent all hosts with this attribute by a single rule. Accordingly, if the number of hosts satisfies  $U \leq 2^{W-1}$ , then  $\sum_{i \in [1, |A|]} 2^{\lceil \log_2(n_i) \rceil} < \sum_{i \in [1, |A|]} 2 \cdot n_i = 2 \cdot |U| \leq 2^W$  and there exists a representation with a single rule for each policy value. In other words, we can always guarantee the existence of a representation with a single rule for each attribute by allocating addresses that are longer by a single bit than the minimal length required for the uniqueness of the addresses.

We generalize this analysis to describe a bound that applies also for smaller amount of redundancy (with  $U > 2^{W-1}$ ).

**Property 6.** *A dimension  $\alpha$ , with a set of attributes  $A$ , defined on a set of hosts  $U$  that satisfies  $|U| \cdot (1 + 2^{-k}) \leq 2^W$  can always be represented by at most  $(k + 1) \cdot |A|$  rules with the allocation of hosts to  $W$ -bit addresses.*

*Proof.* Consider a value  $n_i$  with  $j = \lceil \log_2(n_i) \rceil$ . For  $k \in [0, j - 1]$  there exists  $h_i \geq n_i$  that is a multiple of  $2^{j-k}$  satisfying  $h_i - n_i < 2^{j-k}$ . Such a value must satisfy  $|h_i|_b \leq j + 1 - (j - k) = k + 1$  since either  $h_i = 2^{j+1}$  or  $h_i \in [2^j, 2^{j+1} - 1]$  with a binary representation composed of  $(j - k)$  last bits of 0. Such  $h_i$  satisfies  $h_i/n_i < (1 + 2^{j-k}/2^j) = (1 + 2^{-k})$ . If the number of available addresses enables to consider for some fixed  $k$  the value of  $h_i$  for all attributes, all hosts with a common attribute can be represented by at most  $k + 1$  rules and the total number of required rules to represent  $\alpha$  is at most  $(k + 1) \cdot |A|$ .  $\square$

## VI. ALPACA WILDCARD ALGORITHM

In this section, we present an algorithm that generates wildcard rules by optimizing the output of the prefix solution. To illustrate the algorithm, we use the same example in Table II. For clarity, we use *host group* to refer the set of hosts with the same attributes along all dimensions; we use *rules* and *patterns* interchangeably to refer the compact ACIP representation of host groups and attributes.

We revisit our example. The prefix solution uses 1 pattern for the (CS, Student) group and 2 patterns for each of the other groups, because it views the size of a host group as the sum of powers-of-twos (e.g.,  $5 = 1 + 4$ ), each of which corresponds to a prefix pattern. Hence, it uses  $2 + 2 = 4$  prefix rules to represent Faculty attribute. But if we assign  $\{0111, 010*\}$  to EE Faculty and  $\{0110, 11**\}$  to CS Faculty, then we can compress these patterns to a single pattern  $*1**$  to represent Faculty attribute. The key observation is that *if two host groups share common attribute(s), it is beneficial to*

	CS	EE
Faculty	5	3
Student	2	6

(a) Group sizes

⇒

Id	(Dept, Role, Value)
1	(CS, F, 1)
2	(EE, F, 1)
3	(EE, S, 2)
4	(CS, S, 2)
5	(EE, F, 2)
6	(EE, S, 4)
7	(CS, F, 4)

(b) Initial nodes

Fig. 2. Create nodes from input of Table II(a).

assign them similar patterns that can be compressed to reduce the number of rules for the common attribute.

In the design of our algorithm we are inspired by ideas described in BitWeaving [14]. This novel approach was shown to achieve a significant reduction of semantically-equivalent representations of packet classifiers. It relies on two techniques named bit swapping and bit merging. In particular, using wildcard rules, in bit merging there is an effort to merge rules whose matching patterns differ by a single bit. Similarly, in our suggested graph approach, we try to merge two nodes describing host groups aligned on a large number of dimensions. Inherent differences between the two approaches exist: BitWeaving finds an efficient semantically-equivalent representation of a given single classifier. An input address allocation affects the starting point of that technique. On the other hand, Alpaca assumes multiple dimensions, each can be seen as a mapping from hosts to different attributes. This is equivalent to a simultaneous optimization of multiple classifiers through the selection of the address allocation. Unlike our approach that assumes a simple single-field address as the input to the classification, BitWeaving is more general in this aspect and supports multiple-field classification. A comparison with BitWeaving for the scope of this paper is performed in Section VIII.

We now describe the highlights of the wildcard algorithm.

**Potential compression.** Our first task is to find out all the potential compression of patterns among host groups. Starting with the output of prefix solution, which uses the sum of power-of-two terms to denote the size of a host group, we map every power-of-two term to a *node*. The node saves the value of the term and copies the attributes of the host group. For example, we can create two nodes for the CS Faculty group: (CS, Faculty, 1) and (CS, Faculty, 4), as the group size  $5 = 1 + 4$ . Figure 2(b) shows the full list of nodes.

Two nodes can be compressed if their values are equal and they share some common attribute(s). The result of compression is a new node that (1) has a value equal the sum of the values of the two nodes, (2) “inherits” the shared attributes and (3) has  $\emptyset$  attributes for other dimensions. For example, (CS, Faculty, 1) and (EE, Faculty, 1) can be compressed into a new node ( $\emptyset$ , Faculty, 2). We call the new node a *super-node* and the two original nodes *sub-nodes*. The compression suggests that we could use the super-node instead of listing two sub-nodes individually to represent their common attributes. In Figure 3 nodes appear in white background and super-nodes in gray.

A super-node can be compressed with other nodes, as long

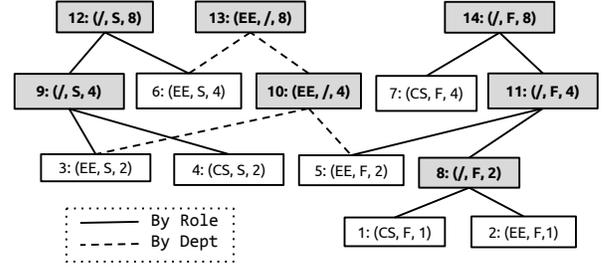


Fig. 3. The compression graph: a node has an id, attributes and a value. Colored nodes are super-nodes.

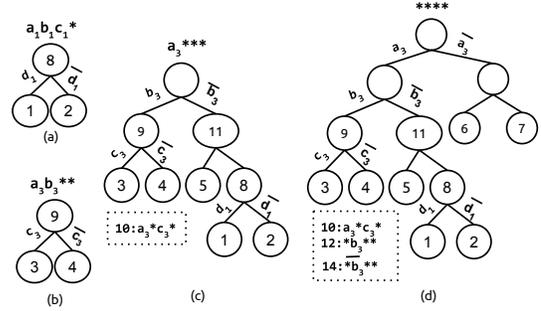


Fig. 4. Flip bits to compress nodes

as they share the same attributes (except  $\emptyset$ ). But a node cannot be compressed twice on the same dimension, *i.e.*, once (CS, Faculty, 1) and (EE, Faculty, 1) are compressed, neither of them can be compressed with other nodes that have Faculty attribute. We repeat the compression until no new super-nodes can be produced. We plot graphs to denote the compression relationship by creating edges from sub-nodes to their super-nodes.

In the graph, a node with value  $2^k$  can be assigned a wildcard pattern with exactly  $k$  wildcards, which represent  $2^k$  hosts. As we work on the output of the prefix solution, the initial nodes (Figure 2(b)) should be assigned prefix patterns.

**Compressible patterns.** Two patterns are compressible if they negate at exactly one bit, *e.g.*,  $*00*$  and  $*10*$  are compressed into  $**0*$  [14]. When two sub-nodes (of a super-node) are assigned compressible patterns (*e.g.*,  $*00*$  and  $*10*$ ), the resulting pattern (*e.g.*,  $**0*$ ) can be assigned to the super-node to achieve a reduction of one rule in representing the common attribute, where we can use the compressed pattern instead of listing two patterns independently. In the example, we can use the pattern for ( $\emptyset$ , Faculty, 2) to represent CS attribute rather than two patterns for (CS, Faculty, 1) and (EE, Faculty, 1). Therefore, our goal is to assign patterns to nodes to *maximize the number of pairs of sub-nodes with compressible patterns*, *i.e.*, the total reduction in the number of rules to represent attributes.

**Key idea: flip one bit.** Let  $a_i b_i c_i d_i$  be the pattern assigned to the  $i$ -th node, where  $a_i, b_i, c_i, d_i \in \{0, 1, *\}$ . Consider Node 1 (CS, Faculty, 1) and Node 2 (EE, Faculty, 1).  $a_1 b_1 c_1 d_1$  and  $a_2 b_2 c_2 d_2$  are compressible if they negate at one bit, *i.e.*,  $a_1 b_1 c_1 d_1 = \bar{a}_2 b_2 c_2 d_2$ , or  $a_2 b_2 c_2 d_2 = a_1 b_1 c_1 d_1$ , or  $a_2 b_2 \bar{c}_2 d_2 = a_1 b_1 c_1 d_1$ . Our key idea to enable compression is to *choose a bit* (*e.g.*,

Dept	
p	a
0110	CS
001*	CS
11**	CS
0111	EE
10**	EE
0*0*	EE

Role	
p	a
*1**	Faculty
*0**	Student

Fig. 5. Wildcard rule-sets.

$a, b, c$  or  $d$ ) to flip. If we flip  $d$ , the compressed pattern  $a_1b_1c_1*$  (or  $a_2b_2c_2*$ ) can be assigned to the super-node ( $\emptyset$ , Faculty, 2), *i.e.*, Node 8. As a result,  $a_8b_8c_8d_8 = a_1b_1c_1*$ . We plot the equality in Figure 4(a).

Similarly, we can compress patterns of Node 3 and Node 4 by flipping  $c$  ( $d_3 = d_4 = *$ ) as shown in Figure 4(b). We can further compress (1) Node 5 and Node 3 by flipping  $b$  (as  $c_3$  is flipped before) and (2) Node 5 and Node 8 by flipping  $c$  (shown in Figure 4(c)). However, we are unable to compress Node 6 ( $a_6b_6c_6d_6 = \bar{a}_3b_3**$ ) and Node 10 ( $a_{10}b_{10}c_{10}d_{10} = a_3*c_3*$ ), because their patterns do not match. We finish the procedure by compressing Node 6 and Node 9, Node 7 and Node 11 (Figure 4(d)). To translate the results to patterns, we set all variables, *i.e.*,  $a_3, b_3, c_3, d_1$ , to 0.

To summarize, with the idea of bit flipping, we construct *equality and inequality* between the bits of patterns, *i.e.*,  $a, b, c, d$ , assigned to nodes. For the patterns of each pair of sub-nodes, if the equality (or inequality) is not determined before, we choose the last possible bit to flip. When there is no such a bit (*i.e.*, all the patterns negating at one bit are already used), then we choose to flip more than one bit until the resulting patterns do not overlap with any used ones. After checking all pairs of sub-nodes, we obtain the full equality and inequality. The final step is to set all free bit variables to 0.

We can represent each attribute with rules given the pattern assignment (Figure 5). For example, to represent Student, we can use  $*0**$  for super-node ( $\emptyset$ , Student, 8), *i.e.*, Node 12. Similarly, we use  $10**$ ,  $0*0*$  and  $0111$  to represent EE. In total, we need  $3 + 3 + 1 + 1 = 8$  rules to represent all the attributes, whereas prefix solution needs 9 or more rules (Table II(c)).

**Processing order of sub-nodes.** The order we use to process sub-nodes matters, as the compression of one pair of sub-nodes may restrict the compression of another (due to the equality and inequality between bits). The algorithm calculates the *order values* for super-node  $n$  as the total number of super-nodes in the tree rooted at  $n$  in the graph. For example, the tree rooted at ( $\emptyset$ , S, 4) only contains one super-node (*i.e.*, itself); the tree rooted at ( $\emptyset$ , F, 8) contains three. Super-nodes are sorted according to their order values and examined one by one. When examining one super-node, we process all the pairs of sub-nodes in its tree. If the compression failed for one pair (*i.e.*, we cannot find a bit to flip), we roll back all previous compressions of sub-node pairs in the tree and continue to examine the next super-node in the sorted list; if compressions of all pairs of sub-nodes succeed, we remove these sub-nodes from the trees of other super-nodes, re-calculate order values of the affected super-nodes and sort again.

## VII. SUPPORTING UPDATES AND PRACTICAL ISSUES

In this section we discuss the overhead of the approach. We also explain how to easily handle changes in the association of hosts to attributes.

### A. Overhead and Practical Issues

Taking into account host attributes in the address allocation process has an inherent overhead. A change in a host attribute can lead to a severe increase in the memory required to represent policies given an existing address assignment, suitable for the previous attributes. Such changes can occur too often eliminating the possibility to reallocate addresses upon any such small change. More specifically, we would like to minimize the impact on any other hosts while sometimes due to a limited address space, changing the address of a specific host necessitates changing the address of some others. In particular, consider a policy that takes into account a location of a host as a dimension. Based on an estimated frequency for updates in a dimension (*e.g.*, a location is changed more often than a role of a host), a potential improvement for the allocation would be not to consider all dimensions identically. Similarly, the attributes along one dimension might not be fully available during the allocation process. This might lead to a non-optimal allocation. In this case, a potential generalization can be to find an allocation given a restricted possible subset of attributes for a host dimension, where the specific value might not be known. Next we discuss how to efficiently deal with updates in host attributes.

### B. Handle Changes in Host Attributes

Our algorithms proposed so far support address allocation given the attributes of each host. In practice, attributes of a host may change over time (*e.g.*, the department of the corresponding user might change), or new attributes may be added (*e.g.*, a new department may be created). In handling changes, a key consideration is ensuring that only the IP addresses of impacted hosts are modified to the extent possible.

We employ two techniques to handle changes in attributes. First, to handle growth in the number of hosts that have a certain attribute, we introduce *slack*, and budget for more hosts than actually exist. While earlier in Section V-D we took advantage of redundant addresses to improve representation cost, here the focus is to handle updates and the discussion is not limited to the prefix algorithm nor to a single dimension. A straight-forward solution is to provision for a growth in the number of hosts corresponding to a given attribute by a fixed percentage (*e.g.*, 10%), though information about projected trends could be used when available. For example, a university can estimate the number of hosts in the coming semester based on the number of newly admitted students.

Second, to handle growth in the number of attributes along each dimension, we introduce a “ghost” attribute for each dimension (an additional attributes with which no host is currently associated) and decide the group sizes for combinations of ghost and real attributes (*e.g.*, the number of hosts with ghost Department and Student, or the number of hosts with ghost Department and ghost Role).

	CS (9)	EE (8)	Ghost_Dept (6)
Faculty (6)	3	1	2
Student (11)	4	5	2
Ghost_Role (6)	2	2	2

↓

	CS (16)	EE (8)	Ghost_Dept (6)
Faculty (8)	5	1	2
Student (16)	9	5	2
Ghost_Role (6)	2	2	2

TABLE III  
AN EXAMPLE OF SLACK

Given the input with slacked group sizes and ghost attributes, Alpaca algorithms compute ACIP allocation. When the updates only occur for the existing attributes, we change the addresses of the affected hosts to unused ACIP from the patterns computed for their new attribute. In the case that the provisioned slack of a group is exhausted, we partition the address space of the associated ghost groups, whose attributes are either ghost attributes or attributes of the exhausted group, and allocate part of the space to the exhausted one. For example, if the ACIPs of (Student, CS) are used up, we could partition the address space of (Student, GhostDept), (GhostRole, CS) or (GhostRole, GhostDept) and assign new space to (Student, CS). When the updates involve a new attribute in one dimension, *e.g.*, Department, we run Alpaca algorithms on the address space for the ghost attribute to split the space into two parts: one for the new attribute and the other for the ghost attribute. Afterwards, the addresses of affected hosts are changed accordingly.

**Benefits of slack and ghost attributes.** The above two techniques offer another important advantage: *further compacting network policies beyond the optimal solution*. Consider an example where there are 7 CS hosts and 7 EE hosts. Alpaca needs at least 3 rules for each attribute, as  $7 = 4 + 2 + 1$  (Section V). With slack, we can round 7 to 8, thus allowing Alpaca to use only 1 rule per attribute. In fact, if we round the group size for every attribute to the nearest power-of-two upper bound, we at most double the used addresses.

We create extra hosts with “mix-matched” attributes such that the number of hosts for every attribute is power-of-two (Algorithm 3). Let  $p_a$  be the target power-of-two and  $g_a$  be the number of hosts for the attribute  $a$ . We choose attribute  $v_i$  from  $i$ -th dimension such that  $p_{v_i} > g_{v_i}, \forall i \in [1, M]$ , and create  $h = \min_i \{p_{v_i} - g_{v_i}\}$  hosts with attribute  $v_1, \dots, v_M$ . When all attributes in a dimension reach their target power-of-two (*i.e.*,  $p_a = g_a$ ), we use the ghost attribute as default, assuming its target power-of-two is infinite. We repeat the procedure until all attributes reach their target power-of-two (except ghost attributes). Consider the example in Table III. The numbers of hosts for CS, Faculty and Student should be rounded to 16, 8 and 16. We create 2 CS Faculty hosts in the first iteration and create 5 CS Student afterwards.

For more complex updates that involve additions of new dimensions, it may be desirable to recompute IP allocations from scratch. However, we make several points. First, such scenarios are relatively infrequent. We envision that Alpaca algorithms are run with a conservative set of dimensions, even

---

### Algorithm 3 Slack algorithm

---

```

while True do
  for Dimension  $d_i$  do
     $v_i =$  ghost attribute
    for Attribute  $a \in d_i$  do
      if  $p_a > g_a$  then
         $v_i = a$ 
    if  $\forall i, v_i$  is ghost attribute then
      break
     $h = \min p_{v_i} - g_{v_i}$ 
     $g_{v_i} = g_{v_i} + h, \forall i$ 
    Create  $h$  hosts with attribute  $v_1, \dots, v_M$ 
return
```

---

	#ACLs	Total #Rules	#SPU	#DPU
University B	13	17868	577	624
University C	5027	32401	523	87

TABLE IV  
NETWORK POLICIES OF TWO UNIVERSITIES.

if some of these dimensions are not currently used as part of network policy. Addition of new dimensions is likely to happen over long time-scales — operators typically collect host attribute information using device registration information filled by owners, and introducing new dimensions would require new data collection for registered devices. Second, when such scenarios do occur, it is feasible to temporarily deal with it by splitting the unused address space of other dimensions and introducing less compact classification rules to identify a given set of hosts. Finally, changes in address allocation can be incrementally handled using DHCP.

## VIII. EVALUATION

In this section, we evaluate Alpaca’s effectiveness in producing concise rules under two scenarios: (i) actual policies in existing networks and (ii) futuristic scenarios where operators may express policy along many orthogonal dimensions. For existing settings, we evaluate Alpaca using the network configuration files of University B and University C (Section VIII-A). For futuristic settings, we use the host attribute data obtained from University A (Section VIII-B). Details of both data-sets were presented in Section III.

Overall, our results show that Alpaca can reduce the number of rules by 60% – 68% as compared to the current IP address allocation for multi-table switches. We observe a reduction of 20% - 31% in comparison with schemes relying on existing techniques for rule table compression. Meanwhile, it has the potential to reduce the total number of rules by over an order of magnitude as compared to the traditional single dimensional approaches (*e.g.*, VLAN) in futuristic scenario where the policy is expressed on many dimensions. Our experiments further demonstrate that Alpaca can handle changes in hosts gracefully, with only a small extra number of rules.

Our evaluations explore the performance of both Alpaca variants: Prefix (ALP\_PFX) and Wildcard (ALP\_WC). For comparison purposes we also consider the simple BitSegmentation scheme (BitSeg) as well as schemes based on popular

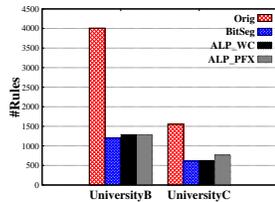


Fig. 6. Optimize network policies.

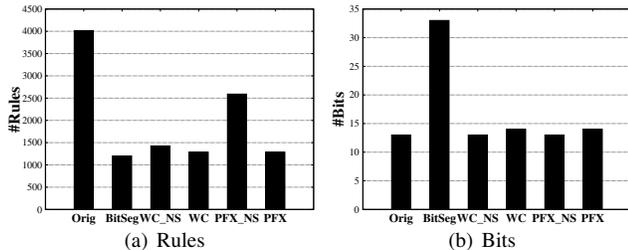


Fig. 7. Benefits of slack. BS, WC and PFX denote Wildcard, Prefix and BitSegmentation schemes. NS indicates variant without slack.

schemes for the compression of rule classifiers [12]–[14]. Unless otherwise mentioned, both our prefix and wildcard algorithms use the algorithm (with prefix extension) in Section VI and the extension with slack in Section VII-B by default. We focus on multi-table architectures.

#### A. Benefits with Existing Policies

We extract the source and destination policy units (SPUs and DPUs) from the low-level configuration files for both University B and University C, as discussed in Section III. Table IV shows the total number of ACL rule-sets, ACL rules across rule-sets, and the number of SPUs and DPUs for both universities. Given a pipeline of tables, we install the classification rules that associate a given IP with its appropriate SPU and DPU in the first two tables, and the actual policy action (*e.g.*, permit or deny) based on the SPUs and DPUs in the last table. We focus on the number of classification rules in the first two tables for a given policy, since the last table is the same in all approaches.

Figure 6 compares the number of rules used by (1) the original IP allocation (Orig), (2) BitSeg, (3) ALP\_PFX and (4) ALP\_WC for University B and University C, respectively. The original IP allocation needs the most rules. BitSeg takes the least, as it uses one rule for each policy unit. Specifically, the number of rules used by BitSeg equals the number of SPUs and DPUs. Both ALP\_PFX and ALP\_WC perform closely to BitSeg, achieving 68% reduction in rule consumption as compared to the original. It confirms that Alpaca can efficiently encode policy units.

**Benefits of slack.** We compare the case with and without slack operations to show the benefits of trading an extra bit for significant reduction in number of rules. Figure 7(a) presents the reduction in the number of rules for University B. We use NS to indicate running Alpaca without slack. While WC\_NS (3rd bar) is competitive with other approaches, PFX\_NS (5th

bar) performs slightly worse, giving a reduction of 35.4%. The reasons are two-fold: for one thing, prefix patterns fundamentally restrict the potential of using fewer rules (as compared to wildcard patterns); for the other, PFX\_NS solutions represent the exact group size of every combination (*i.e.*, each SPU and DPU pair) without any slack. If the group size is not power-of-two, PFX\_NS solutions need many more rules. We bridge the gap by adding slack and rounding group sizes. As a result, PFX offers similar performance to the optimal (*i.e.*, BitSeg). In the remainder of the evaluation, we run slack algorithm before IP allocation by default.

Moving to the number of bits for encoding (Figure 7(b)), we find BitSeg performs the worst, using as many as 34 bits (more than IPv4!). We calculate the least number of bits that can sufficiently number all hosts (13 for University B). While the original allocation uses the least bits, PFX\_NS takes exactly one bit more as the slack algorithm makes use of an extra bit to round group sizes (Section VII-B). Alpaca strikes a balance in both the number of rules and bits, using almost as few rules as BitSeg and one extra bit than the least number of bits.

#### B. Benefits with Futuristic Policies

We demonstrate Alpaca’s capability to support flexible attribute-based policies with a series of experiments on the host information at the CS department of University A (Table I). In the current CS network, operators deploy VLANs to group hosts with the same Role, which is used in most network policies. But they would like to use Security Level, Status and Operating System for access control and have flexible QoS policies defined on Usage, CS\_owned as well. Hence, we examine the cost in terms of rules and address space to support the futuristic scenarios, where policies are defined on attributes along multiple dimensions.

We compare Alpaca with three approaches:

(1) *SingleDim* (*e.g.*, VLAN), which assigns addresses based on a single dimension. SingleDim uses a few rules to represent attributes for one dimension: VLAN uses one rule (the subnet) for each Role attribute; a host is assigned a random address in the subnet corresponding to its Role attribute. However, given a second dimension or more, SingleDim has to enumerate every single host and list their attributes.

(2) *TCAMRazor* [13], which applies an algorithm to minimize the number of prefix rules for attributes, given the SingleDim address assignment.

(3) *BitWeaving* [14], which uses an efficient approach to compute the wildcard rules to represent attributes based on the SingleDim address assignment, as minimizing wildcard rules is NP-hard [17].

We remark that TCAMRazor and BitWeaving minimize the number of rules by assuming rule priority. Both methods generate overlapping rules for different attributes. In contrast, Alpaca generates non-overlapping rules for different attributes, *i.e.*, does not apply rule priority. Below, we show that even without using rule priority, Alpaca significantly outperforms the two compression methods.

**Scale with more dimensions.** We evaluate Alpaca’s encoding efficiency and scalability with increasing number of

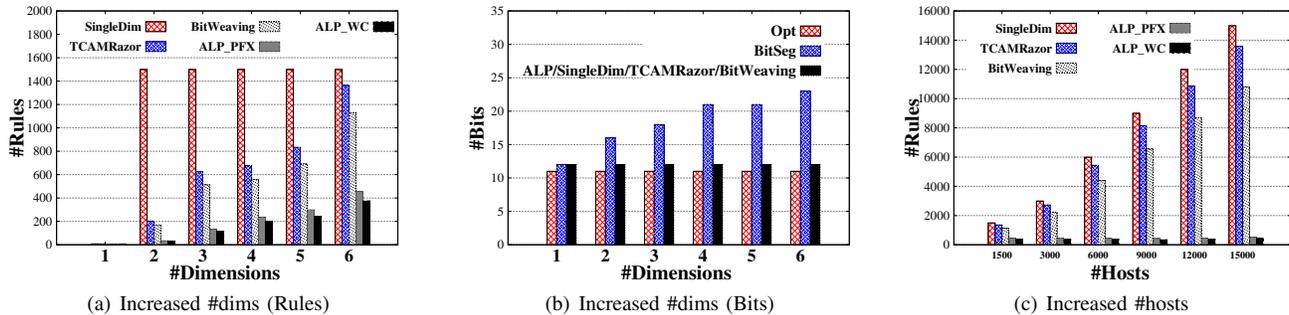


Fig. 8. Comparison with previous approaches. Alpaca simultaneously considers multiple dimensions.

#Dimensions	ORTC	ALP_PFX	ALP_PFX + ORTC
1	8	8	8
2	35	202	24
3	135	627	107
4	237	677	177
5	300	834	227
6	456	1363	316

TABLE V  
#PREFIX RULES: COMPARISON WITH ORTC [12]

dimensions. Six dimensions are chosen (in order): Role, Security Level, Location, Status, CS\_owned and Usage. The initial set of dimensions only contains Role. Then, in each iteration, we add one more dimension to the current set and run Alpaca algorithms to generate classification rules. Figure 8(a) plots the number of rules generated by SingleDim, TCAMRazor and BitWeaving and Alpaca over the six iterations. Given one dimension, all approaches generate a small number of rules. Moving to two dimensions (*i.e.*, Role and Security Level), SingleDim has to potentially enumerate hosts and their attributes, taking as many rules as the number of hosts in the data. The number of rules used by SingleDim is unchanged over the iterations then. TCAMRazor and BitWeaving generates less rules than SingleDim, as they apply compression algorithms for a smaller rule set to represent attributes. Yet, when there are six dimensions, the number of rules (e.g., 1130 wildcard rules) is very close to the number of hosts. In contrast, both Alpaca prefix and wildcard scales well with increasing number of dimensions. Alpaca uses 376 wildcard rules for six dimensions, which are significantly smaller. We show the number of bits in Figure 8(b). Both Alpaca and SingleDim approaches use 12 bits, while the least number of bits (denoted as Opt) is 11 ( $1491 < 2^{11}$ ). BitSeg is infeasible in practice after two dimensions, as it takes more than 16 bits to encode attributes.

**Scale with more hosts.** Our data only covers a single department, but an entire university (with dozens of departments) has more hosts and enterprises can be even larger in sizes. We examine Alpaca’s scalability with more hosts, by synthesizing the host information. We copy each host 2 to 10 times and obtain the scaled-up host data. Figure 8(c) compares several approaches to classify hosts in 6 dimensions. Alpaca scales well, using 436 wildcard rules or 528 prefix rules for around 15000 hosts. Its performance is very stable,

due to the use of aggregated patterns (*e.g.*, wildcard or prefix matches) to classify groups of hosts. The number of hosts does not impact its performance. In contrast, SingleDim potentially needs 15000 rules to enumerate every host; it does not scale to larger networks. The compression algorithms do not help much: TCAMRazor and BitWeaving need 13589 and 10821 rules, respectively, because the single dimension based allocation does not help the aggregation on other dimensions.

In Table V we compare the Alpaca prefix algorithm (ALP\_PFX) with ORTC [12] for the same parameters as in Figure 8(a). For a single dimension and fixed address assignment, ORTC finds an optimal representation with a minimal number of (possibly overlapping) prefix rules. For the ORTC we again assume the SingleDim address assignment. We also examine applying ORTC on the address assignment found by ALP\_PFX. We can see that ORTC cannot compete ALP\_PFX since its address assignment was selected without considering all dimensions. With 5 dimensions for instance, ALP\_PFX and ORTC require 834 and 300 rules, respectively. Their combination reduces the number to only 227, a reduction of 72.8% and 24.3%, respectively. More generally, for 2-6 dimensions, combining the two schemes achieves a reduction of 20.7%-31.4% in comparison with ORTC, the prefix optimal scheme for a fixed address assignment.

**Encode different sets of dimensions.** While Alpaca performs well for three dimensions: Role, Security Level and Location (as shown above), we are curious about its performance on a different set of three dimensions, such as Role, CS\_owned and Operating System. Hence, we fix the number of dimensions to encode and run the algorithm on various sets of dimensions. Figure 9(a) shows the performance of Alpaca to encode seven sets of three dimensions. We do observe the fluctuation: the number of rules generated by Alpaca ranges from 36 to 117 for wildcard case and 44 to 135 for prefix case. Upon closer examination, we find out that the performance is highly correlated with the possible combinations of attributes. Specifically, for the set of dimension {Role, Security Level, Location}, there are 80 combinations of attributes which at least one host is associated with; for the set {Status, CS\_owned, OS}, there are only 22 combinations. Given increasing numbers of combinations of attributes, Alpaca is more likely to generate many rules.

**Update the assignment for new hosts.** We divide hosts

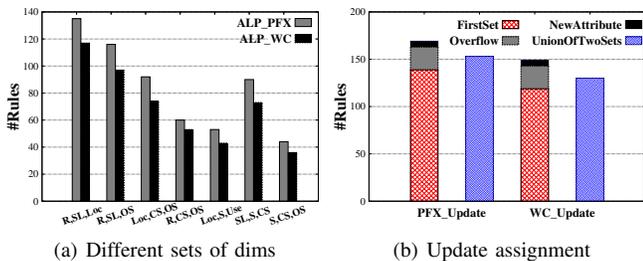


Fig. 9. Property of Alpaca algorithm

into two equal-sized sets based on their created time and run Alpaca to encode four dimensions: Role, Security Level, Status and Usage. We use the first set for the initial hosts and the second set for the newly added hosts. For the first set, Alpaca provisions slack and creates ghost attributes for all dimensions. The second set not only inserts more hosts with the existing attributes but also introduces 4 more new attributes in Security Level. To assign addresses to new hosts with existing attributes, Alpaca uses the slack in the corresponding group. But if the group size is insufficient, Alpaca has to “steal” flow space from the related ghost groups (Section VII-B). To handle the new attributes, Alpaca splits the address space of ghost attributes in the same dimension as well. In our evaluation (Figure 9(b)), the first set (left red bar) uses 139 prefix rules or 119 wildcard rules to represent the four dimensions. Fixing the assignment for the first set, we calculate the extra rules needed to handle the second set. The extra rules come from two parts: (1) hosts with new attributes and (2) overflowed group sizes. As a result, we need an extra 6 rules for the new attributes and 24 rules for the overflowed groups. The overhead is very small compared to 153 prefix rules or 130 wildcard rules for the union of the two sets, where the assignment is computed from scratch without any incremental updates.

## IX. CONCLUSION

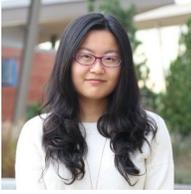
In this paper, we have made three contributions. First, we show the importance and feasibility of considering attributes in IP address allocation. Second, we present the Alpaca system, and two algorithms which cope well with constraints on the IP address space, enterprise churn, and heterogeneity in group sizes. When evaluated with configuration data from several universities, ALP\_WC and ALP\_PFX reduce the number of rules by 50% – 68% and 60% – 68% respectively for multi-table switches. Further, the algorithms have the potential to reduce the total number of rules by over an order of magnitude compared to single dimension based IP address allocation schemes. This can in turn lower the barriers for network administrators to express richer policies involving multiple dimensions. While promising, our results are only a first step. In the future, we hope to build and deploy an actual prototype, as well as evaluate the system with more networks and richer data-sets.

## X. ACKNOWLEDGEMENTS

This work was supported in part by the NSF Career Award No. 0953622 and a Google Research Award.

## REFERENCES

- [1] N. Kang, O. Rottenstreich, S. G. Rao, and J. Rexford, “Alpaca: Compact network policies with attribute-carrying addresses,” in *ACM CoNEXT*, 2015.
- [2] M. Yu *et al.*, “A survey of virtual LAN usage in campus networks,” *IEEE Communications Magazine*, vol. 49, no. 7, pp. 98–103, 2011.
- [3] B. Pfaff *et al.*, “The design and implementation of Open vSwitch,” in *USENIX NSDI*, 2015.
- [4] M. Casado *et al.*, “Rethinking enterprise network control,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 4, pp. 1270–1283, 2009.
- [5] S. K. Fayazbakhsh *et al.*, “Enforcing network-wide policies in the presence of dynamic middlebox actions using FlowTags,” in *NSDI*, 2014.
- [6] N. McKeown *et al.*, “OpenFlow: Enabling innovation in campus networks,” *ACM SIGCOMM CCR*, vol. 38, no. 2, pp. 69–74, 2008.
- [7] P. Bosshart *et al.*, “P4: Programming protocol-independent packet processors,” *ACM SIGCOMM CCR*, vol. 44, no. 3, pp. 87–95, 2014.
- [8] P. Bosshart *et al.*, “Forwarding metamorphosis: Fast programmable match-action processing in hardware for SDN,” in *ACM SIGCOMM*, 2013.
- [9] R. Ozdag, “Intel®Ethernet Switch FM6000 Series-Software Defined Networking,” *Intel Corporation*, 2012.
- [10] M. Appelman and M. D. Boer, “Performance analysis of OpenFlow hardware,” University of Amsterdam, Tech. Rep., Feb 2012, <http://www.delaat.net/rp/2011-2012/p18/report.pdf>.
- [11] D. Y. Huang, K. Yocum, and A. C. Snoeren, “High-fidelity switch models for software-defined network emulation,” in *ACM HotSDN*, 2013.
- [12] R. Draves, C. King, S. Venkatachary, and B. Zill, “Constructing optimal IP routing tables,” in *IEEE INFOCOM*, 1999.
- [13] C. R. Meiners, A. X. Liu, and E. Torng, “TCAM Razor: A systematic approach towards minimizing packet classifiers in TCAMs,” *IEEE/ACM Trans. Netw.*, vol. 18, no. 2, pp. 490–500, 2010.
- [14] C. R. Meiners, A. X. Liu, and E. Torng, “BitWeaving: A non-prefix approach to compressing packet classifiers in TCAMs,” *IEEE/ACM Trans. Netw.*, vol. 20, no. 2, pp. 488–500, 2012.
- [15] R. Wei, Y. Xu, and H. J. Chao, “Block permutations in boolean space to minimize TCAM for packet classification,” in *IEEE INFOCOM*, 2012.
- [16] C. R. Meiners, A. X. Liu, and E. Torng, “Topological transformation approaches to TCAM-based packet classification,” *IEEE/ACM Trans. Netw.*, vol. 19, no. 1, pp. 237–250, 2011.
- [17] R. McGeer and P. Yalagandula, “Minimizing rulesets for TCAM implementation,” in *IEEE INFOCOM*, 2009.
- [18] S. Suri, T. Sandholm, and P. R. Warkhede, “Compressing two-dimensional routing tables,” *Algorithmica*, vol. 35, no. 4, pp. 287–300, 2003.
- [19] D. L. Applegate *et al.*, “Compressing rectilinear pictures and minimizing access control lists,” in *ACM-SIAM SODA*, 2007, pp. 1066–1075.
- [20] O. Rottenstreich *et al.*, “Optimal In/Out TCAM encodings of ranges,” *IEEE/ACM Trans. Netw.*, vol. 24, no. 1, pp. 555–568, 2016.
- [21] O. Rottenstreich and J. Tapolcai, “Optimal rule caching and lossy compression for longest prefix matching,” *IEEE/ACM Trans. Netw.*, 2017.
- [22] K. Kogan *et al.*, “Exploiting order independence for scalable and expressive packet classification,” *IEEE/ACM Trans. Netw.*, vol. 24, no. 2, pp. 1251–1264, 2016.
- [23] O. Rottenstreich *et al.*, “Compressing forwarding tables,” in *IEEE Infocom*, 2013.
- [24] C. R. Meiners, A. X. Liu, E. Torng, and J. Patel, “Split: Optimizing space, power, and throughput for TCAM-based classification,” in *ACM/IEEE ANCS*, 2011.
- [25] X. Jin, L. E. Li, L. Vanbever, and J. Rexford, “SoftCell: Scalable and flexible cellular core network architecture,” in *ACM CoNEXT*, 2013.
- [26] R. MacDavid *et al.*, “Concise encoding of flow attributes in SDN switches,” in *ACM Symposium on SDN Research (SOSR)*, 2017.
- [27] S. Donovan and N. Feamster, “NetAssay: Providing new monitoring primitives for network operators,” in *ACM HotNets*, 2014.
- [28] T. Benson, A. Akella, and D. A. Maltz, “Mining policies from enterprise network configuration,” in *ACM IMC*, 2009.
- [29] R. Wang, D. Butnariu, and J. Rexford, “Openflow-based server load balancing gone wild,” in *USENIX Hot-ICE*, 2011.



**Nanxi Kang** is a Software Engineer in Databricks, Inc. She received Master and Ph.D. degree from the Computer Science department in Princeton University, Princeton, US in 2013 and 2016 respectively. She is a recipient of Wu Fellowship in Princeton and Google Anita Borg Memorial Scholarship. She has won the Best paper award at EuroSys conference.



**Ori Rottenstreich** is a Postdoctoral Research Fellow at the department of Computer Science, Princeton university. He received the B.S. in Computer Engineering (summa cum laude) and Ph.D. degree from the Electrical Engineering department of the Technion, Haifa, Israel in 2008 and 2014, respectively. He is a recipient of the Rothschild Yad-Hanadiv postdoctoral fellowship, the Google Europe PhD Fellowship in Computer Networking, the Andrew Viterbi graduate fellowship, the Jacobs-Qualcomm fellowship, the Intel graduate fellowship and the

Gutwirth Memorial fellowship. He also received the Best Paper Runner Up Award at the IEEE Infocom 2013 conference.



**Sanjay G. Rao** is an Associate Professor in the School of Electrical and Computer Engineering at Purdue University, where he leads the Internet Systems Laboratory. His research spans network management, cloud computing, and Internet video distribution. He received a B.Tech in Computer Science and Engineering from the Indian Institute of Technology, Madras, and the Ph.D from the School of Computer Science, Carnegie Mellon University. He has been a Visiting Researcher at Google, AT&T Research and Princeton University. He is a recipient

of the NSF Career award, and won the ACM SIGMETRICS Test of Time Award for his work on End System Multicast (peer-to-peer video streaming). He has served on the Technical Program Committees of conferences including ACM Sigcomm, Usenix NSDI, and ACM CoNext, has served as the Area technical program chair of IEEE Infocom, is currently an Editor for the IEEE/ACM Transactions on Networking, and was the technical program co-chair of the INM/WREN workshop held in conjunction with NSDI 2010.



**Jennifer Rexford** is the Gordon Y.S. Wu Professor of Engineering and the Chair of Computer Science at Princeton University. Before joining Princeton in 2005, she worked for eight years at AT&T Labs–Research. Jennifer received her BSE degree in electrical engineering from Princeton University in 1991, and her PhD degree in electrical engineering and computer science from the University of Michigan in 1996. She is co-author of the book “Web Protocols and Practice” (Addison-Wesley, May 2001). She served as the chair of ACM SIGCOMM from 2003

to 2007. Jennifer was the 2004 winner of ACM’s Grace Murray Hopper Award for outstanding young computer professional. She is an ACM Fellow (2008), and a member of the American Academy of Arts and Sciences (2013) and the National Academy of Engineering (2014).