Topological Transformation Approaches to TCAM-Based Packet Classification*

Chad R. Meiners Alex X. Liu[†] Eric Torng Department of Computer Science and Engineering Michigan State University East Lansing, MI 48824, U.S.A. {meinersc, alexliu, torng}@cse.msu.edu

Abstract-Several range reencoding schemes have been proposed to mitigate the effect of range expansion and the limitations of small capacity, large power consumption, and high heat generation of TCAM-based packet classification systems. However, they all disregard the semantics of classifiers and therefore miss significant opportunities for space compression. In this paper, we propose new approaches to range reencoding by taking into account classifier semantics. Fundamentally different from prior work, we view reencoding as a topological transformation process from one colored hyperrectangle to another where the color is the decision associated with a given packet. Stated another way, we reencode the entire classifier by considering the classifier's decisions rather than reencode only ranges in the classifier ignoring the classifier's decisions as prior work does. We present two orthogonal, yet composable, reencoding approaches, domain compression and prefix alignment. Our techniques significantly outperform all previous reencoding techniques. In comparison with prior art, our experimental results show that our techniques achieve at least 5 times more space reduction in terms of TCAM space for an encoded classifier and at least 3 times more space reduction in terms of TCAM space for a reencoded classifier and its transformers. This, in turn, leads to improved throughput and decreased power consumption.

Index Terms—Hardware-based Packet Classification, Ternary Content Addressable Memory (TCAM), Range Encoding.

I. INTRODUCTION

Packet classification is the core mechanism that enables many networking devices, such as routers and firewalls, to perform services such as packet filtering, virtual private networks (VPNs), network address translation (NAT), quality of service (QoS), load balancing, traffic accounting and monitoring, differentiated services (Diffserv), etc. The basic classification problem is to compare each packet with a list of predefined rules and find the first (i.e., highest priority) rule that the packet matches. Table I shows an example classifier of two rules. The format of these rules is based upon the format used in Access Control Lists on Cisco routers.

Packet classification is often a performance bottleneck for routers as they need to classify every packet. Achieving wire speed packet classification has long been a networking

† Alex X. Liu is the corresponding author of this paper. Email: alexliu@cse.msu.edu. Telephone: +1 (517) 353-5152. Fax: +1 (517) 432-1061.

Rule	Src. IP	Dest. IP	Src. Port	Dest. Port	Prot.	Action	
r_1	1.2.3.0/24	192.168.0.1	[1,65534]	[1,65534]	TCP	accept	
r_2	*	*	*	*	*	discard	
TABLE I							

AN EXAMPLE PACKET CLASSIFIER

goal. Although software-based packet classification has been extensively studied [27], using Ternary Content Addressable Memories (TCAMs) to perform hardware-based packet classification has become the de facto industrial standard [13].

A traditional random access memory chip receives an address and returns the content of the memory at that address. A TCAM chip, however, works in a reverse manner: it receives content and returns the address of the *first* entry where the content lies in the TCAM in constant time (*i.e.*, a few dozen CPU cycles). Exploiting this hardware feature, TCAM-based packet classifiers store a rule in each entry as an array of 0's, 1's, or *'s (*don't-care* values). A packet header (*i.e.*, a search key) matches an entry if and only if their corresponding 0's and 1's match. Given a search key to a TCAM, the hardware circuits compare the key with all its occupied entries in parallel and return the index (or sometimes the content, depending on chip configuration,) of the first matching entry.

Unfortunately, TCAM-based solutions may not scale up to meet the classification needs of the rapidly growing Internet where classifiers are growing rapidly in size. First, current TCAMs have limited capacity. The largest available TCAM chip has a capacity of only 72 megabits (Mb). Furthermore, the well known range expansion problem exacerbates the problem of limited capacity. In a typical rule, the three fields of source and destination IP addresses and protocol type are specified as prefixes where all the *s are at the end of the ternary string, so the fields can be directly stored in a TCAM. However, the other two fields of source and destination port numbers are specified in ranges (*i.e.*, integer intervals), which need to be mapped to often many prefixes before being stored in a TCAM. This can lead to a large increase in the number of TCAM entries needed to encode a rule. For example, 30 prefixes are needed to represent the range [1, 65534], so $30 \times 30 = 900$ TCAM entries are required to represent the single rule r_1 in Table I. Second, TCAM chip size growth has been and will likely continue to be slow due to their extremely high circuit density. Finally, even if larger TCAM chips were available, their deployment may be limited due to their high power consumption, large footprints, and high cost. TCAM chips consume lots of power and generate lots of

¹The preliminary version of this paper titled "Topological Transformation Approaches to Optimizing TCAM-Based Packet Processing Systems" was published in the proceedings of the International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS), pages 73-84, Seattle, Washington, June 2009.

heat because every memory access searches the entire active memory in parallel, and TCAM power consumption grows linearly with the number of ternary bits searched in each memory access [30]. Power constrains deployed TCAM chip size when systems designers must obey a "power budget", *e.g.*, TCAM components may use 10% of an entire board's power budget. Likewise, TCAM chips occupy 6 times (or more) board space than an equivalent SRAM which leads to TCAMs having high costs, even in large quantities. Due to these issues, the most popular TCAM chips in 2004 were the 1Mb and 2Mb chips even though a 36Mb TCAM chip was then available [1].

Range reencoding schemes have been proposed to improve the scalability of TCAMs, primarily by mitigating the effect of range expansion [5], [6], [13], [18], [21], [22], [29], [31]. The basic idea is to first reencode a classifier into another classifier that requires less TCAM space and then reencode each packet correspondingly such that the decision made by the reencoded classifier for the reencoded packet is the same as the decision made by the original classifier for the original packet. Range reencoding has two possible benefits: rule width compression so that narrower TCAM entries can be used and rule number compression so that fewer TCAM entries can be used.

We observe that all previous reencoding schemes suffer from one fundamental limitation: they all ignore the decision associated with each rule and thus the classifier's decision for each packet. Disregarding classifier semantics leads all previous techniques to miss significant opportunities for space compression. Fundamentally different from prior work, we view reencoding as a topological transformation process from one colored hyperrectangle to another where the color is the decision associated with a given packet. Topological transformation allows us to reencode the entire classifier as opposed to reencoding only the ranges in a classifier. Furthermore, we also view reencoding as a classification process that can be implemented with small TCAM tables, which enables fast packet reencoding. We present two orthogonal, yet composable, reencoding approaches, domain compression and prefix alignment. In domain compression, we transform a given colored hyperrectangle, which represents the semantics of a given classifier, to the smallest possible "equivalent" colored hyperrectangle. This leads to both optimal rule width *compression* as well as rule number compression. In prefix alignment, on the other hand, we strive for rule number compression only by transforming a colored hyperrectangle to an equivalent "prefix-friendly" colored hyperrectangle where the ranges align well with prefix boundaries, minimizing the costs of range expansion.

Domain Compression: In most packet classifiers, many coordinates (*i.e.*, values) within a field domain are equivalent. The idea of domain compression is to reencode the domain so as to eliminate as many redundant coordinates as possible. This leads to both rule width and rule number compression. From a geometric perspective, domain compression "squeezes" a colored hyperrectangle as much as possible. For example, consider the colored rectangle in Figure 1(A) that represents the classifier in Figure 1(H). In field F_1 represented by the X-axis, all values in $[0,7] \cup [66,99]$ are equivalent; that is, for any

 $y \in F_2$ and any $x_1, x_2 \in [0,7] \cup [66,99]$, packets (x_1, y) and (x_2, y) have the same decision. Therefore, when reencoding F_1 , we can map all values in $[0,7] \cup [66,99]$ to a single value, say 0. By identifying such equivalences along all dimensions, the rectangle in Figure 1(A) is reencoded to the one in Figure 1(D), whose corresponding classifier is shown in Figure 1(I). Figures 1(B) and (C) show the two transforming tables for F_1 and F_2 , respectively. We use "a" as a shorthand for "accept" and "d" as a shorthand for "discard".



Fig. 1. Example of topological transformations

Prefix Alignment: In prefix alignment, we "shift", "shrink", or "stretch" ranges by transforming the domain of each field to a new "prefix-friendly" domain so that the majority of the reencoded ranges either are prefixes or can be expressed by a small number of prefixes. This will reduce the costs of range expansion and leads to rule number compression with a potentially small loss in rule width compression. For example, consider the packet classifier in Figure 1(I), whose corresponding rectangle is in Figure 1(D). Range expansion will yield 5 prefix rules because interval [1,2] or [01,10]cannot be combined into one prefix. However, by transforming the rectangle in Figure 1(D) to the one in Figure 1(G), the range expansion of the resulting classifier, as shown in Figure 1(J), will have 3 prefix rules because [2, 3] is expanded to 1*. Figures 1(E) and (F) show the two transforming tables for F_1 and F_2 , respectively.

Our domain compression and prefix alignment techniques have several nice properties. First, they are powerful in reducing TCAM space. In our experiments on real-world and synthetic classifiers, they achieved at least 5 times more space reduction with transformers excluded and at least 3 times more space reduction with transformers included in comparison with current state-of-the-art reencoding techniques. Second, they can be easily implemented on existing hardware by using TCAM to perform reencoding. Third, not only are they composable, they can also be composed with many other TCAM optimization and reencoding schemes proposed in prior work because the reencoded classifier produced by domain compression contains range rules and the prefix alignment technique can take any prefix classifier as its input. Fourth, because they allow the use of smaller TCAM chips, they can lead to improved throughput and decreased power consumption even though more TCAM searches are needed to classify a packet.

The rest of the paper proceeds as follows. We start by reviewing previous work in Section II and formally defining relevant terms in Section III. In Section IV, we give an overview of our topological transformation approaches. In Sections V and VI, we present the technical details of the two topological transformation approaches. We discuss implementation issues in Section VII. Experimental results are presented in Section VIII and performance modeling results are presented in Section IX. We draw conclusions in Section X.

II. RELATED WORK

Prior work in optimizing TCAM-based packet classification systems fall into three broad categories: circuit modification, classifier compression, and range reencoding.

Circuit Modification: Spitznagel *et al.* proposed adding comparators at each entry level to better accommodate range matching [24]. While this research direction is important, such solutions are hard to deploy due to high cost [13].

Classifier Compression: These optimizations convert a given packet classifier to another semantically equivalent classifier that requires fewer TCAM entries. The schemes in [3], [8], [26] focus on one-dimensional and two dimensional packet classifiers. The redundancy removal algorithms in [15]–[17] can reduce TCAM usage by eliminating all the redundant rules in a packet classifier. In [7], Dong et al. proposed schemes to reduce range expansion by repeatedly expanding or trimming ranges to prefix boundaries without changing the number of bits used to represent each dimension. They ensure correctness by using core effective region algorithms in [15]. In essence, they insert the new rule before the rule being modified and check if the new rule is redundant. In contrast, our prefix alignment algorithm mitigates range expansion by intelligently adding bits to a given dimension to increase the number of prefix boundaries. In [19] Meiners, et al. proposed a greedy algorithm that finds locally minimal solutions along each field and combines these solutions into a smaller equivalent packet classifier. In [20], Meiners et al. proposed the first algorithm that can compress a given classifier into a non-prefix ternary classifier.

Range Reencoding: Previous range reencoding schemes fall into two categories: those that only consider rule set size, often at the expense of rule width [5], [6], [13], [18], [31] and those that attempt to both compress rule set size and rule width [4], [21], [22], [29]. In [18], Liu proposed a scheme that allocates specific TCAM column bits to represent ranges in a manner similar to Lakshman and Stiliadis' software bitmap classification method [12]. Lakshminarayan et al. [13] proposed a scheme called fence encoding, which encodes interval ranges as a range of unary numbers. Fence encoding has an expansion factor of one, meaning all ranges can be encoded with one string, but the number of unary bits required for each rule is prohibitive. To reduce rule width, Lakshminarayan et al. proposed DIRPE, which compresses the width of fence encodings at the expense of a larger expansion factor. Bremler-Barr and Hendler [5] proposed SRGE, which utilizes the structural properties of binary reflected gray codes to reduce range expansion without increasing rule width. Lunteren and Engbersen proposed a hierarchy of three methods, P^2C , that can be used to compress both rule number and rule width [29]. Two methods guarantee an expansion factor of one but have potentially larger rule widths. The third method has the best rule width compression at the cost of expansion factors greater than one. Bremmel-Bar *et al.* [4] purpose concrete algorithms for the P^2C hierarchy. Pao *et al.* proposed a prefix inclusion method (PIC) that achieves better rule width compression than P^2C [21], [22]. Che *et al.* [6], [31] and Pao *et al.* [21], [22] propose using TCAMs to reencode packets.

Reencoding has been used in software based packet classification. Lakshman and Stiliadis proposed to reencode each field's value into a bitmap that specifies the containment relationship among values and rules [12]. Given a reencoded packet, this method uses customized parallel AND gates to perform an intersection of these bitmaps and ultimately find the first matching rule. Srinivasan et al. proposed an encoding method called *cross-producting* that assigns a unique number to each disjoint range within a classifier field and constructs a lookup table for the cross product of the numbers associated with each field [25]. Gupta and McKeown proposed Recursive Flow Classification (RFC) [10], an optimized version of the cross-producting scheme that uses recursive cross-producting tables to reduce the space requirements of regular cross producting tables. Furthermore, they map disjoint ranges that are contained by the same set of rules into a single value. RFC's mapping tables use a weaker equivalence relation than our domain compression technique, so they do not achieve as much compression as we do. Unfortunately, these software based reencoding methods are difficult to deploy because the required RAM to perform the reencoding is extremely large. By using TCAMs to perform reencoding, we overcome this memory issue.

III. FORMAL DEFINITIONS

We now formally define the concepts of fields, packets, and packet classifiers. A *field* F_i is a variable of finite length (*i.e.*, of a finite number of bits). The domain of field F_i of w bits, denoted $D(F_i)$, is $[0, 2^w - 1]$. A *packet* over the d fields F_1, \dots, F_d is a d-tuple (p_1, \dots, p_d) where each p_i $(1 \le i \le d)$ is an element of $D(F_i)$. Packet classifiers usually check the following five fields: source IP address, destination IP address, source port number, destination port number, and protocol type. The lengths of these packet fields are 32, 32, 16, 16, and 8, respectively. We use Σ to denote the set of all packets over fields F_1, \dots, F_d . It follows that Σ is a finite set and $|\Sigma| = |D(F_1)| \times \dots \times |D(F_d)|$, where $|\Sigma|$ denotes the number of elements in set Σ and $|D(F_i)|$ denotes the number of elements in set $D(F_i)$.

A rule has the form $\langle predicate \rangle \rightarrow \langle decision \rangle$. A $\langle predicate \rangle$ defines a set of packets over the fields F_1 through F_d , and is specified as $F_1 \in S_1 \land \cdots \land F_d \in S_d$ where each S_i is a subset of $D(F_i)$ and is specified as either a prefix or a nonnegative integer interval. A prefix $\{0,1\}^k \{*\}^{w-k}$ with k leading 0s or 1s for a packet field of length w denotes the integer interval $[\{0,1\}^k \{0\}^{w-k}, \{0,1\}^k \{1\}^{w-k}]$. For example, prefix 01^{**} denotes the interval [0100, 0111]. A rule $F_1 \in S_1 \land \cdots \land F_d \in S_d \rightarrow \langle decision \rangle$ is a prefix rule if and only if each S_i is represented as a prefix.

A packet matches a rule if and only if the packet matches the predicate of the rule. A packet (p_1, \dots, p_d) matches a predicate $F_1 \in S_1 \land \dots \land F_d \in S_d$ if and only if the condition $p_1 \in S_1 \land \dots \land p_d \in S_d$ holds. We use DS to denote the set of possible values that $\langle decision \rangle$ can be. Typical elements of DS include accept, discard, accept with logging, and discard with logging.

A sequence of rules $\langle r_1, \dots, r_n \rangle$ is *complete* if and only if for any packet p, there is at least one rule in the sequence that p matches. To ensure that a sequence of rules is complete and thus a packet classifier, the predicate of the last rule is usually specified as $F_1 \in D(F_1) \wedge \cdots \in F_d \in AD(F_d)$. A packet classifier \mathbb{C} is a sequence of rules that is complete. The size of \mathbb{C} , denoted $|\mathbb{C}|$, is the number of rules in \mathbb{C} . A packet classifier \mathbb{C} is a *prefix packet classifier* if and only if every rule in \mathbb{C} is a prefix rule.

Two rules in a packet classifier may *overlap*; that is, a single packet may match both rules. Furthermore, two rules in a packet classifier may *conflict*; that is, the two rules not only overlap but also have different decisions. Packet classifiers typically resolve such conflicts by employing a first-match resolution strategy where the decision for a packet p is the decision of the first (i.e., highest priority) rule that p matches in \mathbb{C} . The decision that packet classifier \mathbb{C} makes for packet p is denoted $\mathbb{C}(p)$.

We can think of a packet classifier \mathbb{C} as defining a many-toone mapping function from Σ to DS. Two packet classifiers \mathbb{C}_1 and \mathbb{C}_2 are *equivalent*, denoted $\mathbb{C}_1 \equiv \mathbb{C}_2$, if and only if they define the same mapping function from Σ to DS; that is, for any packet $p \in \Sigma$, we have $\mathbb{C}_1(p) = \mathbb{C}_2(p)$. A rule is *redundant* in a classifier if and only if removing the rule does not change the semantics of the classifier.

In a typical packet classifier rule, the fields of source IP, destination IP, and protocol type are specified in prefix format, which can be directly stored in TCAMs; however, the remaining two fields of source port and destination port are specified as ranges (i.e., non-negative integer intervals), which are typically converted to prefixes before being stored in TCAMs. This leads to range expansion, the process of converting a non-prefix rule to prefix rules. In range expansion, each field of a rule is first expanded separately. The goal is to find a minimum set of prefixes such that the union of the prefixes corresponds to the range. For example, if one 3-bit field of a rule is the range [1, 6], a corresponding minimum set of prefixes would be 001, 01*, 10*, 110. The worst-case range expansion of a w-bit range results in a set containing 2w-2prefixes [11]. The next step is to compute the cross product of the set of prefixes for each field, resulting in a potentially large number of prefix rules.

IV. TOPOLOGICAL TRANSFORMATION

Given a *d*-dimensional classifier \mathbb{C} over fields F_1, \dots, F_d , a topological transformation process produces two separate components. The first component is a set of *transformers* $\mathbb{T} = \{\mathbb{T}_i \mid 1 \leq i \leq d\}$ where transformer \mathbb{T}_i transforms $D(F_i)$ into a new domain $D'(F_i)$. Together, the set of transformers \mathbb{T} transforms the original packet space Σ into a new packet space Σ' . The second component is a transformed *d*-dimensional classifier \mathbb{C}' over packet space Σ' such that for any packet $(p_1, \dots, p_d) \in \Sigma$, the following condition holds:

$$\mathbb{C}(p_1,\cdots,p_d)=\mathbb{C}'(\mathbb{T}_1(p_1),\cdots,\mathbb{T}_d(p_d))$$

Each of the *d* transformers \mathbb{T}_i and the transformed packet classifier \mathbb{C}' are implemented in TCAM.

The TCAM space needed by our transformation approach is measured by the total TCAM space needed by the d + 1tables: $\mathbb{C}', \mathbb{T}_1, \dots, \mathbb{T}_d$. We define the space used by a classifier or transformer in a TCAM as the number of entries (*i.e.*, rules) multiplied by the width of the TCAM in bits: $space = \# of entries \times TCAM width$. Although TCAMs can be configured with varying widths, they do not allow arbitrary widths. The width of a TCAM typically can be set at 40, 80, 160, and 320 bits (per entry). The primary goal of the transformation approach is to produce $\mathbb{C}', \mathbb{T}_1, \cdots, \mathbb{T}_d$ such that the TCAM space needed by these d + 1 TCAM tables is much smaller than the TCAM space needed by the original classifier \mathbb{C} . Most previous reencoding approaches ignore the space required by the transformers and only focus on the space required by the transformed classifier \mathbb{C}' . Note that we can implement the table for the protocol field using SRAM if desired since the field has only 8 bits.

There are two natural architectures for storing the d + 1TCAM tables $\mathbb{C}', \mathbb{T}_1, \dots, \mathbb{T}_d$: the multi-lookup architecture and the pipelined-lookup architecture.

In the multi-lookup architecture, we store all the d+1 tables in one TCAM chip. For each table, we prepend a $\lceil \log(d+1) \rceil$ *table ID* bit string to every entry. Figure 2 illustrates the packet classification process using the multi-lookup architecture when d = 2. Suppose we use the table IDs 00, 01, and 10 for the three tables \mathbb{C}' , \mathbb{T}_1 , and \mathbb{T}_2 , respectively. Given a packet (p_1, p_2) , we first concatenate \mathbb{T}_1 's table ID 01 with p_1 and use the resulting bit string $01|p_1$ as the search key for the TCAM. Let p_1' denote the search result. Second, we concatenate \mathbb{T}_2 's table ID 10 with p_2 and use the resulting bit string $10|p_2$ as the search key for the TCAM. Let p_2' denote the search result. Third, we concatenate the table ID 00 of \mathbb{C}' with p_1' and p_2' and use the resulting bit string $00|p_1'|p_2'$ as the search key for the TCAM. The search result is the final decision for the given packet (p_1, p_2) .



Fig. 2. Multi-lookup

There are two natural pipelined-lookup architectures: parallel pipelined-lookup and chained pipelined-lookup. In both, we store the d+1 tables in d+1 separate TCAMs, so table IDs are no longer needed. In the parallel pipelined-lookup architecture, the d transformer tables \mathbb{T} , laid out in parallel, form a twoelement pipeline with the transformed classifier \mathbb{C}' . Figure 3 illustrates the packet classification process using the parallel pipelined-lookup architecture when d = 2. Given a packet (p_1, p_2) , we send p_1 and p_2 , in parallel over separate buses, to \mathbb{T}_1 and \mathbb{T}_2 , respectively. Then, the search result $p_1'|p_2'$ is used as a key to search on \mathbb{C}' . This second search result is the final decision for the given packet (p_1, p_2) . Figure 4 illustrates the packet classification process using the chained pipelined-lookup architecture when d = 2. We focus primarily on the parallel pipelined-lookup architecture as this allows us to minimize latency.



The main advantage of the multi-lookup architecture is that it can be easily deployed since it requires minimal modification of existing TCAM-based packet processing systems. Its main drawback is a modest slowdown in packet processing throughput because d + 1 TCAM searches are required to process a d-dimensional packet. In contrast, the main advantage of the two pipelined-lookup architectures is high packet processing throughput. Their main drawback is that the hardware needs to be modified to accommodate d + 1 TCAM chips (or d chips if SRAM is used for the protocol field). We present a performance modeling analysis of the parallel pipelinedlookup and multi-lookup architectures in Section IX.

V. DOMAIN COMPRESSION

We now describe our *domain compression* technique. The basic idea is to simplify the logical structure of a classifier by mapping the domain of each field $D(F_i)$ to the smallest possible domain $D'(F_i)$. We implement domain compression by exploiting the equivalence classes that any classifier \mathbb{C} defines on the domain of each of its fields. Domain compression is especially powerful because it contributes to both rule width compression, which allows us to use 40 bit TCAM entries instead of 160 bit TCAM entries, and rule number compression because each transformed rule r' in classifier \mathbb{C}' will contain fewer equivalence classes than the original rule r did in classifier \mathbb{C} . Through domain compression and redundancy removal, \mathbb{C}' typically has far fewer rules than \mathbb{C} did, something no other reencoding scheme can achieve.

Our domain compression algorithm consists of three steps: (1) computing equivalence classes, (2) constructing transformer \mathbb{T}_i for each field F_i , and (3) constructing the transformed classifier \mathbb{C}' .

A. Step 1: Compute Equivalence Classes

We first formally define the equivalence relation that classifier \mathbb{C} defines on each field domain and the resulting equivalence classes. We use the notation Σ_{-i} to denote the set of all (d-1)-tuple packets over the fields $(F_1, \dots, F_{i-1}, F_{i+1}, \dots, F_d)$ and p_{-i} to denote an element of Σ_{-i} . Then we use $\mathbb{C}(p_i, p_{-i})$ to denote the decision that packet classifier \mathbb{C} makes for the packet p that is formed by combining $p_i \in D(F_i)$ and p_{-i} .

Definition 5.1 (Equivalence Class): Given a packet classifier \mathbb{C} over fields F_1, \dots, F_d , we say that $x, y \in D(F_i)$ for $1 \leq i \leq d$ are equivalent with respect to \mathbb{C} if and only if $\mathbb{C}(x, p_{-i}) = \mathbb{C}(y, p_{-i})$ for any $p_{-i} \in \Sigma_{-i}$. It follows that \mathbb{C} partitions $D(F_i)$ into *equivalence classes*. We use the notation $\mathbb{C}\{x\}$ to denote the equivalence class that x belongs to as defined by classifier \mathbb{C} .

In domain compression, we compress every equivalence class in each domain $D(F_i)$ to a single point in $D'(F_i)$. The crucial tool of our domain compression algorithm is the Firewall Decision Diagram (FDD) [9]. A Firewall Decision Diagram (FDD) with a decision set DS and over fields F_1, \dots, F_d is an acyclic and directed graph that has the following five properties: (1) There is exactly one node that has no incoming edges. This node is called the *root*. The nodes that have no outgoing edges are called *terminal* nodes. (2) Each node v has a label, denoted F(v), such that

$$F(v) \in \begin{cases} \{F_1, \cdots, F_d\} & \text{if } v \text{ is a nonterminal node,} \\ DS & \text{if } v \text{ is a terminal node.} \end{cases}$$

(3) Each edge $e: u \to v$ is labeled with a nonempty set of integers, denoted I(e), where I(e) is a subset of the domain of u's label (*i.e.*, $I(e) \subseteq D(F(u))$). (4) A directed path from the root to a terminal node is called a *decision path*. No two nodes on a decision path have the same label. (5) The set of all outgoing edges of a node v, denoted E(v), satisfies the following two conditions: (i) Consistency: $I(e) \cap I(e') = \emptyset$ for any two distinct edges e and e' in E(v). (ii) Completeness: $\bigcup_{e \in E(v)} I(e) = D(F(v))$. Two nodes v and v' in an FDD are *isomorphic* if and only if v and v' satisfy one of the following two conditions: (1) both v and v' are terminal nodes with identical labels; (2) both v and v' are nonterminal nodes and there is a one-to-one correspondence between the outgoing edges of v and the outgoing edges of v' such that every pair of corresponding edges have identical labels and they both point to the same node.

We define a *full-length ordered FDD* as an FDD where in each decision path all fields appear in the same order. For simplicity, we use the term "FDD" to mean "full-length ordered FDD" if not otherwise specified. Given a classifier \mathbb{C} , the FDD construction algorithm in [14] can convert it to an equivalent full-length ordered FDD f.

After an FDD f is constructed, we can reduce f's size by merging isomorphic subgraphs. A full-length ordered FDD f is *reduced* if and only if it satisfies the following two conditions: (1) no two nodes in f are isomorphic; (2) no two nodes have more than one edge between them. A reduced FDD is essentially a canonical representation for packet classifiers.

The first step of our domain compression algorithm is to convert a given d-dimensional packet classifier \mathbb{C} to d equivalent reduced FDDs f_1 through f_d where the root of FDD f_i is labeled by field F_i . Figure 5(a) shows an example packet classifier over two fields F_1 and F_2 where the domain of each field is [0,63]. Figures 5(b) and (c) show the two FDDs f_1 and f_2 , respectively. The FDDs f_1 and f_2 are almost reduced except that the terminal nodes are not merged together for illustration purposes.

The crucial observation is that each edge out of reduced FDD f_i 's root node corresponds to one equivalence class of domain $D(F_i)$. For example, consider the classifier in Figure 5(a) and the corresponding FDD f_1 in Figure 5(b). Obviously,



Fig. 5. Example of domain compression

for any p_1 and p_1' in $[7, 11] \cup [16, 19] \cup [39, 40] \cup [43, 60]$, we have $\mathbb{C}(p_1, p_2) = \mathbb{C}(p_1', p_2)$ for any p_2 in [0,63], so it follows that $\mathbb{C}\{p_1\} = \mathbb{C}\{p_1'\}$.

Theorem 5.1 (Equivalence Class Theorem): For any packet classifier \mathbb{C} over fields F_1, \dots, F_d and an equivalent reduced FDD f_i rooted at an F_i node v, the labels of v's outgoing edges are all the equivalence classes over field F_i as defined by \mathbb{C} .

B. Step 2: Construct Transformers

Given a packet classifier \mathbb{C} over fields F_1, \dots, F_d and the d equivalent reduced FDDs f_1, \cdots, f_d where the root node of f_i is labeled F_i , we compute transformer \mathbb{T}_i as follows. Let v be the root of f_i with m outgoing edges e_1, \dots, e_m . First, for each edge e_j out of v, we choose one of the ranges in e_j 's label to be a representative label, which we call the landmark. By Theorem 5.1, all the ranges in e_i 's label belong to the same equivalence class, so any one of them can be chosen as the landmark. For each equivalence class, we choose the range that intersects the fewest number of rules in \mathbb{C} as the landmark breaking ties arbitrarily. We then sort edges in the increasing order of their landmarks. We use L_j and e_j to denote the landmark range and the corresponding edge in sorted order where edge e_1 has the smallest valued landmark L_1 and edge e_m has the largest valued landmark L_m . Our transformer \mathbb{T}_i then maps all values in e_j 's label to value j where $1 \le j \le m$. For example, in Figures 5(b) and (c), the greyed ranges are chosen as the landmarks of their corresponding equivalence classes, and Figures 5(d) and (e) show transformers \mathbb{T}_1 and \mathbb{T}_2 that result from choosing those landmarks.

C. Step 3: Construct Transformed Classifier

We now construct transformed classifier \mathbb{C}' from classifier \mathbb{C} using transformers \mathbb{T}_i for $1 \leq i \leq d$ as follows. Let $F_1 \in S_1 \wedge \cdots \wedge F_d \in S_d \rightarrow \langle decision \rangle$ be an original rule in \mathbb{C} . The domain compression algorithm converts $F_i \in S_i$ to $F_i' \in S_i'$ such that $S_i' = \{j | 0 \leq j \leq m - 1 \wedge L_j \cap S_i \neq \emptyset\}$. Stated another way, we replace range S_i with range $[a, b] \subseteq D'(F_i)$ where a is the smallest number in [0, m-1] such that $L_a \cap S_i \neq \emptyset$. Note, it is possible no landmark ranges intersect range S_i ; in this case a and b are undefined and $S_i' = \emptyset$. For a converted rule $r' = F_1' \in S_1' \wedge \cdots \wedge F_d' \in S_d' \rightarrow \langle decision \rangle$ in \mathbb{C}' , if there exists $1 \leq i \leq d$ such that $S_i' = \emptyset$, we delete this converted rule r' from \mathbb{C}' .

Consider the rule $F_1 \in [7, 60] \land F_2 \in [10, 58] \rightarrow discard$ in the example classifier in Figure 5(a). For field F_1 , the five landmarks are the five greyed intervals in 5(b), namely [0,0], [1,6], [7,11], [12,15], and [63, 63]. Among these five landmarks, [7,60] overlaps with [7,11] and [12,15], which are mapped to 2 and 3 respectively by transformer \mathbb{T}_1 . Thus, $F_1 \in [7,60]$ is converted to $F_1' \in [2,3]$. Similarly, for field F_2 , [10,58] overlaps with only one of F_2 's landmarks, [10,19], which is mapped to 3 by F_2 's mapping table. Thus, $F_2 \in [10,58]$ is converted to $F_2' \in [3,3]$.

We now prove that \mathbb{C}' together with \mathbb{T} is semantically equivalent to \mathbb{C} .

Theorem 5.2: Consider any classifier \mathbb{C} and the resulting transformers \mathbb{T} and transformed classifier \mathbb{C}' . For any packet $p = (p_1, \dots, p_d)$, we have

 $\mathbb{C}(p_1,\cdots,p_d)=\mathbb{C}'(\mathbb{T}_1(p_1),\cdots,\mathbb{T}_d(p_d)).$

Proof: For each field F_i for $1 \leq i \leq d$, consider p's field value p_i . Let $L(p_i)$ be the landmark range for $\mathbb{C}\{p_i\}$. We set $x_i = \min(L(p_i))$. We now consider the packet $x = (x_1, \dots x_d)$ and the packets $x(j) = (x_1, \dots x_{j-1}, p_j, \dots, p_d)$ for $0 \leq j \leq d$; that is, in packet x(j), the first j fields are identical to packet x and the last d - j fields are identical to packet p. Note x(0) = p and x(d) = x. We now show that $\mathbb{C}(p) = \mathbb{C}(x)$. This follows from $\mathbb{C}(x(0)) = \mathbb{C}(x(1)) = \dots = \mathbb{C}(x(d))$. Each equality follows from the fact that x_j and p_j belong to the same equivalence class within $D(F_j)$.

Let r be the first rule in \mathbb{C} that packet x matches. We argue that p' will match the transformed rule $r' \in \mathbb{C}'$. Consider the conjunction $F_i \in S_i$ of rule r. Since x matches rule r, it must be the case that $x_i \in S_i$. This implies that $L(p_i) \cap S_i \neq i$ \emptyset . Thus, by our construction $p_i' = \mathbb{T}_i(p_i) = \mathbb{T}_i(x_i) \in S_i'$. Since this holds for all fields F_i , packet p' matches rule r'. We also argue that packet p' will not match any rule before transformed rule $r' \in \mathbb{C}'$. Suppose packet p' matches some rule $r_1' \in \mathbb{C}'$ that occurs before rule r'. This implies that for each conjunction $F_i \in S_i$ of the corresponding rule $r_1 \in \mathbb{C}$ that $L(p_i) \cap S_i \neq \emptyset$. However, this implies that $x_i \in S_i$ since if any point in $L(p_i)$ is in S_i , then all points in $L(p_i)$ are in S_i . It follows that x matches rule $r_1 \in \mathbb{C}$, contradicting our assumption that rule r was the first rule that x matches in \mathbb{C} . Thus, it follows that p' cannot match rule r_1' . It then follows that r' will be the first rule in \mathbb{C} that p' matches and the theorem follows.

VI. PREFIX ALIGNMENT

We now describe our *prefix alignment* approach. The basic idea is to "shift", "shrink", or "stretch" ranges by transforming the domain of each field to a new "prefix-friendly" domain so that the majority of the reencoded ranges either are prefixes or can be expressed by a small number of prefixes. This will reduce the costs of range expansion with perhaps a small penalty in rule width.

We first solve the special case where \mathbb{C} has only one field F. We develop an optimal solution using dynamic programming techniques. We then use this solution as a building block to perform prefix alignment on multi-dimensional classifiers. Finally, we compose domain compression and prefix alignment together.

A. Prefix Alignment Overview

The one-dimensional prefix alignment problem is equivalent to the following "cut" problem. Consider the three ranges [0, 12], [5, 15], and [0, 15] over domain $D(F_1) = [0, 15]$ in classifier \mathbb{C} in Figure 6(A), and suppose the transformed domain $D'(F_1) = [00, 11]$ in binary format. Because $D'(F_1)$ has a total of 4 elements, we want to identify three cut points $0 \le x_1 < x_2 < x_3 \le 15$ such that if $[0, x_1] \in D(F_1)$ transforms to $00 \in D'(F_1)$, $[x_1+1, x_2] \in D(F_1)$ transforms to $01 \in D'(F_1)$, $[x_2+1, x_3] \in D(F_1)$ transforms to $10 \in D'(F_1)$, and $[x_3 + 1, 15] \in D(F_1)$ transforms to $11 \in D'(F_1)$, the range expansion of the transformed ranges will have as few rules as possible. For this simple example, there are two families of optimal solutions: those with x_1 anywhere in [0, 3], $x_2 = 4$, and $x_3 = 12$, and those with $x_1 = 4$, $x_2 = 12$, and x_3 anywhere in [13, 15]. For the first family of solutions, range [0, 12] is transformed to $[00, 10] = 0 * \cup 10$, range [5, 15] is transformed to [10, 11] = 1*, and range [0, 15] is transformed to [00, 11] = **. In the second family of solutions, range [0, 12] is transformed to [00, 01] = 0*, range [5, 15]is transformed to $[01, 11] = 01 \cup 1*$, and range [0, 15]is transformed to [00, 11] = **. The classifier \mathbb{C}' in Figure 6(A) shows the three transformed ranges using the first family of solutions. In both examples, the range expansion of the transformed ranges only has 4 prefix rules while the range expansion of the original ranges has 7 prefix rules.



Fig. 6. Example of 1-D prefix alignment

We now illustrate how to compute an optimal solution using a divide and conquer strategy. We first observe that we can divide the original problem into two subproblems by choosing the middle cut point. We next observe that a cut point should be the starting or ending point of a range, if possible, in order to reduce range expansion. Suppose the target domain $D'(F_1)$ is $[0, 2^b - 1]$. We first need to choose the middle cut point $x_{2^{b-1}}$, which will divide the problem into two subproblems with target domains $[0, 2^{b-1} - 1] = 0 \{*\}^{b-1}$ and $[2^{b-1}, 2^{b} - 1] = 0 \{*\}^{b-1}$ $1 = 1 \{ * \}^{b-1}$ respectively. Consider the example in Figure 6(A), the x_2 cut point partitions [0, 15] into $[0, x_2]$, which transforms to prefix 0*, and $[x_2+1, 15]$, which transforms to prefix 1*. The second observation implies either $x_2 = 4$ or $x_2 = 12$. Suppose we choose $x_2 = 4$; that is, we choose the dashed line in Figure 6(A). This produces two subproblems where we need to identify the x_1 cut point in the range [0, 4]and the x_3 cut point in [5, 15]. In the two subproblems, we include each range trimmed to fit the restricted domain. For example, ranges [0, 12] and [0, 15] are trimmed to [0, 4] in the first subproblem. In the second subproblem, ranges [5, 15] and [0, 15] are trimmed to [5, 15] while range [0, 12] is trimmed to [5, 12]. We must maintain each trimmed range even if there may be duplicates. In the first subproblem, the choice of x_1 is immaterial since both trimmed ranges span the entire restricted domain. In the second subproblem, the range [5, 12] dictates that $x_3 = 12$ is the right choice.

We represent this divide and conquer process of computing cut points as a binary cut tree. Figure 6(B) depicts the tree where we select $x_2 = 4$ and $x_3 = 12$. This tree also encodes the transformation from the original domain to the target domain: all the values in a terminal node are mapped to the prefix represented by the path from the root to the terminal node. For example, as the path from the root to the terminal node of [0, 4] is 0, all values in $[0, 4] \in D(F_1)$ are transformed to 0*.

In domain compression, we considered transformers that mapped points in $D(F_i)$ to points in $D'(F_i)$. In prefix alignment, we consider transformers that map points in $D(F_i)$ to prefix ranges in $D'(F_i)$. If this is confusing, we can also work with transformers that map points in $D(F_i)$ to points in $D'(F_i)$ with no change in results; however, transformers that map to prefixes more accurately represent the idea of prefix alignment than transformers that map to points. Because we will perform range expansion on \mathbb{C}' before performing any further optimizations including redundancy removal, we can ignore rule order. We can then view a one-dimensional classifier \mathbb{C} as a multiset of ranges S in $D(F_1)$.

B. One-dimensional Prefix Alignment

We next present the technical details of our dynamic programming solution to the prefix alignment problem by addressing four issues.

1) Correctness of Prefix Alignment: We prove that prefix alignment preserves the semantics of the original classifier by first defining the concept of *prefix transformers* and then showing that prefix alignment must be correct when prefix transformers are used.

Given a prefix P, we use min P and max P to denote the smallest and the largest values in P, respectively.

Definition 6.1 (Prefix transformers): A transformer \mathbb{T}_i is an order-preserving prefix transformer from $D(F_i)$ to $D'(F_i)$ for a packet classifier \mathbb{C} if \mathbb{T}_i satisfies the following three properties. (1) (prefix property) $\forall x \in D(F_i), \mathbb{T}_i(x) = P$ where P is a prefix in domain $D'(F_i)$; (2) (order-preserving property) $\forall x, y \in D(F_i), x < y$ implies either $\mathbb{T}_i(x) = \mathbb{T}_i(y)$ or max $\mathbb{T}_i(x) < \min \mathbb{T}_i(y)$; (3) (consistency property) $\forall x, y \in$ $D(F_i), \mathbb{T}_i(x) = \mathbb{T}_i(y)$ implies $\mathbb{C}\{x\} = \mathbb{C}\{y\}$.

The following Lemma 6.1 and Theorem 6.1 easily follow from the definition of prefix transformers.

Lemma 6.1: Given any prefix transformer \mathbb{T}_i for a field F_i , for any $a, b, x \in D(F_i)$, $x \in [a, b]$ if and only if $\mathbb{T}_i(x) \subseteq [\min \mathbb{T}_i(a), \max \mathbb{T}_i(b)]$.

Theorem 6.1 (Prefix Alignment Theorem): Given a packet classifier \mathbb{C} over fields F_1, \dots, F_d , and d prefix transformers $T = \{\mathbb{T}_i \mid 1 \leq i \leq d\}$, and the classifier \mathbb{C}' constructed by replacing any range [a,b] over field F_i $(1 \leq i \leq d)$ by the range $[\min \mathbb{T}_i(a), \max \mathbb{T}_i(b)]$, the condition $\mathbb{C}(p_1, \dots, p_d) = \mathbb{C}'(\mathbb{T}_1(p_1), \dots, \mathbb{T}_d(p_d))$ holds.

2) Find Candidate Cut Points: We next identify candidate cut points using the concept of atomic ranges. For any multiset of ranges S (a multiset may have duplicate entries) and any range x over domain $D(F_1)$, we use S@x to denote the set of ranges in S that contain x.

Definition 6.2 (Atomic Range Set): Given a multiset S of ranges, the union of which constitute a range denoted $\bigcup S$,

and a set of ranges S', S' is the atomic range set of S if and only if the following four conditions hold: (1) (coverage property) $\bigcup S = \bigcup S'$; (2) (disjoint property) $\forall x, y \in S'$, $x \cap y = \emptyset$; (3) (atomicity property) $\forall x \in S$ and $\forall y \in S'$, $x \cap y \neq \emptyset$ implies $y \subseteq x$; (4) (maximality property) $\forall x, y \in S'$ and $\max x + 1 = \min y$ implies $S@x \neq S@y$.

For any multiset of ranges S, there is a unique atomic range set of S, which we denote as AR(S). Because of the maximality property of atomic range set, the candidate cut points correspond to the end points of ranges in AR(S). We now show how to compute S-start points and S-end points. For any range $[x, y] \in S$, define the points x - 1 and y to be S-end points, and define the points x and y + 1 to be S-start points. Note that we ignore x - 1 if x is the minimum element of $\bigcup S$ and y + 1 if y is the maximum element of $\bigcup S$. Let (s_1, \dots, s_m) and (e_1, \dots, e_m) be the ordered list of S-start points and S-end points. It follows that for $1 \le i \le m - 1$ that $s_i \le e_i = s_{i+1} - 1$. Thus, $AR(S) = \{[s_1, e_1], \dots, [s_m, e_m]\}$.

For example, if we consider the three ranges in classifier \mathbb{C} in example Figure 6(A), range [0, 12] creates S-start point 13 and S-end point 12, range [5, 15] creates S-end point 4 and S-start point 5, and range [0, 15] creates no S-start points or S-end points. Finally, 0 is an S-start point and 15 is an S-end point. This leads to $AR(S) = \{[0, 4], [5, 12], [13, 15]\}$.

3) Choose Target Domain Size: We next choose the number of bits b used to encode domain $D'(F_1)$. This value b imposes constraints on legal prefix transformers. Consider $S = \{[0,4], [0,7], [0,12], [0,15]\}$ with $AR(S) = \{[0,4], [5,7], [8,12], [13,15]\}$. If b = 2, then the only legal prefix transformer maps [0,4] to 00, [5,7] to 01, [8,12] to 10, and [13,15] to 11. If b = 3, there are many more legal prefix transformers including one that maps [0,4] to 000, [5,7] to 001, [8,12] to 01*, and [13,15] to 1 **. In this case, the second prefix transformer is superior to this first prefix transformer.

We include b as an input parameter to our prefix alignment problem. We initialize b as $\lceil \log_2 |AR(S)| \rceil$, the smallest possible value, and compute an optimal prefix alignment for this value of b. We then increment b and repeat until no improvement is seen. We choose a linear search as opposed to a binary search because computing the optimal solution for b bits requires an optimal solution for b - 1 bits.

4) Choose Optimal Cut Points: We now show how to compute the optimal cut points given b bits. We view a onedimensional classifier \mathbb{C} as a multiset of ranges S in $D(F_1)$ and formulate the prefix alignment problem as follows: Given a multiset of ranges S over field F_1 and a number of bits b, find prefix transformer \mathbb{T}_1 such that the range expansion of the transformed multiset of ranges S' has the minimum number of prefix rules and $D'(F_1)$ can be encoded using only b bits.

We present an optimal solution using dynamic programming. Given a multiset of ranges S, we first compute AR(S). Suppose there are m atomic ranges R_1, \dots, R_m with S-start points s_1 through s_m and S-end points e_1 through e_m sorted in increasing order. For any S-start point s_x and S-end point s_y where $1 \le x \le y \le m$, we define $S \cap [x, y]$ to be the multiset of ranges from S that intersect range $[s_x, s_y]$; furthermore, we assume that each range in $S \cap [x, y]$ is trimmed so that its start point is at least s_x and its end point is at most s_y . We then define a collection of subproblems as follows. For every $1 \le x \le y \le m$, we define a prefix alignment problem PA(x, y, b) where the problem is to find a prefix transformer \mathbb{T}_1 for $[s_x, e_y] \subseteq D(F_1)$ such that the range expansion of $(S \cap [x, y])'$ has the smallest possible number of prefix rules and the transformed domain $D'(F_1)$ can be encoded in b bits. We use cost(x, y, b) to denote the number of prefix rules in the range expansion of the optimal $(S \cap [x, y])'$. The original prefix alignment problem then corresponds to PA(1, m, b)where b can be arbitrarily large.

The prefix alignment problem obeys the optimal substructure property. For example, consider PA(1, m, b). As we employ the divide and conquer strategy to locate a middle cut point that will establish what the prefixes $0{*}^{b-1}$ and $1{*}^{b-1}$ correspond to, there are m-1 choices of cut points to consider: namely e_1 through e_{m-1} . Suppose the optimal cut point is e_k where $1 \le k \le m - 1$. Then the optimal solution to PA(1, m, b) will build upon the optimal solutions to subproblems PA(1, k, b-1) and PA(k+1, m, b-1). That is, the optimal transformer for PA(1, m, b) will simply append a 0 to the start of all prefixes in the optimal transformer for PA(1, k, b - 1) and a 1 to the start of all prefixes in the optimal transformer for PA(k+1, m, b-1). Moreover, cost(1, m, b) = cost(1, k, b - 1) + cost(k + 1, m, b - 1) -|S@[1,m]|. We subtract |S@[1,m]| in the above cost equation because ranges that include all of $[s_1, e_m]$ are counted twice, once in cost(1, k, b - 1) and once in cost(k + 1, m, b - 1). However, as $[s_1, e_k]$ transforms to $0\{*\}^{b-1}$ and $[s_{k+1}, e_m]$ transforms to $1\{*\}^{b-1}$, the range $[s_1, e_m]$ can be expressed by one prefix $\{*\}^b = 0\{*\}^{b-1} \cup 1\{*\}^{b-1}$.

Based on this analysis, Theorem 6.2 shows how to compute the optimal cuts and binary cut tree. As stated earlier, the optimal prefix transformer \mathbb{T}_1 can then be computed from the binary cut tree.

Theorem 6.2: Given a multiset of ranges S with |AR(S)| = m, cost(l,r,b) for any $b \ge 0, 1 \le l \le r \le m$ can be computed as follows. For any $1 \le l < r \le m$, and $b \ge 0$:



Note that we set cost(k, k, 0) to |S@[k, k]| for the convenience of the recursive case. The interpretation is that with a 0-bit domain, we can allow only a single value in $D'(F_1)$; this single value is sufficient to encode the transformation of an atomic interval.

C. Multi-Dimensional Prefix Alignment

We now consider multi-dimensional prefix alignment. Unfortunately, while we can optimally solve the one-dimensional problem, there are complex interactions between the dimensions that complicate the multi-dimensional problem. In particular, the total range expansion required for each rule is the product of the range expansion required for each field. Thus, there may be complex tradeoffs where we sacrifice one field of a rule but align another field so that the costs do not multiply. The complexity of the multi-dimensional prefix alignment problem is currently unknown.

We present a hill-climbing solution where we iteratively apply our one-dimensional prefix alignment algorithm one field at a time. Because the range expansion of one field affects the numbers of ranges that appear in the other fields, we run prefix alignment for each field more than once. We stop when running prefix alignment in each field fails to improve the solution. More precisely, for a classifier \mathbb{C} over fields F_1, \ldots, F_d , we first create d identity prefix transformers $\mathbb{T}_1^0, \ldots, \mathbb{T}_d^0$. We define a *multi-field prefix alignment iteration* k as follows. For i from 1 to d, generate the optimal prefix transformer \mathbb{T}_i^k assuming the prefix transformers for the other fields are $\{\mathbb{T}_1^{k-1}, \ldots, \mathbb{T}_{i-1}^{k-1}, T_{i+1}^{k-1}, \ldots, T_d^{k-1}\}$. Our iterative solution starts at k = 1 and preforms successive multi-field prefix alignment iterations until no improvement is found for any field.

D. Composing with Domain Compression

While domain compression and prefix alignment can be used individually, they can be easily combined to achieve superior compression. Given a classifier \mathbb{C} over fields F_1, \ldots, F_d , we first perform domain compression resulting in a transformed classifier \mathbb{C}' and d transformers $\mathbb{T}_1^{dc}, \ldots, \mathbb{T}_d^{dc}$; then, we perform prefix alignment on the classifier \mathbb{C}' resulting in a transformed classifier \mathbb{C}'' and d transformers $\mathbb{T}_1^{pa}, \ldots, \mathbb{T}_d^{pa}$. To combine the two transformation processes into one, we merge each pair of transformers \mathbb{T}_i^{dc} and \mathbb{T}_i^{pa} into one transformer \mathbb{T}_i for $1 \leq i \leq d$. We apply the optimal algorithm in [26] to compute the minimum possible transformers \mathbb{T}_i for 1 < i < d. When running prefix alignment after domain compression, computing the atomic ranges and candidate cut points is unnecessary because each point $x \in D'(F_i)$ for $1 \leq i \leq d$ belongs to its own equivalence class in $D'(F_i)$ which implies [x, x] is an atomic range.

VII. DISCUSSION

A. TCAM Update

We now discuss strategies for handling the TCAM updates. In most applications, such as router ACLs and firewalls, the rule sets are relatively static. Therefore, we propose using the bank mechanism in TCAMs to handle rule list updates. TCAMs are commonly configured into a series of row banks. Each bank can be individually enabled or disabled to determine whether or not its entries will be included in the TCAM search. We propose storing the compressed transformers and classifier before update in the active banks and the ones after

A	Range encoding scheme	Direct	direct range expansion
\mathbb{C}	packet classifier	$A(\mathbb{C})$	reencoded classifier
$W(A(\mathbb{C}))$	width of rules in $A(\mathbb{C})$	$ A(\mathbb{C}) $	# rules in $A(\mathbb{C})$
$TW(A(\mathbb{C}))$	minimum TCAM entry width for $A(\mathbb{C})$	$B(A(\mathbb{C}))$	$TW(A(\mathbb{C})) \times A(\mathbb{C}) $, <i>i.e.</i> , total bits of $A(\mathbb{C})$

Fig. 7. Summary of notation

update in the disabled banks. Once the writing is finished, we activate the banks containing the new transformers and compressed classifier and deactivate the banks containing the old ones. This technique takes advantage of the fact that each TCAM in the pipeline has uniform capacity that should be able to hold all the encoders and the encoded TCAM table. Therefore, there normally is adequate free space within the pipeline to store the original classifier and the updated classifier across each stage.

In some applications, there may be more frequent updates of the rule set. Fortunately, such updates are typically the insertion of new rules to the top or front of the classifier or the deletion of recently added rules. We are not aware of any applications that require frequent updates involving rules at arbitrary locations in a classifier. We can support this update pattern by chaining the TCAM chips in our proposed architecture after a small TCAM chip of normal width (160 bits), which we call the "hot" TCAM chip. When a new rule comes, we add the rule to the top of the hot TCAM chip. When a packet comes, we first use the packet as the key to search in the hot chip. If the packet has a match in the hot chip, then the decision of the first matching rule is the decision of the packet. Otherwise, we feed the packet to the TCAM chips in our architecture described as above to find the decision for the packet. Note that this method can support insertions of rules before the default rule by giving the default rule a unique decision and using the hot chip's match only if the default rule is matched.

Although the lookup on the hot TCAM chip adds a constant delay to per packet latency, throughput will not be affected because we use pipelining. Using batch updating, we only need to run our topological transformation algorithms to recompute the TCAM lookup tables when the hot chip is about to fill up. Note, we may not include specific rules when running topological transformation if they are likely to be deleted in the near future. Instead, we run topological transformation on the remainder of the classifier and retain these likely to be deleted rules in the hot TCAM chip.

B. Rule Logging

Packet classifiers sometimes allow rule logging; that is, recording the packets that match some particular rules. Our algorithm handles rule logging by assigning each rule that is logged a unique decision. Our experiments show that even when all rules in a classifier have unique decisions, our algorithm still achieves significant TCAM space reduction.

VIII. EXPERIMENTAL RESULTS

We evaluate the effectiveness and efficiency of our topological transformation approaches on both real-world and synthetic packet classifiers. Although our two approaches can be used independently, they are much more effective when used together. We primarily report results for both techniques used together. When a distinction is needed, we use the label DC + PA when reporting results obtained using both techniques combined and the label DC when reporting results obtained using only domain compression. In all cases, we preprocess each classifier by running the redundancy removal algorithm in [17].

A. Effectiveness

1) Evaluation Methodology: Given a TCAM range encoding algorithm A and a classifier \mathbb{C} , let $A(\mathbb{C})$ denote the reencoded classifier, $W(A(\mathbb{C}))$ denote the number of bits to represent each rule in $A(\mathbb{C})$, $TW(A(\mathbb{C}))$ denote the minimum TCAM entry width for storing $A(\mathbb{C})$ given choices 40, 80, 160, or 320, $|A(\mathbb{C})|$ denote the number of rules in $A(\mathbb{C})$, and $B(A(\mathbb{C})) = TW(A(\mathbb{C})) \times |A(\mathbb{C})|$, which represents the total number of TCAM bits required to store $A(\mathbb{C})$. The main goal of TCAM optimization algorithms is to minimize $B(A(\mathbb{C}))$. We use *Direct* to denote the direct range expansion algorithm, so $B(Direct(\mathbb{C}))$ represents the baseline we compare against, $W(Direct(\mathbb{C})) = 104$, $TW(Direct(\mathbb{C})) = 160$, and $B(Direct(\mathbb{C})) = 160 \times |Direct(\mathbb{C})|$. Figure 7 summarizes our notation.

For any A and \mathbb{C} , we measure overall effectiveness by the compression ratio $CR(A(\mathbb{C})) = \frac{B(A(\mathbb{C}))}{B(Direct(\mathbb{C}))}$. To isolate the factors that contribute to the success of our approaches at compressing classifiers, we define the *Rule Number Ratio* of A on \mathbb{C} to be $RNR(A(\mathbb{C})) = \frac{|A(\mathbb{C})|}{|\mathbb{C}|}$, which is often referred to as expansion ratio, and the *Rule Width Ratio* of A on \mathbb{C} to be $RWR(A(\mathbb{C})) = \frac{W(A(\mathbb{C}))}{104}$. When we consider a set of classifiers S where |S| denotes the number of classifiers in S, we generalize our metrics as follows. Average compression ratio of A for S is $CR(A(S)) = \frac{\sum_{\mathbb{C} \in S} CR(A(\mathbb{C}))}{|S|}$, average rule number ratio of A for S is $RNR(A(S)) = \frac{\sum_{\mathbb{C} \in S} RNR(A(\mathbb{C}))}{|S|}$, and average rule width ratio of A for S is $RWR(A(S)) = \frac{\sum_{\mathbb{C} \in S} RWR(A(\mathbb{C}))}{|S|}$.

We use RL to denote a set of 40 real-world packet classifiers that we performed experiments on. RL is chosen from a larger set of real-world classifiers obtained from various network service providers, where the classifiers range in size from a handful of rules to thousands of rules. We eliminated structurally similar classifiers from RL because similar classifiers exhibited similar results. We created RLby randomly choosing a single classifier from each set of structurally similar classifiers. We then split RL into two groups, RLa and RLb where $RNR(Direct(\mathbb{C})) \leq 4$ for all $\mathbb{C} \in RLa$ and $RNR(Direct(\mathbb{C})) > 40$ for all $\mathbb{C} \in RLb$. We have no classifiers where $4 < RNR(Direct(\mathbb{C})) \leq 40$. It turns out |RLa| = 26 and |RLb| = 14. By separating these classifiers into two groups, we can determine how well our techniques work on classifiers that do suffer significantly from range expansion as well as those that do not. Figure 8 shows the accumulated percentage graph of atomic intervals for each field for the classifiers in RL, and Figure 9 shows the accumulated percentage graphs of classifier sizes in RLbefore and after direct range expansion.

Because packet classifiers are considered confidential due to security concerns making it difficult to acquire a large number of real-world classifiers, we generated a set of synthetic classifiers SYN with the number of rules ranging from 250 to 8000 using Singh *et al.*'s [23] model of synthetic rules. The predicate of each rule has five fields: source IP, destination IP, source port, destination port, and protocol. We also performed experiments on TRS, a set of 490 classifiers produced by Taylor&Turner's Classbench [28]. These classifiers were generated using the parameter files downloaded from Taylor's web site http://www.arl.wustl.edu/~det3/ClassBench/index.htm. To represent a wide range of classifiers, we chose a uniform sampling of the allowed values for the parameters of smoothness, address scope, and application scope.

To stress test the sensitivity of our algorithms to the number of decisions in a classifier, we created a set of classifiers RL_U (and thus RLa_U and RLb_U) by replacing the decision of every rule in each classifier by a unique decision. Similarly, we created the set SYN_U . Thus, each classifier in RL_U (or SYN_U) has the maximum possible number of distinct decisions. Such classifiers might arise in the context of rule logging where the system monitors the frequency that each rule is the first matching rule for a packet.

2) Results on real-world and synthetic classifiers: Table II shows the average compression ratio, rule size ratio, and rule number ratio for our algorithm on all eight data sets. Figure 10 shows the accumulated percentage graphs for the compression ratios of our combined techniques for both RL and RL_U with and without transformers, and Figure 11 shows the accumulated percentage graphs for the compressions ratios of our combined techniques for both RL and RL_U with and without transformers, and Figure 11 shows the accumulated percentage graphs for the compressions ratios of our combined techniques for each field in RL. Note that the data with transformers depicts the true space savings of our methods, but most previous range encoding papers focus only on the data without transformers. Figure 12 and Figure 13 show the accumulated percentrage graphs of our combined techniques on RL and RL_U for rule number ratio and rule width ratio, respectively.

	compression			rule size	rule number		
	DC	w.o. T	with T		w.o. T	with T	
RL	11.8%	4.5%	13.8%	15.9%	36.1%	126.0%	
RL_U	29.6%	9.8%	20.8%	19.2%	77.0%	183.0%	
RLa	17.8%	6.8%	20.7%	20.4%	38.7%	105.2%	
RLa_U	44.6%	14.9%	31.3%	23.6%	82.7%	161.7%	
RLb	0.6%	0.1%	0.9%	7.5%	31.2%	164.7%	
RLb_U	1.7%	0.4%	1.1%	11.0%	66.4%	222.6%	
SYN	0.7%	0.6%	2.5%	10.4%	2.7%	11.8%	
SYN_U	13.4%	9.3%	12.4%	16.0%	43.9%	58.9%	
TRS	6.2%	1.0%	2.7%	15.7%	9.7%	23.3%	

TABLE II Average compression ratio, rule width ratio, and rule number ratio for our algorithm on 9 data sets. Average compression ratio and rule number ratios are reported with transformers included and excluded.

Our algorithm achieves significant compression on both



Fig. 12. Rule number ratios of RL and RL_U

Fig. 13. Rule width ratios of RL and RL_U

real-world and synthetic classifiers. On RL, our algorithm achieves an average compression ratio of 13.8% if we count TCAM space for transformers and 4.5% if we do not. These savings are attributable to both rule width and rule number compression. The average rule width compression ratio is 15.9%, which means that a typical encoded classifier only requires 17 bits, instead of 104 bits, to store a rule. However, the actual savings that rule width compression contributes to average compression ratio is only 25% because the encoded classifiers will use 40 bit wide TCAM entries, the smallest possible TCAM widths (two classifiers in RL_U require an 80 bit wide TCAM entry). In comparison, direct range expansion would use 160 bit wide TCAM entries. That is, $TW(A(\mathbb{C})) =$ 40 for all but two classifiers in RL_U . The remaining savings is due to rule number compression. Note that the average rule number compression ratio without transformers is 36.1%; that is, domain compression and redundancy removal eliminate an average of 63.9% of the rules from our real-life classifier sets. In comparison, the goal of all other reencoding schemes is an average rule number compression ratio without transformers of 100%. Our algorithm performs well on all of our other data sets too. For example, for Taylor's rule set TRS, we achieve an average compression ratio of 2.7% with transformers included and 1.0% with transformers excluded. Note that prefix alignment is an important component of our algorithm because it reduces the average compression ratio without transformers for RL from 11.8% to 4.5%.

3) Sensitivity to classifier efficiency: Our algorithm is effective for both efficiently specified classifiers and inefficiently specified classifiers. The efficiently specified classifiers in RLaexperience relatively little range expansion; the inefficiently specified classifiers in RLb experience significant range expansion. Not surprisingly, our algorithm provides roughly 20 times better compression for RLb than for RLa with average compression ratios of 0.9% and 20.7%, respectively. In both sets, TCAM width compression contributes approximately 25% savings. The difference is rule number compression. Whereas our algorithm achieves relatively similar average rule number ratios of 38.7% and 31.2% without transformers for RLa and RLb, respectively, these rule number ratios have significantly different impacts on the final compression ratios given that all the efficiently specified classifiers in RLa have modest range expansion while all the inefficiently specified classifiers in RLb have tremendous range expansion.

4) Sensitivity to number of unique decisions: Our algorithm's effectiveness is only slightly diminished as we increase the number of unique decisions in a classifier. In the extreme case where we assign each rule a unique decision in RL_U , our algorithm achieves an average compression ratio of 20.8% with transformers included and 9.8% with transformers excluded; and on SYN_U , our algorithm achieves an average compression ratio of 12.4% with transformers included and 9.3% with transformers excluded. In particular, the TCAM width used by each classifier is unaffected. Rule number ratio compression is worse for RL_U , but the rule number ratio without transformers is still less than 100% for all our data sets with unique decisions.

5) Comparison with state-of-the-art results: Our algorithm outperforms all existing reencoding schemes by at least a factor of 3.11 including transformers and by at least a factor of 5.54 excluding transformers. We first consider the width of TCAM entries. Our algorithm uses 40 bit wide TCAM entries for all but 2 classifiers in RL_U whereas the smallest TCAM width achieved by prior work is 80 bits [21], [22]. Therefore, on TCAM entry width, our algorithm is 2 times better than the best known result. Next, we consider the number of TCAM entries. Excluding TCAM entries for transformers, the best rule number ratio that any other method can achieve on RL is 100% whereas we achieve 36.1%. Therefore, excluding TCAM entries for transformers, our algorithm is at least $5.54 \ (= \ 2 \times 100\%/36.1\%)$ times better than the optimal TCAM reencoding algorithm that does not consider classifier semantics. In comparison with PIC [21], [22], the best previous TCAM-based reencoding algorithm, the transformers in PIC use at least the same number of TCAM entries as our algorithm because our domain compression technique may map multiple intervals to one decision whereas PIC maps each interval to a unique decision. Thus, including TCAM entries for transformers, the best average rule number ratio that PIC can achieve on RL is 195.7% (= 123.3% - 27.6% + 100%). Therefore, including TCAM entries for transformers, our algorithm is at least 3.11 (= $2 \times 195.7\%/126.0\%$) times better than PIC.

B. Efficiency

We implemented our algorithms on the Microsoft .Net framework 2.0 and performed our experiments on a desktop PC running Windows XP with 3G memory and a single 3.4 GHz Pentium D processor. On RL, the minimum, mean, median, and maximum running time is 0.003, 37.642, 0.079, and 1093.308 seconds; on RL_U , the minimum, mean, median, and maximum running time is 0.006, 1540.934, 0.203, and 54604.311 seconds. Table III shows running time of some representative classifiers in RL and RL_U . The last two classifiers in the table run slower due to the overhead of building the FDDs. This overhead can be alleviated by caching the FDDs for reuse during updates.

Size	Time (seconds)	Time (seconds) Unique Decisions
511	0.40	1.92
1308	1.00	6.59
1365	14.51	80.91
1794	26.85	273.35
2331	42.33	355.52
3928	0.64	4.86
4004	117.01	3234.90
7652	1093.31	54604.31

TABLE III RUNNING TIME ON 5 CLASSIFIERS IN RL and RL_U

IX. PERFORMANCE MODELING

We now assess the impact that our two topological transformation schemes (parallel pipelined-lookup using 6 TCAM chips and multi-lookup using 1 TCAM chip) will have on power, latency, and throughput. We compare our topological transformation schemes against direct range expansion. Because we cannot build actual devices, we use Agrawal and Sherwood's power, latency, and throughput models for TCAM chips [2]. To our best knowledge, Agrawal and Sherwood's TCAM models are the only publicly available models and have become widely adopted. To derive meaningful power results, we need much larger classifiers than the largest available classifier in RL. Rather than make large synthetic classifiers, we consider hypothetical classifiers whose direct range expansion fits exactly within standard TCAM chip sizes ranging from 1Mbit to 72bit. We further assume that when topological transformation is applied to these hypothetical classifiers, the resulting compression ratio will be 15%. To account for updates and because we do not know how the bits will be allocated to each of the 5 transformers and reencoded classifier, we conservatively assume that each transformer and the reencoded classifier will have a size that is 15% of the direct expansion classifier.

A. Model

For power and latency, we use Agrawal and Sherwood's TCAM modeling tool which takes as input a TCAM configuration (size in bits, row width, fabrication size, number of banks, and number of row dividers) and produces as output the power consumed for each access as well as the latency of each access. We compute throughput using the resulting latency. With respect to latency, the dominant factor is the priority encoder. Specifically, the depth of the priority encoder circuit increases as the number of banks searched increases. Thus, it is possible to perform searches on two small TCAMs with shallow priority encoder circuits in much less than double the time required to perform one search on a much larger TCAM. In summary, our modeling results demonstrate that the 6 chip configuration significantly improves throughput and power and is essentially neutral on latency whereas the 1 chip configuration significantly improves power while suffering some loss in latency and throughput.



Fig. 14. Power, latency, and throughput by size for 0.18 µm technology with 16 banks and 4 row dividers

	Average					
	Power		Latency		Throughput	
	6 Chip	1 Chip	6 Chip	1 Chip	6 Chip	1 Chip
1 Mb	122.0 %	20.6 %	100.7 %	304.6 %	200.4 %	32.8 %
2 Mb	91.7 %	15.6 %	92.0 %	300.0 %	238.1 %	33.3 %
4.5 Mb	66.8 %	11.5 %	100.0 %	342.9 %	233.3 %	29.2 %
9 Mb	50.3 %	8.8 %	112.5 %	375.0 %	200.0 %	26.7 %
18 Mb	40.5 %	7.2 %	97.0 %	317.4 %	226.8 %	31.5 %
36 Mb	34.8 %	6.3 %	63.5 %	205.2 %	341.1 %	48.7 %
72 Mb	31.8 %	5.8 %	32.3 %	103.4 %	662.9 %	96.7 %

 TABLE IV

 Power, latency, and throughput ratios

B. Power

For any classifier C, let $\mathbb{P}(A(C))$ represent the nanojoules needed to classify one packet using the given scheme. For the two topological transformation schemes, we include the power consumed by the transformers. For one classifier C, we define the power ratio of algorithm A as $\frac{\mathbb{P}(A(C))}{\mathbb{P}(Direct(C))}$. For a set of classifiers S, we define the average power ratio of algorithm Aover S to be $\frac{\sum_{C \in S} \frac{\mathbb{P}(A(C))}{|S|}}{|S|}$. The extrapolated average power ratios are displayed in Table IV and Figure 14 (a).

The modeling results clearly demonstrate that topological transformation results in a significant improvement in power usage per search. The reason for the savings is that even though we perform more searches, each search is performed on a much smaller TCAM chip.

C. Latency

For any classifier C, let $\mathbb{L}(A(C))$ represent the total number of nanoseconds required to classify a packet using the given scheme. For topological transformation with 6 chips, this is the time required to perform 5 lookups in parallel for the transformer tables plus the final lookup on the reencoded classifier. For topological transformation with one chip, it is the time time required to perform all 6 lookups sequentially. For one classifier C, we define the latency ratio of algorithm Aas $\frac{\mathbb{L}(A(C))}{\mathbb{L}(Direct(C))}$. For a set of classifiers S, we define the average latency ratio for algorithm A over S to be $\frac{\sum_{C \in S} \frac{\mathbb{L}(A(C))}{\|S\|}}{|S|}$. The extrapolated average latencies are included in Figure 14 (b) and Table IV.

The modeling results demonstrate that topological transformation does not significantly increase latency, particularly if we use the 6 chip configuration. In fact, this configuration can even reduce latency. As we discussed earlier, the reason is that even though we must perform some searches sequentially, each search is faster because it is performed on a much smaller TCAM chip with a much shallower priority encoder. As a result, even the 1 chip configuration is only 2-4 times slower than direct range expansion.

D. Throughput

For any classifier C, let $\mathbb{T}(A(C))$ represent the number of packets per second that can be classified using the given scheme. For topological transformation with 6 chips, this is the minimum throughput of any of the 6 TCAM chips. For topological transformation with 1 TCAM chip, this is essentially the inverse of latency because there is no pipelining. For one classifier C, we define the throughput ratio of algorithm A as $\frac{\mathbb{T}(A(C))}{\mathbb{T}(Direct(C))}$. For a set of classifiers S, we define the average throughput ratio for algorithm A over S to be $\frac{\sum_{C \in S} \frac{\mathbb{T}(A(C))}{\mathbb{T}(Direct(C))}}{|S|}$. The extrapolated average throughputs are included in Figure 14 (c) and Table IV.

The modeling results demonstrate that topological transformation significantly improves throughput if we use the 6 chip configuration. The reason for the throughput increase is the use of the pipeline and the use of smaller and thus faster TCAM chips. The throughput of the 1 chip configuration is significantly reduced because there is no pipeline; however, the throughput is better than 16.6% because it again uses smaller, faster TCAM chips.

X. CONCLUSIONS AND FUTURE WORK

We make three major contributions in this paper. First, we propose a novel topological view of the TCAM reencoding process where we consider the semantics of the packet classifier. Second, we present two techniques, domain compression and prefix alignment, for realizing such a view. These techniques are not only composable, they can be composed with other TCAM optimization and reencoding schemes. Third, we implemented our algorithms and conducted extensive experiments on both real-life and synthetic packet classifiers. The experimental results show that our techniques achieve at least 5.54 times more space reduction with transformers excluded and at least 3.11 times more space reduction with transformers included.

Our work opens up new problems for future research. One problem is to find an optimal choice of landmarks for each equivalence class in the domain compression technique that leads to the smallest final classifier. Another is to find an optimal solution to the multi-dimensional prefix alignment problem or prove it is NP-hard. We also plan to study more potential combinations of our techniques with other TCAM optimization and reencoding schemes.

Acknowledgement

This material is based in part upon work supported by the National Science Foundation under Grant Numbers CNS-0716407 and CNS-0916044. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- A guide to search engines and networking memory. http://www.linleygroup.com/pdf/NMv4.pdf.
- [2] B. Agrawal and T. Sherwood. Modeling tcam power for next generation network devices. In *Proceedings of the IEEE International Symposium* on *Performance Analysis of Systems and Software*, pages 120–129, 2006.
- [3] D. A. Applegate, G. Calinescu, D. S. Johnson, H. Karloff, K. Ligett, and J. Wang. Compressing rectilinear pictures and minimizing access control lists. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, January 2007.
- [4] A. Bremler-Barr, D. Hay, D. Hendler, and R. Roth. Layered interval codes for tcam based classification. In *Proceedings of the IEEE Infocom*, 2009.
- [5] A. Bremler-Barr and D. Hendler. Space-efficient TCAM-based classification using gray coding. In Proceedings of the 26th Annual IEEE Conference on Computer Communications (Infocom), May 2007.
- [6] H. Che, Z. Wang, K. Zheng, and B. Liu. DRES: Dynamic range encoding scheme for tcam coprocessors. *IEEE Transactions on Computers*, 57(7):902–915, July 2008.
- [7] Q. Dong, S. Banerjee, J. Wang, D. Agrawal, and A. Shukla. Packet classifiers in ternary CAMs can be smaller. In *Proceedings of the ACM Sigmetrics*, pages 311–322, 2006.
- [8] R. Draves, C. King, S. Venkatachary, and B. Zill. Constructing optimal IP routing tables. In *Proceedings of the IEEE INFOCOM*, pages 88–97, 1999.
- [9] M. G. Gouda and A. X. Liu. Firewall design: consistency, completeness and compactness. In *Proceedings of the 24th IEEE International Conference on Distributed Computing Systems (ICDCS-04)*, pages 320– 327, March 2004.
- [10] P. Gupta and N. McKeown. Packet classification on multiple fields. In Proceedings of the ACM SIGCOMM, pages 147–160, 1999.
- [11] P. Gupta and N. McKeown. Algorithms for packet classification. *IEEE Network*, 15(2):24–32, 2001.
- [12] T. V. Lakshman and D. Stiliadis. High-speed policy-based packet forwarding using efficient multi-dimensional range matching. In *Proceedings of the ACM SIGCOMM*, pages 203–214, 1998.
- [13] K. Lakshminarayanan, A. Rangarajan, and S. Venkatachary. Algorithms for advanced packet classification with ternary CAMs. In *Proceedings* of the ACM SIGCOMM, pages 193 – 204, August 2005.
- [14] A. X. Liu and M. G. Gouda. Diverse firewall design. In Proceedings of the International Conference on Dependable Systems and Networks (DSN-04), pages 595–604, June 2004.
- [15] A. X. Liu and M. G. Gouda. Complete redundancy detection in firewalls. In Proceedings of the 19th Annual IFIP Conference on Data and Applications Security, LNCS 3654, pages 196–209, August 2005.
- [16] A. X. Liu and M. G. Gouda. Complete redundancy removal for packet classifiers in tcams. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, in press.
- [17] A. X. Liu, C. R. Meiners, and Y. Zhou. All-match based complete redundancy removal for packet classifiers in TCAMs. In *Proceedings* of the 27th Annual IEEE Conference on Computer Communications (Infocom), April 2008.

- [18] H. Liu. Efficient mapping of range classifier into Ternary-CAM. In Proceedings of the Hot Interconnects, pages 95–100, 2002.
- [19] C. R. Meiners, A. X. Liu, and E. Torng. TCAM Razor: A systematic approach towards minimizing packet classifiers in TCAMs. In *Proceed*ings of the 15th IEEE Conference on Network Protocols (ICNP), pages 266–275, October 2007.
- [20] C. R. Meiners, A. X. Liu, and E. Torng. Bit weaving: A non-prefix approach to compressing packet classifiers in TCAMs. Technical Report MSU-CSE-09-1, Department of Computer Science and Engineering, Michigan State University, January 2009.
- [21] D. Pao, Y. Li, and P. Zhou. Efficient packet classification using TCAMs. *Computer Networks*, 50(18):3523–3535, 2006.
- [22] D. Pao, P. Zhou, B. Liu, and X. Zhang. Enhanced prefix inclusion coding filter-encoding algorithm for packet classification with ternary content addressable memory. *IET Computers & Digital Techniques*, 1(5):572– 580, September 2007.
- [23] S. Singh, F. Baboescu, G. Varghese, and J. Wang. Packet classification using multidimensional cutting. In *Proceedings of the ACM SIGCOMM*, pages 213–224, 2003.
- [24] E. Spitznagel, D. Taylor, and J. Turner. Packet classification using extended TCAMs. In *Proceedings of the 11th IEEE International Conference on Network Protocols (ICNP)*, pages 120–131, November 2003.
- [25] V. Srinivasan, G. Varghese, S. Suri, and M. Waldvogel. Fast and scalable layer four switching. In *Proceedings of the ACM SIGCOMM*, pages 191–202, 1998.
- [26] S. Suri, T. Sandholm, and P. Warkhede. Compressing two-dimensional routing tables. *Algorithmica*, 35:287–300, 2003.
- [27] D. E. Taylor. Survey & taxonomy of packet classification techniques. ACM Computing Surveys, 37(3):238–275, 2005.
- [28] D. E. Taylor and J. S. Turner. Classbench: A packet classification benchmark. In *Proceedings of the IEEE Infocom*, March 2005.
- [29] J. van Lunteren and T. Engbersen. Fast and scalable packet classification. *IEEE Journals on Selected Areas in Communications*, 21(4):560–571, 2003.
- [30] F. Yu, T. V. Lakshman, M. A. Motoyama, and R. H. Katz. SSA: A power and memory efficient scheme to multi-match packet classification. In *Proceedings of the Symposium on Architectures for Networking and Communications Systems (ANCS)*, pages 105–113, October 2005.
- [31] K. Zheng, H. Che, Z. Wang, B. Liu, and X. Zhang. DPPC-RE: TCAMbased distributed parallel packet classification with range encoding. *IEEE Transactions on Computers*, 55(8):947–961, August 2006.



Chad R. Meiners received his Ph.D. in Computer Science at Michigan State University in 2009. He is currently a research associate in the Department of Computer Science and Engineering at Michigan State University. His research interests include networking, algorithms, and security.



Alex X. Liu received his Ph.D. degree in computer science from the University of Texas at Austin in 2006. He is currently an assistant professor in the Department of Computer Science and Engineering at Michigan State University. He received the IEEE & IFIP William C. Carter Award in 2004 and an NSF CAREER award in 2009. His research interests focus on networking, security, and dependable systems.

Eric Torng received his Ph.D. degree in computer science from Stanford University in 1994. He is currently an associate professor and graduate director in the Department of Computer Science and Engineering at Michigan State University. He received an NSF CAREER award in 1997. His research interests include algorithms, scheduling, and networking.

