ClassBench: A Packet Classification Benchmark

David E. Taylor and Jonathan S. Turner, *Fellow, IEEE*

***Abstract*—Packet classification is an enabling technology for next generation network services and often a performance bottleneck in high-performance routers. The performance and capacity of many classification algorithms and devices, including TCAMs, depend upon properties of filter sets and query patterns. Despite the pressing need, no standard performance evaluation tools or filter sets are publicly available. In response to this problem, we present *ClassBench*, a suite of tools for benchmarking packet classification algorithms and devices. *ClassBench* includes a *Filter Set Generator* that produces synthetic filter sets that accurately model the characteristics of real filter sets. Along with varying the size of the filter sets, we provide high-level control over the composition of the filters in the resulting filter set. The tool suite also includes a *Trace Generator* that produces a sequence of packet headers to exercise packet classification algorithms with respect to a given filter set. Along with specifying the relative size of the trace, we provide a simple mechanism for controlling locality of reference. While we have already found *ClassBench* to be very useful in our own research, we seek to eliminate the significant access barriers to realistic test vectors for researchers and initiate a broader discussion to guide the refinement of the tools and codification of a formal benchmarking methodology. (The *ClassBench* tools are publicly available at the following site: http://www.arl.wustl.edu/~det3/ClassBench/.)**

***Index Terms*—Communication systems, computer network performance, packet switching, packet classification.**

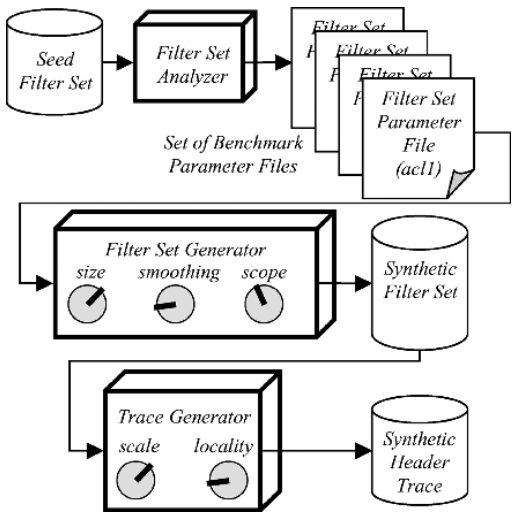

Fig. 1. Block diagram of the *ClassBench* tool suite. The synthetic *Filter Set Generator* has size, smoothing, and scope adjustments which provide high-level, systematic mechanisms for altering the size and composition of synthetic filter sets. The set of benchmark *parameter files* model real filter sets and may be refined over time. The *Trace Generator* provides adjustments for trace size and locality of reference.

I. Introduction

DEPLOYMENT of next generation network services hinges on the ability of Internet infrastructure to perform packet classification at physical link speeds. A packet classifier must compare header fields of every incoming packet against a set of filters in order to assign a flow identifier that is used to apply security policies, application processing, and quality-of-service guarantees. Typical packet classification filter sets have fewer than a thousand filters and reside in enterprise firewalls or edge routers. As network services continue to migrate into the network core, it is anticipated that filter sets could swell to tens of thousands of filters or more. The most common type of packet classification examines the packet header fields comprising the standard IP 5-tuple. A packet classifier searches for the highest priority filter or set of filters matching the packet where each filter specifies a prefix of the IP source and destination addresses, an exact match or wildcard for the transport protocol number, and ranges for the source and destination port numbers for TCP and UDP packets.

As reported in Section III, it has been observed that real filter sets exhibit a considerable amount of structure. In response, several algorithmic techniques have been developed which exploit filter set structure to accelerate search time or reduce storage requirements [1]. Consequently, the performance of these approaches are subject to the structure of filter sets. Likewise, the capacity and efficiency of the most prominent packet classification solution, ternary content addressable memory (TCAM), is also subject to the characteristics of filter sets [1]. Despite the influence of filter set composition on the performance of packet classification search techniques and devices, no publicly available benchmarking tools or filter sets exist for standardized performance evaluation. Due to security and confidentiality issues, access to large, real filter sets has been limited to a small subset of the research community. Some researchers in academia have gained access to filter sets through confidentiality agreements, but are unable to distribute those filter sets. Furthermore, performance evaluations using real filter sets are restricted by the size and structure of the sample filter sets.

In order to facilitate future research and provide a foundation for a meaningful benchmark, we present *ClassBench*, a publicly available suite of tools for benchmarking packet classification algorithms and devices. As shown in Fig. 1, *ClassBench* consists of three tools: a *Filter Set Analyzer*, *Filter Set Generator*, and *Trace Generator*. The general approach of *ClassBench* is to

Manuscript received March 29, 2005; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor M. Crovella. This work was supported in part by the National Science Foundation under Grant ANI-9813723.

D. E. Taylor is with Exegy Inc., St. Louis, MO 63118 USA (e-mail: david@exegy.com).

J. S. Turner is with the Applied Research Laboratory, Washington University, Saint Louis, MO 63130 USA (e-mail: jon.turner@wustl.edu).

Digital Object Identifier 10.1109/TNET.2007.893156

construct a set of benchmark *parameter files* that specify the relevant characteristics of real filter sets, generate a synthetic filter set from a chosen *parameter file* and a small set of high-level inputs, and generate a sequence of packet headers to probe the synthetic filter set using the *Trace Generator*. *Parameter files* contain various statistics and probability distributions that guide the generation of synthetic filter sets. The *Filter Set Analyzer* tool extracts the relevant statistics and probability distributions from a seed filter set and generates a *parameter file*. This provides the capability to generate large synthetic filter sets which model the structure of a seed filter set. In Section IV, we discuss the statistics and probability distributions contained in the *parameter files* that drive the synthetic filter generation process.

The *Filter Set Generator* takes as input a *parameter file* and a few high-level parameters. In addition the filter set *size* parameter, the *smoothing* and *scope* parameters provide high-level control over the composition of the filter set, abstracting the user from the low-level statistics and distributions contained in the *parameter files*. The *smoothing* adjustment provides a structured mechanism for introducing new address aggregates which is useful for modeling filter sets significantly larger than the filter set used to generate the *parameter file*. The *scope* adjustment provides a biasing mechanism to favor more or less specific filters during the generation process. These adjustments and their affects on the resulting filter sets are discussed in Section V. Finally, the *Trace Generator* tool examines the synthetic filter set, then generates a sequence of packet headers to exercise the filter set. Like the *Filter Set Generator*, the trace generator provides adjustments for scaling the size of the trace as well as the locality of reference of headers in the trace. These adjustments are described in detail in Section VI.

We highlight previous performance evaluation efforts by the research community as well as related benchmarking activity of the IETF in Section II. It is our hope that this work initiates a broader discussion which will lead to refinement of the tools, compilation of a standard set of *parameter files*, and codification of a formal benchmarking methodology. Its value will depend on its perceived clarity and usefulness to the interested community.

- *Researchers* seeking to evaluate new classification algorithms relative to alternative approaches and commercial products.
- *Classification product vendors* seeking to market their products with convincing performance claims over competing products.
- *Classification product customers* seeking to verify and compare classification product performance on a uniform scale.[1]

II. Related Work

Extensive work has been done in developing benchmarks for many applications and data processing devices. Benchmarks are used extensively in the field of computer architecture to evaluate microprocessor performance. In the field of computer communications, the Internet Engineering Task Force (IETF) has several working groups exploring network performance measurement. Specifically, the IP Performance Metrics (IPPM) working group was formed with the purpose of developing standard metrics for Internet data delivery [2]. The Benchmarking Methodology Working Group (BMWG) seeks to make measurement recommendations for various internetworking technologies [3]. These recommendations address metrics and performance characteristics as well as collection methodologies.

The BMWG specifically attacked the problem of measuring the performance of forwarding information base (FIB) routers [4] and also produced a methodology for benchmarking firewalls [5]. The methodology contains broad specifications such as: the firewall should contain at least one rule for each host, tests should be run with various filter set sizes, and test traffic should correspond to rules at the "end" of the filter set. *ClassBench* complements efforts by the IETF by providing the necessary tools for generating test vectors with high-level control over filter set and input trace composition. The Network Processor Forum (NPF) has also initiated a benchmarking effort [6]. Currently, the NPF has produced benchmarks for switch fabrics and route lookup engines. To our knowledge, there are no current efforts by the IETF or the NPF to provide a benchmark for multiple field packet classification.

In the absence of publicly available packet filter sets, researchers have exerted much effort in order to generate realistic performance tests for new algorithms. Several research groups obtained access to real filter sets through confidentiality agreements. Gupta and McKeown obtained access to 40 real filter sets and extracted a number of useful statistics which have been widely cited [7]. Feldmann and Muthukrishnan composed filter sets based on *NetFlow* packet traces from commercial networks [8]. Several groups have generated synthetic 2-D filter sets consisting of source-destination address prefix pairs by randomly selecting address prefixes from publicly available route tables [8]–[10]. Baboescu and Varghese also generated synthetic 2-D filter sets by randomly selecting prefixes from publicly available route tables, but added refinements for controlling the number of zero-length prefixes (wildcards) and prefix nesting [11], [12]. A simple technique for appending randomly selected port ranges and protocols from real filter sets in order to generate synthetic five-dimensional filter sets is also described [11]. Baboescu and Varghese also introduced a scheme for using a sample filter set to generate a larger synthetic five-dimensional filter set [13]. This technique replicates filters by changing the IP prefixes while keeping the other fields unchanged. While these techniques address some aspects of scaling filter sets in size, they lack high-level mechanisms for adjusting filter set composition which is crucial for evaluating algorithms that exploit filter set characteristics.

Woo provided strong motivation for a packet classification benchmark and initiated the effort by providing an overview of filter characteristics for different environments (ISP Peering Router, ISP Core Router, Enterprise Edge Router, etc.) [14]. Based on high-level characteristics, Woo generated large synthetic filter sets, but provided few details about how the filter sets were constructed. The technique also does not provide controls for varying the composition of filters within the filter set. Nonetheless, his efforts provide a good starting point for con-

[1]In order to facilitate broader discussion, we make the *ClassBench* tools and 12 *parameter files* publicly available at the following site: http://www.arl.wustl.edu/~det3/ClassBench/.

structing a benchmark capable of modeling various application environments for packet classification. Sahasranaman and Buddhikot used the characteristics compiled by Woo in a comparative evaluation of a few packet classification techniques [15].

Stanford's Packet Lookup and Classification Simulator (PALAC) [16] tools provide a framework for comparative performance evaluation of various IP lookup and packet classification algorithms. The Classifier Description Language (CDL) module of PALAC gerenates a synthetic route table or filter set based on input parameters controlling the number of filters and the number of fields per filter. Alternatively, PALAC allows IPMA table snapshots to be used for algorithm evaluation. PALAC also includes traffic generation, statistics collection, and classifier update modules. The *ClassBench* tool suite may be used in conjunction with frameworks such as PALAC to explore the effects of filter set size and composition on packet classifier performance.

III. Analysis of Real Filter Sets

Recent efforts to identify better packet classification techniques have focused on leveraging the characteristics of real filter sets for faster searches. While lower bounds for the general multi-field searching problem have been established, observations made in recent packet classification work offer enticing new possibilities to provide significantly better performance. The focus of this section is to identify and understand the impetus for the observed structure of filter sets and to develop metrics and characterizations of filter set structure that aid in generating synthetic filter sets. We performed a battery of analyses on 12 real filter sets provided by Internet Service Providers (ISPs), a network equipment vendor, and other researchers working in the field. The filter sets range in size from 68 to 4557 entries and utilize one of the following formats: access control list (ACL), firewall (FW), and IP chain (IPC). Due to confidentiality concerns, the filter sets were provided without supporting information regarding the types of systems and environments in which they are used. We are unable to comment on "where" in the network architecture the filter sets are used. Nonetheless, the following analysis provide useful insight into the structure of real filter sets. We observe that various useful properties hold regardless of filter set size or format. Due to space constraints, we are unable to fully elaborate on our analysis, but a more complete discussion of this work is available in technical report form [17].

A. Understanding Filter Composition

Many of the observed characteristics of filter sets arise due to the administrative policies that drive their construction. The most complex packet filters typically appear in firewall and edge router filter sets due to the heterogeneous set of applications supported in these environments. Firewalls and edge routers typically implement security filters and network address translation (NAT), and they may support additional applications such as virtual private networks (VPNs) and resource reservation. Typically, these filter sets are created manually by a system administrator using a standard management tool such as CiscoWorks VPN/Security Management Solution (VMS) [18] and Lucent Security Management Server (LSMS) [19]. Such tools conform to a model of filter construction which views a filter as specifying the communicating subnets and the application or set of applications. Hence, we can view each filter as having two major components: an address prefix pair and an application specification. The address prefix pair identifies the communicating subnets by specifying a source address prefix and a destination address prefix. The application specification identifies a specific application session by specifying the transport protocol, source port number, and destination port number. A set of applications may be identified by specifying ranges for the source and destination port numbers.

TABLE I
Distribution of Filters Over the Five Port Classes for Source and Destination Port Range Specifications; Values Given as Percentage (%) of Filters in the Filter Set

Port	*WC*	*HI*	*LO*	*AR*	*EM*
Source	78.08	6.60	0.92	0.42	13.99
Destination	40.39	6.18	0.06	4.33	49.04

B. Application Specifications

We analyzed the application specifications in the 12 filter sets in order to corroborate previous observations as well as extract new, potentially useful characteristics.

1) Protocol: For each of the filter sets, we examined the unique protocol specifications and the distribution of filters over the set of unique values. Filters specified one of nine protocols or the wildcard. The most common protocol specification was TCP (49%), followed by UDP (27%), the wildcard (13%), and ICMP (10%). The following protocols were specified by less than 1% of the filters: General Routing Encapsulation (GRE), Open Shortest Path First (OSPF) Interior Gateway Protocol (IGP), Enhanced Interior Gateway Routing Protocol (EIGRP), IP Encapsulating Security Payload (ESP) for IPv6, IP Authentication Header (AH) for IPv6, IP Encapsulation within IP (IPE).

2) Port Ranges: Next, we examined the port ranges specified by filters in the filter sets and the distribution of filters over the unique values. In order to observe trends among the various filter sets, we define five classes of port ranges:

- WC: wildcard;
- HI: ephemeral user port range [1024:65535];
- LO: well-known system port range [0:1023];
- AR: arbitrary range;
- EM: exact match.

Motivated by the allocation of port numbers, the first three classes represent common specifications for a port range. The last two classes may be viewed as partitioning the remaining specifications based on whether or not an exact port number is specified. We computed the distribution of filters over the five classes for both source and destination ports for each filter set. Table I shows the combined distribution for all filter sets. We observe some interesting trends in the raw data. With rare exception, the filters in the ACL filter sets specify the wildcard for the source port. A majority of filters in the ACL filters specify an exact port number for the destination port. Source port specifications in the other filter sets are also dominated by the wildcard, but a considerable portion of the filters specify an exact port number. Destination port specifications in the

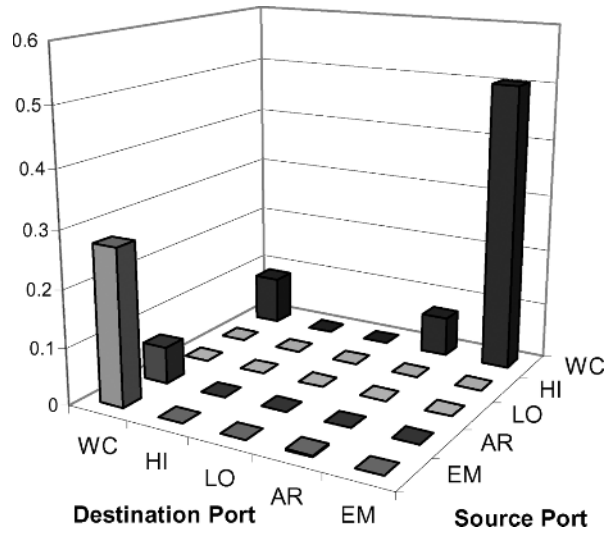

Fig. 2. *Port Pair Class Matrix* for TCP, filter set fw4.

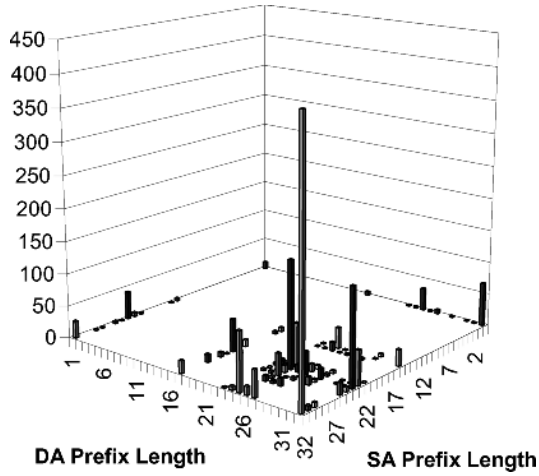

Fig. 3. Prefix length distribution for address prefix pairs in filter set ipc1.

other filter sets share the same trend, however the distribution between the wildcard and exact match is a bit more even. Only one filter set contained filters specifying the LO port class for either the source or destination port range.

3) Port Pair Class: As previously discussed, the structure of source and destination port range pairs is a key point of interest for both modeling real filter sets and designing efficient search algorithms. We can characterize this structure by defining a *Port Pair Class* (PPC) for every combination of source and destination port class. For example, WC-WC if both source and destination port ranges specify the wildcard, AR-LO if the source port range specifies an arbitrary range and the destination port range specifies the set of well-known system ports. As shown in Fig. 2, a convenient way to visualize the structure of *Port Pair Classes* is to define a *Port Pair Class Matrix* where rows share the same source port class and columns share the same destination port class. For each filter set, we examined the *PPC Matrix* defined by filters specifying the same protocol. For all protocols except TCP and UDP, the *PPC Matrix* is trivial—a single spike at WC/WC. Fig. 2 shows the *PPC Matrix* defined by filters specifying the TCP protocol in filter set fw4.

C. Address Prefix Pairs

A filter identifies communicating hosts or subnets by specifying a source and destination address prefix, or address prefix pair. The speed and efficiency of several longest prefix matching and packet classification algorithms depend upon the number of unique prefix lengths and the distribution of filters across those unique values. We find that a majority of the filter sets specify fewer than 15 unique prefix lengths for either source or destination address prefixes. The number of unique source/destination prefix pair lengths is typically less than 32, which is small relative to the filter set size and the number of possible combinations, 1024. For example, the largest filter set contained 4557 filters, 11 unique source address prefix lengths, 3 unique destination address lengths, and 31 unique source/destination prefix pair lengths.

Next, we examine the distribution of filters over the unique address prefix pair lengths. Note that this study is unique in that previous studies and models of filter sets utilized independent distributions for source and destination address prefixes. Real filter sets have unique prefix pair distributions that reflect the types of filters contained in the filter set. For example, fully specified source and destination addresses dominate the distribution for filter set *ipc1* shown in Fig. 3. There are very few filters specifying a 24-bit prefix for either the source or destination address, a notable difference from backbone route tables which are dominated by class C address prefixes (24-bit network address) and their aggregates. Finally, we observe that while the distributions for different filter sets are sufficiently different from each other a majority of the filters in the filter sets specify prefix pair lengths around the "edges" of the distribution. This implies that, typically, one of the address prefixes is either fully specified or wildcarded.

By considering the prefix pair distribution, we characterize the *size* of the communicating subnets specified by filters in the filter set. Next, we would like to characterize the relationships among address prefixes and the amount of address space covered by the prefixes in the filter set. Consider a binary tree constructed from the IP source address prefixes of all filters in the filter set. From this tree, we could completely characterize the data structure by determining a conditional branching probability for each node. For example, assume that an address prefix is generated by traversing the tree starting at the root node. At each node, the decision to take to the 0 path or the 1 path exiting the node depends upon the branching probability at the node. As shown in Fig. 4, $p\{0|11\}$ is the probability that the 0 path is chosen at level 2 given that the 1 path was chosen at level 0 and the 1 path was chosen at level 1. Such a characterization is overly complex, hence we employ suitable metrics that capture the important characteristics while providing a more concise representation.

We begin by constructing two binary tries from the source and destination prefixes in the filter set. Note that there is one level in the tree for each possible prefix length 0 through 32 for a total of 33 levels. For each level in the tree, we compute the probability that a node has one child or two children. Nodes

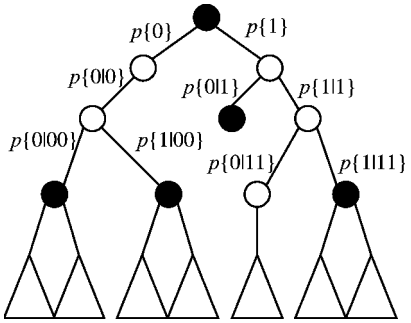

Fig. 4. Example of complete statistical characterization of address prefixes.

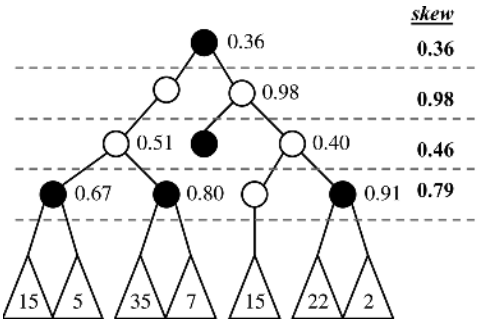

Fig. 5. Example of skew computation for the first four levels of an address trie. Shaded nodes denote a prefix specified by a single filter. Subtrees denoted by triangles with associated weight.

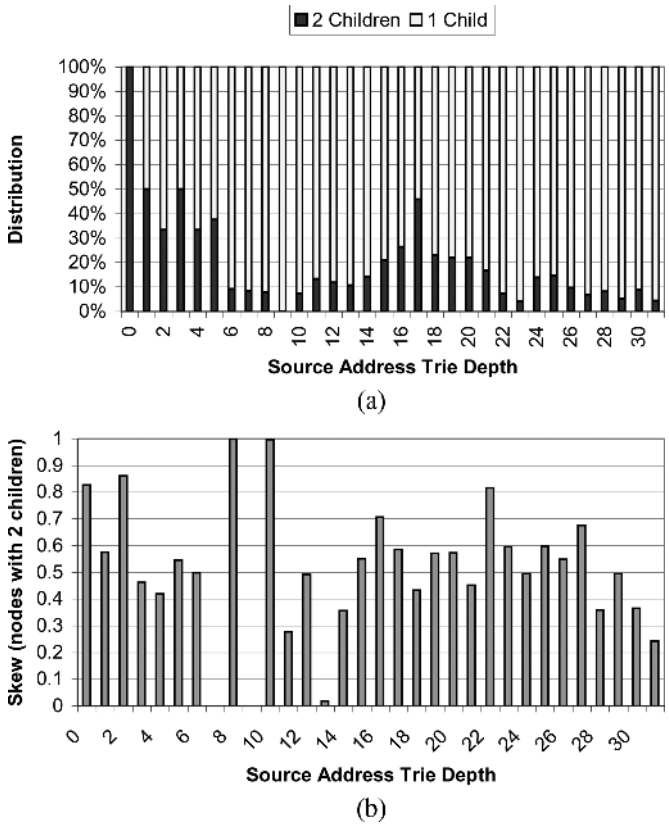

Fig. 6. Source address branching probability and skew for filter set acl5. (a) Source address branching probability; average per level. (b) Source address skew; average per level for nodes with two children.

with no children are excluded from the calculation. We refer to this distribution as the *Branching Probability*. For nodes with two children, we compute *skew*, which is a relative measure of the "weights" of the left and right subtrees of the node. Subtree weight is defined to be the number of filters specifying prefixes in the subtree, not the number of prefixes in the subtree. This definition of weight accounts for popular prefixes that occur in many filters. Let *heavy* be the subtree with the largest weight and let *light* be the subtree with equal or less weight, thus

$$\text{skew} = 1 - \frac{\text{weight(light)}}{\text{weight(heavy)}}. \tag{1}$$

Consider the following example: given a node k with two children at level m, assume that 10 filters specify prefixes in the 1-subtree of node k (the subtree visited if the next bit of the address is 1) and 25 filters specify prefixes in the 0-subtree of node k. The 1-subtree is the *light* subtree, the 0-subtree is the *heavy* subtree, and the skew at node k is 0.6. We compute the average skew for all nodes with two children at level m, record it in the distribution, and move on to level $(m+1)$. We provide and example of computing skew for the first four levels of an address trie in Fig. 5.

The result of this analysis is two distributions for each address trie, a *branching probability* distribution and a *skew* distribution. We plot these distributions for the source address prefixes in filter set acl5 in Fig. 6. In Fig. 6(a), note that a significant portion of the nodes in levels zero through five have two children, but the amount generally decreases as we move down the trie. The increase at level 16 and 17 is a notable exception. This implies that there is a considerable amount of branching near the "top" of the trie, but the paths generally remain contained as we move down the trie. In Fig. 6(b), we observe that skew for nodes with two children hovers around 0.5, thus the one subtree tends to contain prefixes specified by twice as many filters as the other subtree. Note that skew is not defined at levels where all nodes have one child. Also note that levels containing nodes with two children may have an average skew of zero (completely balanced subtrees), but this is rare. Finally, this definition of skew provides an anonymous measure of address prefix structure, as it does not preserve address prefix values.

Branching probability and skew characterize the structure of the individual source and destination address prefixes; however, it does not capture their interdependence. It is possible that some filters in a filter set match flows contained within a single subnet, while others match flows between different subnets. In order to capture this characteristic of a seed filter set, we measure the "correlation" of source and destination prefixes. In this context, we define correlation to be the probability that the source and destination address prefixes continue to be the same for a given prefix length. This measure is only valid within the range of address bits specified by both address prefixes. Additional details regarding the "correlation" metric and results from real filter sets may be found in [17].

D. Scope

Next we seek to characterize the *specificity* of the filters in the filter set. Filters that are more specific cover a small set of

possible packet headers while filters that are less specific cover a large set of possible packet headers. The number of possible packet headers covered by a filter is characterized by its *tuple* specification. To be specific, we consider the standard 5-tuple as a vector containing the following fields:

- $t[0]$: source address prefix length $[0 \ldots 32]$;
- $t[1]$: destination address prefix length $[0 \ldots 32]$;
- $t[2]$: source port range width, the number of port numbers covered by the range $[0 \ldots 2^{16}]$;
- $t[3]$: destination port range width, the number of port numbers covered by the range $[0 \ldots 2^{16}]$;
- $t[4]$: protocol specification, Boolean value denoting whether or not a protocol is specified [0, 1].

We define a new metric, *scope*, to be the logarithmic measure of the number of possible packet headers covered by the filter. Using the definition above, we define a filter's 5-tuple *scope* as follows:

$$\begin{aligned}\text{scope} &= \lg\left\{\left(2^{32-t[0]}\right) \times \left(2^{32-t[1]}\right)\right. \\ &\qquad \left. \times t[2] \times t[3] \times \left(2^{8(1-t[4])}\right)\right\} \\ &= (32 - t[0]) + (32 - t[1]) + (\lg t[2]) \\ &\quad + (\lg t[3]) + 8\,(1 - t[4]) \end{aligned} \tag{2}$$

Thus, *scope* is a measure of filter specificity on a scale from 0 to 104. The average 5-tuple scope for our 12 filter sets ranges from 56 to 24. We note that filters in the ACL filter sets tend to have narrower scope, while filters in the FW filter sets tend to have wider scope.

E. Additional Fields

An examination of real filter sets reveals that additional fields beyond the standard 5-tuple are relevant. In 10 of the 12 filter sets that we studied, filters contain matches on TCP flags or ICMP type numbers. In most filter sets, a small percentage of the filters specify a nonwildcard value for the flags, typically less then two percent. There are notable exceptions, as approximately half the filters in filter set *ipc1* contain nonwildcard flags. We argue that new services and administrative policies will demand that packet classification techniques scale to support additional fields beyond the standard 5-tuple. Matches on ICMP type number and other higher-level header fields are likely to be exact matches. There may be other types of matches that more naturally suit the application, such as arbitrary bit masks on TCP flags.

IV. Parameter Files

Given a real filter set, the *Filter Set Analyzer* generates a *parameter file* that contains statistics and probability distributions that allow the *Filter Set Generator* to produce a synthetic filter set that retains the relevant characteristics of the original filter set. We chose the statistics and distributions to include in the *parameter file* based on thorough analysis of 12 real filter sets and several iterations of the *Filter Set Generator* design. Note that *parameter files* also provide complete anonymity of addresses in the original filter set. By reducing confidentiality concerns, we seek to remove the significant access barriers to realistic test vectors for researchers and promote the development of a benchmark set of *parameter files*. There still exists a need for a large sample space of real filter sets from various application environments. We have generated a set of 12 *parameter files* which are publicly available along with the *ClassBench* tool suite.

Parameter files include the following entries.[2]

- *Protocol* specifications and the distribution of filters over those values.
- *Port Pair Class Matrix* for each unique protocol specification in the filter set
- *Flags* specifications for each protocol and a distribution of filters over those values.
- *Arbitrary port range* specifications and a distribution of filters over those values for both the source and destination port fields.
- *Exact port number* specifications and a distribution of filters over those values for both the source and destination port fields.
- *Prefix pair length* distribution for each *Port Pair Class Matrix*.
- *Address prefix branching and skew* distributions for both source and destination address prefixes.
- *Address prefix correlation* distribution.
- *Prefix nesting thresholds* for both source and destination address prefixes.

Parameter files represent prefix pair length distributions using a combination of a total prefix length distribution and source prefix length distributions for each specified total length[3] as shown in Fig. 7. The total prefix length is simply the sum of the prefix lengths for the source and destination address prefixes. As we will demonstrate in Section V-B, modeling the total prefix length distribution allows us to easily bias the generation of more or less specific filters based on the *scope* input parameter. The source prefix length distributions associated with each specified total length allow us to model the prefix pair length distribution, as the destination prefix length is simply the difference of the total length and the source length.

The number of unique address prefixes that match a given packet is an important property of real filter sets and is often referred to as *prefix nesting*. We found that if the *Filter Set Generator* is ignorant of this property, it is likely to create filter sets with significantly higher prefix nesting, especially when the synthetic filter set is larger than the filter set used to generate the *parameter file*. Given that prefix nesting remains relatively constant for filter sets of various sizes, we place a limit on the prefix nesting during the filter generation process. The *Filter Set Analyzer* computes the maximum prefix nesting for both the source and destination address prefixes in the filter set and records these statistics in the *parameter file*. The *Filter Set Generator* retains these prefix nesting properties in the synthetic filter set, regardless of size. We discuss the process of generating address prefixes and retaining prefix nesting properties in Section V.

[2]We avoid an exhaustive discussion of *parameter file* contents and format details; interested readers and potential users of *ClassBench* may find a discussion of *parameter file* format in the documentation provided with the tools.

[3]We do not need to store a source prefix distribution for total prefix lengths that are not specified by filters in the filter set.

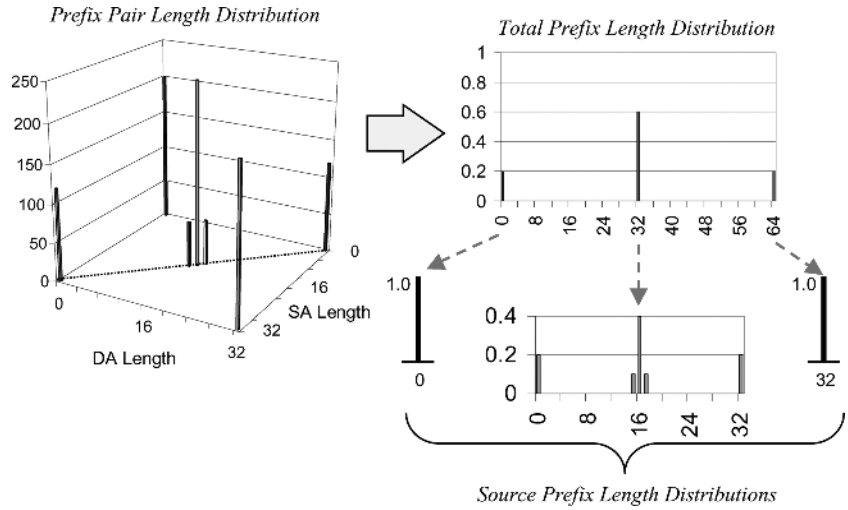

Fig. 7. *Parameter files* represent prefix pair length distributions using a combination of a total prefix length distribution and source prefix length distributions for each nonzero total length.

V. Synthetic Filter Set Generation

The *Filter Set Generator* is the cornerstone of the *ClassBench* tool suite. Perhaps the most succinct way to describe the synthetic filter set generation process is to walk through the pseudocode shown in Fig. 8. The first step in the filter generation process is to read the statistics and distributions from the *parameter file*. Next, we get the four high-level input parameters:

- *size*: target size for the synthetic filter set;
- *smoothing*: controls the number of new address aggregates (prefix lengths);
- *port scope*: biases the tool to generate more or less specific port range pairs;
- *address scope*: biases the tool to generate more or less specific address prefix pairs.

We refer to the *size* parameter as a "target" size because the generated filter set may have fewer filters. This is due to the fact that it is possible for the *Filter Set Generator* to produce a filter set containing redundant filters, thus the final step in the process removes the redundant filters. The generation of redundant filters stems from the way the tool assigns source and destination address prefixes that preserve the properties specified in the *parameter file*. This process will be described in more detail in a moment.

Before we begin the generation process, we apply the *smoothing* adjustment to the prefix pair length distributions[4] (lines 6–10). In order to apply the *smoothing* adjustment, we must iterate over all Port Pair Classes (line 7), apply the adjustment to each total prefix length distribution (line 8) and iterate over all total prefix lengths (line 9), and apply the adjustment to each source prefix length distribution associated with the total prefix length (line 10). We discuss this adjustment and its effects on the generated filter set in Section V-A.

The next set of steps (lines 12–27) generate a *partial* filter for each entry in the `Filters` array. Essentially, we assign all filter fields except the address prefix values. Note that the prefix lengths for both source and destination address *are* assigned. The reason for this approach will become clear when we discuss the assignment of address prefix values in a moment. The first step in generating a *partial* filter is to select a protocol from the `Protocols` distribution (line 14) using a uniform random variable, `rv` (line 13). We chose to select the protocol first because we found that the protocol specification dictates the structure of the other filter fields. Next, we select the protocol flags[5] from the `Flags` distribution associated with the chosen protocol (line 16).

After choosing the protocol and flags, we select a *Port Pair Class*, PPC, from the *Port Pair Class Matrix*, `PPCMatrix`, associated with the chosen protocol (line 18). Note that the selection of the PPC is performed with a random variable that is biased by the *port scope* parameter (line 17). This adjustment allows the user to bias the *Filter Set Generator* to produce a filter set with more or less specific PPCs, where WC-WC (both port ranges wildcarded) is the least specific and EM-EM (both port ranges specify an exact match port number) is the most specific. We discuss this adjustment and its effects on the generated filter set in Section V-B. Given the PPC, we can select the source and destination port ranges from their respective port range distributions associated with each port class (lines 20 and 22). Note that the distributions for port classes WC, HI, and LO are trivial as they define single ranges.

Selecting the address prefix pair lengths is the last step in generating a *partial* filter. We select a total prefix pair length from the distribution associated with the chosen PPC (line 24) using a random variable biased by the *address scope* parameter (line 23). We select a source prefix length from the distribution associated with the chosen PPC and total length (line 26) using a uniform random variable (line 25). Finally, we calculate the destination address prefix length using the chosen total length and source address prefix length (line 27).

After we generate all the *partial* filters, we must assign the source and destination address prefix values. The `AssignSA`

[4]Note that the *scope* adjustments do not add any new prefix lengths to the distributions. It only changes the likelihood that longer or shorter prefix lengths in the distribution are chosen.

[5]Note that the protocol flags field is typically the wildcard unless the chosen protocol is TCP or ICMP.

```
FilterSetGenerator()
   // Read input file and parameters
1  read(parameter file)
2  get(size)
3  get(smoothing)
4  get(port scope)
5  get(address scope)
   // Apply smoothing to prefix pair lengths
6  If smoothing > 0
7     For i : 1 to MaxPortPairClass
8         TotalLengths[i]→smooth(smoothing)
9         For j : 0 to 64
10            SALengths[i][j]→smooth(smoothing)
   // Allocate temporary filter array
11 FilterType Filters[size]
   // Generate partial filters
12 For i : 1 to size
       // Choose an application specification
13     rv = Random()
14     Filters[i].Prot = Protocols→choose(rv)
15     rv = Random()
16     Filters[i].Flags =
         Flags[Filters[i].Prot]→choose(rv)
17     rv = RandomBias(port scope)
18     PPC =
         PPCMatrix[Filters[i].Prot]→choose(rv)
19     rv = Random()
20     Filters[i].SP =
         SrcPorts[PPC.SPClass]→choose(rv)
21     rv = Random()
22     Filters[i].DP =
         DstPorts[PPC.DPClass]→choose(rv)
       // Choose an address prefix pair length
23     rv = RandomBias(address scope)
24     TotalLength =
         TotalLengths[PPC]→choose(rv)
25     rv = Random()
26     Filters[i].SALength =
         SrcLengths[PPC][TotalLength]→choose(rv)
27     Filters[i].DALength =
         TotalLength - Filters[i].SALength
   // Assign address prefix pairs
28 AssignSA(Filters)
29 AssignDA(Filters)
   // Remove redundant filters
30 RemoveRedundantFilters(Filters)
   // Prevent filter nesting
31 OrderNestedFilters(Filters)
32 PrintFilters(Filters)
```

Fig. 8. Pseudocode for *Small Filter Set Generator*.

routine recursively constructs a binary trie using the set of source address prefix lengths in `Filters` and the source address branching probability and skew distributions specified by the *parameter file* (line 28). The recursive process first examines all of the entries in `FilterList`. If an entry has a source prefix length equal to the level of the node, it assigns the node's address to the entry and removes the entry from `FilterList`. The process then distributes the remaining filters to child nodes according to the branching probability and skew for the node's level. Note that we also keep track of the number of prefixes that have been assigned along a path and ensure that the prefix nesting threshold is not exceeded.

Assigning destination address prefix values is symmetric to the process for source address prefixes with one extension. In order to preserve the relationship between source and destination address prefixes in each filter, the `AssignDA` process (line 29) also considers the correlation distribution specified in the *parameter file*. In order to preserve the correlation, `AssignDA` employs a two-phase process of constructing the destination address trie. The first phase recursively distributes filters according to the correlation distribution. When the address prefixes of a particular filter cease to be correlated, it stores the filter in a temporary `StubList` associated with the current tree node. The second phase recursively walks down the tree and completes the assignment process in the same manner as the `AssignSA` process, with the exception that the `StubList` is appended to the `FilterList` passed to the `AssignDA` process prior to processing. Additional details regarding the address prefix assignment process are included in [17].

Note that we do not explicitly prevent the *Filter Set Generator* from generating redundant filters. Identical *partial* filters may be assigned the same source and destination address prefix values by the `AssignSA` and `AssignDA` functions. In essence, this preserves the characteristics specified by the *parameter file* because the number of unique filter field values allowed by the various distributions is inherently limited. Consider the example of attempting to generate a large filter set using a *parameter file* from a small filter set. If we are forced to generate the number of filters specified by the *size* parameter, we face two unfavorable results: 1) the resulting filter set may not model the *parameter file* because we are repeatedly forced to choose values from the tails of the distributions in order to create unique filters or 2) the *Filter Set Generator* never terminates because it has exhausted the distributions and cannot create any more unique filters. With the current design of the *Filter Set Generator*, a user can produce a larger filter set by simply increasing the *size* target beyond the desired size. While this does introduce some variability in the size of the synthetic filter set, we believe this is a tolerable trade-off to make for maintaining the characteristics in the *parameter file* and achieving reasonable execution times for the *Filter Set Generator*.

Thus, after generating a list of *size* synthetic filters, we remove any redundant filters from the list via the `RemoveRedundantFilters` function (line 30). A naïve implementation of this function would require $O(N^2)$ time, where N is equal to $size$. We discuss an efficient mechanism for removing redundant filters from the set in Section V-C. After removing redundant filters from the filter set, we sort the filters in order of increasing scope (line 31). This allows the filter set to be searched using a simple linear search technique, as nested filters will be searched in order of decreasing specificity. An efficient technique for performing this sorting step is also discussed in Section V-C.

A. Smoothing Adjustment

As filter sets scale in size, we anticipate that new address prefix pair lengths will emerge due to subnet aggregation and

segregation. In order to model this behavior, we provide for the introduction of new prefix lengths in a structured manner. Injecting purely random address prefix pair lengths during the generation process neglects the structure of the filter set used to generate the *parameter file*. Using scope as a measure of distance, subnet aggregation and segregation results in new prefix lengths that are "near" to the original prefix length. Consider the address prefix pair length distribution where all filters in the filter set have 16-bit source and destination address prefixes; thus, the distribution is a single "spike." In order to model aggregation and splitting of subnets, new prefix pair lengths should be clustered around the existing spike in the distribution. This structured approach translates "spikes" in the distribution into smoother "hills;" hence, we refer to the process as smoothing.

In order to control the injection of new prefix lengths, we define a *smoothing* parameter which limits the maximum radius of deviation from the original prefix pair length, where radius is measured in the number of bits specified by the prefix pair. Geometrically, this measurement may be viewed as the Manhattan distance from one prefix pair length to another. For convenience, let the *smoothing* parameter be equal to r. We chose to model the clustering using a symmetric binomial distribution. Given the parameter r, a symmetric binomial distribution is defined on the range $[0 : 2r]$, and the probability at each point i in the range is given by

$$p_i = \binom{2r}{i}\left(\frac{1}{2}\right)^{2r}. \tag{3}$$

Note that r is the median point in the range with probability p_r, and r may assume values in the range [0:64]. Once we generate the symmetric binomial distribution from the *smoothing* parameter, we apply this distribution to each specified prefix pair length. The smoothing process involves scaling each "spike" in the distribution according to the median probability p_r, and binomially distributing the residue to the prefix pair lengths within the r-bit radius. When prefix lengths are at the "edges" of the distribution, we simply truncate the binomial distribution. This requires us to normalize the prefix pair length distribution as the last step in the smoothing process.

In order to demonstrate this process, Fig. 9 shows the prefix pair length distribution for a synthetic filter set generated with a *parameter file* specifying 16-bit prefix lengths for all addresses and a smoothing parameter $r = 8$. In practice, we expect that the *smoothing* parameter will be limited to at most 8. In order to demonstrate the effect of smoothing on a real filter set, Fig. 10 shows the prefix pair length distribution for a synthetic filter set of 64 000 filters generated using the ipc1 *parameter file* and smoothing parameter $r = 4$. Note that this synthetic filter set retains the structure of the original filter set shown in Fig. 3 while modeling a realistic amount of address prefix aggregation and segregation.

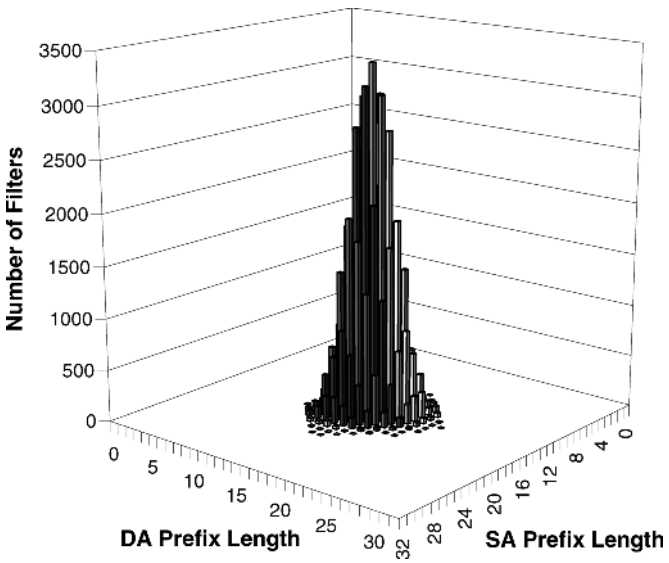

Fig. 9. Prefix pair length distributions for a synthetic filter set of 64 000 filters generated with a *Parameter File* specifying 16-bit prefix lengths for all addresses and smoothing parameter $r = 8$.

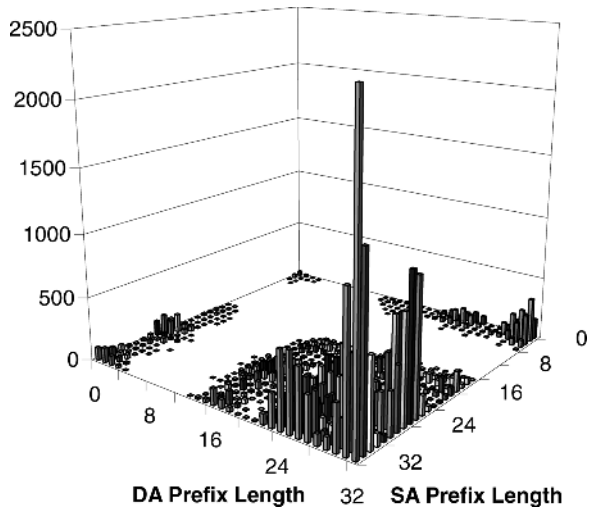

Fig. 10. Prefix pair length distribution for a synthetic filter set of 64 000 filters generated with the ipc1 *Parameter File* with smoothing parameter $r = 4$.

B. Scope Adjustment

As filter sets scale in size and new applications emerge, it is likely that the average scope of the filter set will change. As the number of flow-specific filters in a filter sets increases, the average scope decreases. If the number of explicitly blocked ports for all packets in a firewall filter set increases, then the average scope may increase.[6] In order to explore the performance effects of filter scope, we provide high-level adjustments of the average scope of the synthetic filter set. Two input parameters, *address scope* and *port scope*, allow the user to bias the *Filter Set Generator* to create more or less specific address prefix pairs and port pairs, respectively.

In order to sample from a cumulative distribution, we typically choose a random number uniformly distributed between zero and one, rv_{uni}, then chooses the value covering rv_{uni} in the cumulative distribution. Graphically, this amounts to projecting a horizontal line from the random number on the y axis. The chosen value is the x coordinate of the intersection of the cumulative distribution and the y projection of the random number.

[6]We are assuming a common practice of specifying an exact match on the blocked port number and wildcards for all other filter fields

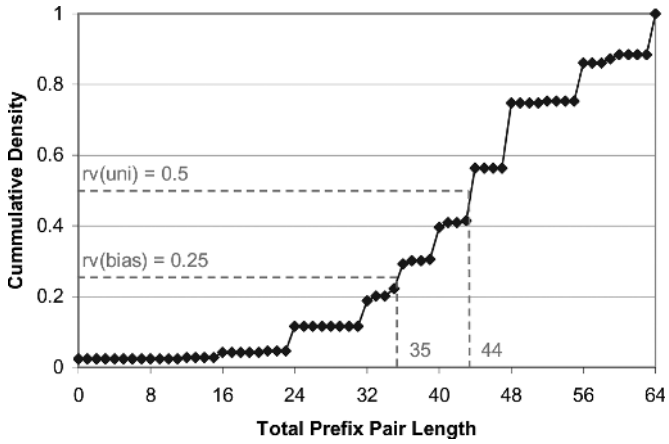

Fig. 11. Example of sampling from a cumulative distribution using a uniform random variable, and a biased random variable. Distribution is for the total prefix pair length associated with the WC-WC port pair class of the acl2 filter set.

In Fig. 11, we shown an example of sampling from a cumulative total prefix pair length distribution with $rv_{\text{uni}} = 0.5$ to choose the total prefix pair length of 44. The *scope* adjustments bias the sampling process to select more or less specific *Port Pair Classes* and prefix pair lengths. We can realize this in two ways: 1) apply the adjustment to the cumulative distribution or 2) bias the random variable used to sample from the cumulative distribution. Consider the case of selecting prefix pair lengths. The first option requires that we recompute the cumulative distribution to make longer or shorter total prefix lengths more or less probable, as dictated by the *address scope* parameter. The second option provides a conceptually simpler alternative. Returning to the example in Fig. 11, if we want to bias the *Filter Set Generator* to produce more specific address prefix pairs, then we want the random variable used to sample from the distribution to be biased to values closer to 1. The reverse is true if we want less specific address prefix pairs. Thus, in order to apply the scope adjustment we simply use a random number generator to choose a uniformly distributed random variable rv_{uni}, apply a biasing function to generate a biased random variable rv_{bias} and sample from the cumulative distribution using rv_{bias}.

While there are many possible biasing functions, we limit ourselves to a particularly simple class of functions. Our chosen biasing function may be viewed as applying a slope s to the uniform distribution as shown in Fig. 12(a). When the slope $s = 0$, the distribution is uniform. The biased random variable corresponding to a uniform random variable on the x axis is equal to the area of the rectangle defined by the value and a line intersecting the y axis at one with a slope of zero. Thus, the biased random variable is equal to the uniform random variable. We can bias the random variable by altering the slope of the line. In order for the biasing function to have a range of [0:1] for random variables in the range [0:1], the slope adjustment must be in the range [−2:2]. For convenience, we define the scope adjustments to be in the range [−1:1], thus the slope is equal to two times the scope adjustment. For nonzero slope values, the biased random variable corresponding to a uniform random variable on the x axis is equal to the area of the trapezoid defined by the value and a line intersecting the point (0.5,1) with a slope

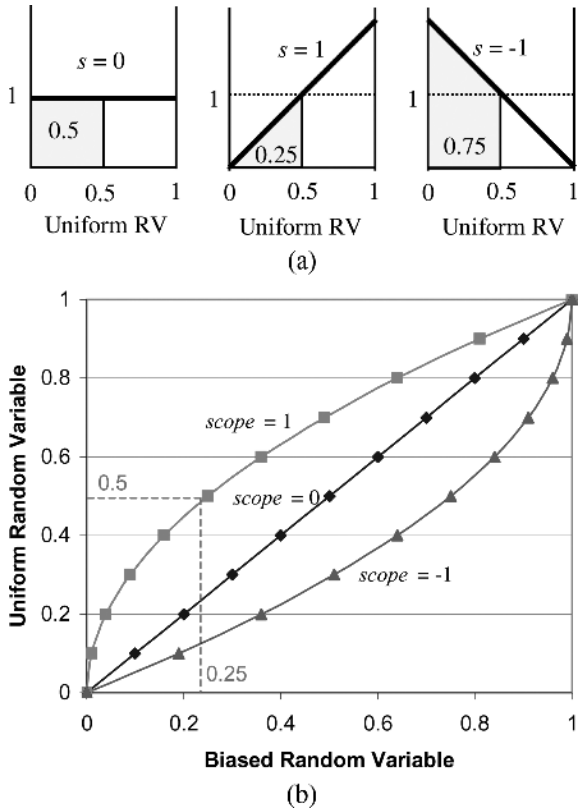

Fig. 12. Scope applies a biasing function to a uniform random variable. (a) Biased random variable is defined by area under line with slope $s = 2\times$scope. (b) Plot of scope biasing function.

of s. The expression for the biased random variable rv_{bias} given a uniform random variable rv_{uni} and a *scope* parameter in the range [−1:1] is

$$rv_{\text{bias}} = rv_{\text{uni}}(\text{scope} \times rv_{\text{uni}}\text{scope} + 1). \tag{4}$$

Fig. 12(b) shows a plot of the biasing function for *scope* values of 0, −1 and 1, as well as an example of computing the biased random variable given a uniform random variable of 0.5 and a *scope* parameter of 1. In this case the rv_{bias} is 0.25. Let us return to the example of choosing the total address prefix length from the cumulative distribution. In Fig. 11, we also show an example of sampling the distribution using the biased random variable, $rv_{\text{bias}} = 0.25$, resulting from applying the biasing function with $scope = 1$. The biasing results in the selection of a less specific address prefix pair, a total length of 35 as opposed to 44.

Positive values of *address scope* bias the *Filter Set Generator* to choose less specific address prefix pairs, thus increasing the average scope of the filter set. Likewise, negative values of *address scope* bias the *Filter Set Generator* to choose more specific address prefix pairs, thus decreasing the average scope of the filter set. The same effects are realized by the *port scope* adjustment by biasing the *Filter Set Generator* to select more or less specific *Port Pair Classes*.

Finally, we show the results of tests assessing the effects of the *address scope* and *port scope* parameters on the synthetic filter sets generated by the *Filter Set Generator* in Fig. 13. Each data point in the plot is from a synthetic filter set containing 16 000

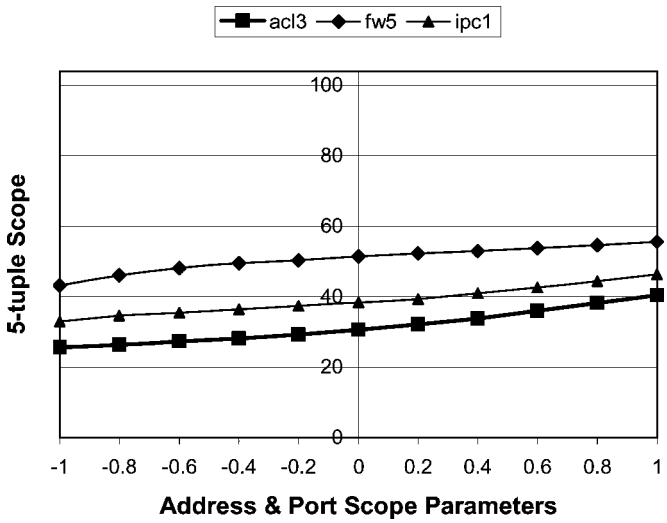

Fig. 13. Average scope of synthetic filter sets consisting of 16 000 filters generated with parameter files extracted from filter sets *acl3*, *fw5*, and *ipc1*, and various values of the scope parameters.

```
TraceGenerator()
    // Generate synthetic packet headers
1   read(FilterSet)
2   get(scale)
3   get(ParetoA)
4   get(ParetoB)
5   Threshold = scale × size(FilterSet)
6   HeaderList Headers()
7   While size(Headers) < Threshold
8         RandFilt = randint(0,size(FilterSet))
9         NewHeader =
            RandomCorner(RandFilt,FilterSet)
10        Copies = Pareto(ParetoA,ParetoB)
11        For i : 1 to Copies
12            Headers→append(NewHeader)
13  Headers→print
```

Fig. 14. Pseudocode for *Trace Generator*.

filters generated from a *parameter file* from filter sets acl3, fw5, or ipc1. For these tests, both scope parameters were set to the same value. Over their range of values, the scope parameters alter the average filter scope by ± 6–± 7.5. We also measured the individual effects of the *address scope* and *port scope* parameters. Over its range of values, the *address scope* alters the average address pair scope by ± 4–± 6. Over its range of values, the *port scope* alters the average port pair scope by ± 1.5–± 2.5. These scope adjustments provide a convenient high-level mechanism for exploring the effects of filter specificity on the performance of packet classification algorithms and devices.

C. Filter Redundancy and Priority

The final steps in synthetic filter set generation are removing redundant filters and ordering the remaining filters in order of increasing scope. The removal of redundant filters may be realized by simply comparing each filter against all other filters in the set; however, this naïve implementation requires $O(N^2)$ time. Such an approach makes execution times of the *Filter Set Generator* prohibitively long for filter sets with more than a few thousand filters. In order to accelerate this process, we first sort the filters into sets according to their tuple specification. We perform this sorting efficiently by constructing a binary search tree of tuple set pointers, using the scope of the tuple as the key for the node. When adding a filter to a tuple set, we search the set for redundant filters. If no redundant filters exist in the set, then we add the filter to the set. If a redundant filter exists in the set, we discard the filter. The time complexity of this search technique depends on the number of tuples created by filters in the filter set and the distribution of filters across the tuples. In practice, we find that this technique provides acceptable performance.

In order to support the traditional linear search technique, filter priority is often inferred by placement in an ordered list. In such cases, the first matching filter is the best matching filter. This arrangement could obviate a filter f_i if a less specific filter $f_j \supset f_i$ occupies a higher position in the list. To prevent this, we order the filters in the synthetic filter set according to scope, where filters with minimum scope occur first. The binary search tree of tuple set pointers makes this ordering task simple. Recall that we use scope as the node key. Thus, we simply perform an in-order walk of the binary search tree, appending the filters in each tuple set to the output list of filters.

VI. Trace Generation

When benchmarking a particular packet classification algorithm or device, many of the metrics of interest such as storage efficiency and maximum decision tree depth may be garnered using the synthetic filter sets generated by the *Filter Set Generator*. In order to evaluate the throughput of techniques employing caching or the power consumption of various devices under load, we must exercise the algorithm or device using a sequence of synthetic packet headers. The *Trace Generator* produces a list of synthetic packet headers that probe filters in a given filter set. Note that we do not want to generate random packet headers. Rather, we want to ensure that a packet header is covered by at least one filter in the *FilterSet* in order to exercise the packet classifier and avoid default filter matches. We experimented with a number of techniques to generate synthetic headers. One possibility is to compute all the d-dimensional polyhedra defined by the intersections of the filters in the filter set, then choose a point in the d-dimensional space covered by the polyhedra. The point defines a packet header. The best-matching filter for the packet header is simply the highest priority filter associated with the polyhedra. If we generate at least one header corresponding to each polyhedra, we fully exercise the filter set. The number of polyhedra defined by filter intersections grows exponentially, and thus fully exercising the filter set quickly becomes intractable. As a result, we chose a method that partially exercises the filter set and allows the user to vary the size and composition of the headers in the trace using high-level input parameters. These parameters control the scale of the header trace relative to the filter set, as well as the locality of reference in the sequence of headers. As we did with the *Filter Set Generator*, we discuss the *Trace Generator* using the pseudocode shown in Fig. 14.

We begin by reading the *FilterSet* (line 1) and getting the input parameters *scale*, *ParetoA*, and *ParetoB* (lines 2–4). The *scale* parameter is used to set a threshold for the size of the list of headers relative to the size of the *FilterSet* (line 5). In this context, *scale* specifies the ratio of the number of headers in the trace to the number of filters in the filter set. The next set of steps continue to generate synthetic headers as long as the size of `Headers` does not exceed the `Threshold` defined by the product of *scale* and the number filters in *FilterSet*.

Each iteration of the header generation loop begins by selecting a random filter in the *FilterSet* (line 8). Next, we must choose a packet header covered by the filter. In the interest of exercising priority resolution mechanisms and providing conservative performance estimates for algorithms relying on filter overlap properties, we would like to choose headers matching a large number of filters. In the course of our analyses, we found the number of overlapping filters is large for packet headers representing the "corners" of filters. Each field of a filter covers a range of values. Choosing a packet header corresponding to a "corner" translates to choosing a value for each header field from one of the extrema of the range specified by each filter field. The `RandomCorner` function chooses a random "corner" of the filter identified by `RandFilt` and stores the header in `NewHeader`.

The last steps in the header generation loop append a variable number of copies of `NewHeader` to the trace. The number of copies, `Copies`, is chosen by sampling from a Pareto distribution controlled by the input parameters, *ParetoA* and *ParetoB* (line 10). In doing so, we provide a simple control point for the locality of reference in the header trace. The Pareto distribution[7] is one of the heavy-tailed distributions commonly used to model the burst size of Internet traffic flows as well as the file size distribution for traffic using the TCP protocol [20]. For convenience, let $a =$*ParetoA* and $b =$*ParetoB*. The probability density function for the Pareto distribution may be expressed as

$$P(x) = \frac{ab^a}{x^{a+1}} \tag{5}$$

where the cumulative distribution is

$$D(x) = 1 - \left(\frac{b}{x}\right)^a. \tag{6}$$

The Pareto distribution has a mean of

$$\mu = \frac{ab}{a-1}. \tag{7}$$

Expressed in this way, a is typically called the shape parameter and b is typically called the scale parameter, as the distribution is defined on values in the interval (b, ∞). The following are some examples of how the Pareto parameters are used to control locality of reference.

- Low locality of reference, short tail: $(a = 10, b = 1)$ most headers will be inserted once.
- Low locality of reference, long tail: $(a = 1, b = 1)$ many headers will be inserted once, but some could be inserted over 20 times.
- High locality of reference, short tail: $(a = 10, b = 4)$ most headers will be inserted four times.

Once the size of the trace exceeds the threshold, the header generation loop terminates. Note that a large burst near the end of the process will cause the trace to be larger than `Threshold`. After generating the list of headers, we write the trace to an output file (line 13).

[7]The Pareto distribution, a power law distribution named after the Italian economist Vilfredo Pareto, is also known as the Bradford distribution.

VII. Discussion

We have already found *ClassBench* to be tremendously valuable in our own research [21]–[23]. The *ClassBench* tools have also been used in a graduate level computer architecture course. Student groups used the tools to evaluate the performance of various packet classification algorithms [7], [21], [25], [26] implemented on an Intel IXP network processor [24].[8]

Acknowledgment

The authors would like to thank E. Spitznagel for contributing his insight to countless discussions on packet classification and assisting in the debugging of the *ClassBench* tools. They also would like to thank V. Srinivasan and W. Eatherton for making real filter sets available for study.

References

[1] D. E. Taylor, "Survey and taxonomy of packet classification techniques," *ACM Comput. Surv.*, vol. 37, no. 5, pp. 238–275, Sep. 2005.
[2] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis, "Framework for IP performance metrics," RFC 2330, May 1998.
[3] S. Bradner and J. McQuaid, "Benchmarking methodology for network interconnect devices," RFC 2544, Mar. 1999.
[4] G. Trotter, "Methodology for forwarding information base (FIB) based router performance," Internet Draft, Jan. 2002.
[5] B. Hickman, D. Newman, S. Tadjudin, and T. Martin, "Benchmarking methodology for firewall performance," RFC 3511, Apr. 2003.
[6] P. Chandra, F. Hady, and S. Y. Lim, "Framework for benchmarking network processors," in *Proc. Netw. Process. Forum*, 2002.
[7] P. Gupta and N. McKeown, "Packet classification on multiple fields," in *Proc. ACM SIGCOMM*, Aug. 1999, pp. 147–160.
[8] A. Feldmann and S. Muthukrishnan, "Tradeoffs for packet classification," in *Proc. IEEE INFOCOM*, Mar. 2000, pp. 1193–1202.
[9] P. Gupta and N. McKeown, "Packet classification using hierarchical intelligent cuttings," in *Proc. Hot Interconnects VII*, Aug. 1999, pp. 27–31.
[10] P. Warkhede, S. Suri, and G. Varghese, "Fast packet classification for 2-D conflict-free filters," in *Proc. IEEE INFOCOM*, 2001, pp. 1434–1443.
[11] F. Baboescu and G. Varghese, "Scalable packet classification," in *Proc. ACM SIGCOMM*, Aug. 2001, pp. 199–210.
[12] F. Baboescu and G. Varghese, "Fast and scalable conflict detection for packet classifiers," in *Proc. IEEE ICNP*, 2002, pp. 270–279.
[13] F. Baboescu, S. Singh, and G. Varghese, "Packet classification for core routers: Is there an alternative to CAMs?," in *Proc. IEEE INFOCOM*, 2003, pp. 53–63.
[14] T. Y. C. Woo, "A modular approach to packet classification: Algorithms and results," in *Proc. IEEE INFOCOM*, Mar. 2000, pp. 1213–1222.
[15] V. Sahasranaman and M. Buddhikot, "Comparative evaluation of software implementations of layer 4 packet classification schemes," in *Proc. IEEE Int. Conf. Netw. Protocols*, 2001, pp. 220–228.
[16] J. Balkman and P. Gupta, *PALAC: Packet Lookup and Classification Simulator*, User's Manual, ver. 4, Rev. 1 Oct. 2000.
[17] D. E. Taylor and J. S. Turner, "ClassBench: A packet classification benchmark," Dept. Comp. Sci. Eng., Washington Univ., St. Louis, Tech. Rep. WUCSE-2004-28, May 2004.
[18] Cisco, CiscoWorks VPN/Security Management Solution, Cisco Systems, Inc., 2004, Tech. Rep..

[8] The *ClassBench* tools and 12 *parameter files* are publicly available at the following site: http://www.arl.wustl.edu/~det3/ClassBench/.

[19] Lucent, Lucent Security Management Server: Security, VPN, and QoS Management Solution, Lucent Technologies Inc., 2004, Tech. Rep..
[20] K. Park, G. Kim, and M. Crovella, "On the effect of traffic self-similarity on network performance," in *Proc. 1997 SPIE Int. Conf. Perform. Contr. Netw. Syst.*, pp. 989–996.
[21] D. E. Taylor and J. S. Turner, "Scalable packet classification using distributed crossproducting of field labels," in *Proc. IEEE INFOCOM*, Mar. 2005, pp. 269–280.
[22] E. Spitznagel, D. Taylor, and J. Turner, "Packet classification using extended TCAMs," in *Proc. IEEE ICNP*, 2003, pp. 120–131.
[23] D. E. Taylor and E. W. Spitznagel, "On using content addressable memory for packet classification," Dept. Comp. Sci. Eng., Washington Univ., Saint Louis, Tech. Rep. WUCSE-2005-9, Mar. 2005.
[24] Intel Corp., IXP2400 Network Processor, Product Brief, Tech. Rep., 2002.
[25] S. Singh, F. Baboescu, G. Varghese, and J. Wang, "Packet classification using multidimensional cutting," in *Proc. ACM SIGCOMM*, 2003, pp. 213–224.
[26] V. Srinivasan, S. Suri, and G. Varghese, "Packet classification using tuple space search," in *Proc. ACM SIGCOMM*, 1999, pp. 135–146.

David E. Taylor received the B.S. and M.S. degrees in electrical and computer engineering in 1998 and 2002, respectively, and the D.Sc. degree in computer engineering in 2004, all from Washington University, St. Louis, MO.

He is a System Architect and the Director of Hardware Engineering at Exegy Inc., St. Louis, MO, where his primary focus is the development of high-performance hybrid computing systems for government intelligence and financial services applications. Prior to joining Exegy, he was a Visiting Assistant Professor in the Department of Computer Science and Engineering, Washington University, where he was also actively involved in computer communications research at the Applied Research Laboratory. His research interests include the design and analysis of scalable searching algorithms and architectures, IP lookup and packet classification algorithms, high-performance reconfigurable hardware systems, programmable routers, and network processors.

Jonathan S. Turner (M'77–SM'88–F'90) received the M.S. and Ph.D. degrees in computer science from Northwestern University, Evanston, IL, in 1979 and 1981, respectively.

He holds the Barbara and Jerome Cox Chair of Computer Science at Washington University, St. Louis, MO, and is Director of the Applied Research Laboratory. The Applied Research Laboratory creates experimental networking technology to validate and demonstrate new research innovations. The Laboratory's current projects center on extensible networking technology with a particular focus on high performance diversified routers. He served as Chief Scientist for Growth Networks, a startup company that developed scalable switching components for Internet routers and ATM switches, before being acquired by by Cisco Systems in early 2000. His primary research interests revolve around the design and analysis of switching systems, with special interest in systems supporting multicast communication. His research interests also include the study of algorithms and computational complexity, with particular interest in the probable performance of heuristic algorithms for NP-complete problems. He has been awarded more than 25 patents for his work on switching systems and has many widely cited publications.

Dr. Turner received the Koji Kobayashi Computers and Communications Award from the IEEE in 1994 and the IEEE Millenium Medal in 2000. He is a Fellow of ACM.