

Compressing IP Forwarding Tables: Towards Entropy Bounds and Beyond

Gábor Rétvári, *Member, IEEE*, János Tapolcai, *Member, IEEE*, Attila Kőrösi, András Majdán, and Zalán Heszberger

Abstract—Lately, there has been an upsurge of interest in compressed data structures, aiming to pack ever larger quantities of information into constrained memory without sacrificing the efficiency of standard operations, like random access, search, or update. The main goal of this paper is to demonstrate how data compression can benefit the networking community by showing how to squeeze the IP Forwarding Information Base (FIB), the giant table consulted by IP routers to make forwarding decisions, into information-theoretical entropy bounds, with essentially zero cost on longest prefix match and FIB update. First, we adopt the state of the art in compressed data structures, yielding a static entropy-compressed FIB representation with asymptotically optimal lookup. Then, we redesign the venerable prefix tree, used commonly for IP lookup for at least 20 years in IP routers, to also admit entropy bounds and support lookup in optimal time and update in nearly optimal time. Evaluations on a Linux kernel prototype indicate that our compressors encode an FIB comprising more than 440 K prefixes to just about 100–400 kB of memory, with a threefold increase in lookup throughput and no penalty on FIB updates.

Index Terms—Data compression, IP forwarding table lookup, prefix tree.

I. INTRODUCTION

DATA compression is widely used in processing large volumes of information. Not just that convenient compression tools are available to curtail the memory footprint of basically any type of data, but these tools also come with theoretical guarantee that the compressed size is indeed minimal, in terms of some suitable notion of *entropy* [2]. Correspondingly, data compression has found its use in basically all aspects of

Manuscript received February 05, 2014; revised September 02, 2014; accepted September 08, 2014; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor A. Bremner-Barr. The work of G. Rétvári was supported by the OTKA/PD-104939 grant. This work is an extended and corrected version of the manuscript “Compressing IP Forwarding Tables: Towards Entropy Bounds and Beyond,” which originally appeared at the ACM SIGCOMM 2013 [1]. The present paper contains a number of crucial revisions to the original text, in particular, a correct statement of the information-theoretical bounds (Proposition 1 and 2), an updated form of the *XBW-b* transform, improved space bounds, and complete proofs.

G. Rétvári, J. Tapolcai, and A. Majdán are with the MTA-BME Future Internet Research Group, High-Speed Networks Laboratory (HSNLab), Department of Telecommunications and Media Informatics, Budapest University of Technology and Economics, 1117 Budapest, Hungary (e-mail: retvari@tmit.bme.hu; tapolcai@tmit.bme.hu; majdan@tmit.bme.hu).

A. Kőrösi and Z. Heszberger are with the MTA-BME Information Systems Research Group, Department of Telecommunications and Media Informatics, Budapest University of Technology and Economics, 1117 Budapest, Hungary (e-mail: korosi@tmit.bme.hu; heszi@tmit.bme.hu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2014.2357051

computation and networking practice, ranging from text or multimedia compression [3] to the very heart of communications protocols [4] and operating systems [5].

Traditional compression algorithms do not admit standard queries, like pattern matching or random access, right on the compressed form, which severely hinders their applicability. An evident workaround is to decompress the data prior to accessing it, but this pretty much defeats the whole purpose. The alternative is to maintain a separate index dedicated solely to navigate the content, but the sheer size of the index can become prohibitive in many cases [6], [7].

It is no surprise, therefore, that the discovery of compressed self-indexes (or, in the context of this paper, *compressed data structures*) came as a real breakthrough [8]. A compressed data structure is, loosely speaking, an entropy-sized index on some data that allows complete recovery of the original content as well as fast queries on it [3], [6]–[13]. What is more, as the compressed form occupies much smaller space than the original representation, and hence is more friendly to CPU cache, the time required to answer a query is often far less than if the data had not been compressed [3], [11]. *Compressed data structures, therefore, turn out one of the rare cases in computer science where there is no space–time tradeoff.*

Researchers and practitioners working with big data were quick to recognize this win–win situation and came up with compressed self-indexes, and accompanying software tools, for a broad set of applications; from compressors for sequential data like bitmaps (RRR), [12] and text documents (CGlimpse [8], wavelet trees [13]); compression front ends to information retrieval systems (MG4J [14], LuceneTransform [15]) and dictionaries (MG [16]); to specialized tools for structured data, like XML/HTML/DOM (XGRIND [17], XBZIPINDEX [18]), graphs (WebGraph [19]), 3-D models (Edgebreaker [20]), genomes and protein sequences (COMRi [21]), multimedia, source and binary program code, formal grammars, etc. [16]. With the advent of replacements for the standard file compression tools (Lzgrep [22]) and generic libraries (libcds [9]), we might be right at the verge of seeing compressed data structures go mainstream.

Curiously, this revolutionary change has gone mostly unnoticed in the networking community, even though this field is just one of those affected critically by skyrocketing volumes of data. A salient example of this trend is the case of the IP Forwarding Information Base (FIB), used by Internet routers to make forwarding decisions, which has been literally deluged by the rapid growth of the routed IP address space lately. Consequently, there has been a flurry of activity to find space-efficient FIB representations [23]–[42], yet very few of these go beyond

ad hoc schemes and compress to information-theoretic limits, let alone come with a convenient notion of FIB entropy. Taking the example of compressing IP FIBs as a use case, our aim in this paper is to popularize compressed data structures to the networking community.

A. FIB Compression

There are hardly any data structures in networking affected as compellingly by the growth of the Internet as the IP FIB. Stored in the line card memory of routers, the FIB maintains an association from every routed IP prefix to the corresponding next-hop, and it is queried on a packet-by-packet basis at line speed (in fact, it is queried *twice* per packet, considering reverse path forwarding check). Lookup in FIBs is not trivial either, as IP's longest prefix match rule requires the most specific entry to be found for each destination address. Moreover, as Internet routers operate in an increasingly dynamic environment [43], the FIB needs to support hundreds of updates to its content each second.

As of 2013, the number of active IPv4 prefixes in the Default Free Zone is more than 440 000 and counting, and IPv6 quickly follows suit [44]. Correspondingly, FIBs continue to expand both in size and management burden. As a quick reality check, the Linux kernel's `fib_trie` data structure [25], when filled with this many prefixes, occupies tens of megabytes of memory, takes several minutes to download to the forwarding plane, and is still heavily debated to scale to multigigabit speeds [45]. Commercial routers suffer similar troubles, aggravated by the fact that line card hardware is more difficult to upgrade than software routers.

Many have argued that mounting FIB memory tax will, sooner or later, present a crucial data-plane performance bottleneck for IP routers [46]. Yet even if the scalability barrier will not prove impenetrable [47], the growth of the IP forwarding table still poses compelling difficulties. Adding further fast memory to line cards boosts silicon footprint, heat production, and power budget, all in all, the CAPEX/OPEX associated with IP network gear, and forces operators into rapid upgrade cycles [48], [49]. Large FIBs also complicate maintaining multiple virtual router instances, each with its own FIB, on the same physical hardware [50] and build up huge control-plane-to-data-plane delay for FIB resets [51].

Several recent studies have identified *FIB aggregation* as an effective way to reduce FIB size, extending the lifetime of legacy network devices and mitigating the Internet routing scalability problem, at least temporarily [48], [49]. FIB aggregation is a technique to transform some initial FIB representation into an alternative form that, supposedly, occupies smaller space but still provides fast lookup. Recent years have seen an impressive reduction in FIB size: From the initial 24 B/prefix (prefix trees [37]), use of hash-based schemes [27], [42], path- and level-compressed multibit tries [25], [26], [39], tree-bitmaps [32], etc., have reduced FIB memory tax to just about 2–4.5 B/prefix [29], [40], [41]. Meanwhile, lookup performance has also improved [25].

The evident questions “Is there an ultimate limit in FIB aggregation?” and “Can FIBs be reduced to fit in fast ASIC SRAM/CPU cache entirely?” have been asked several times before [26], [29], [31], [39]. In order to answer these questions, we need to go beyond conventional FIB aggregation

to find *new compressed FIB data structures that encode to entropy-bounded space and support lookup and update in optimal time*. We coined the term *FIB compression* to mark this ambitious undertaking [52]. Accordingly, this paper is dedicated to the theory and practice of FIB compression.

B. Our Contributions

Our contributions on FIB compression are twofold: Based on the labeled tree entropy measure of Ferragina *et al.* [18], we specify a compressibility metric called FIB entropy, then we propose two entropy-compressed FIB data structures.

Our first FIB encoder, *XBW-b*, is a direct application of the state of the art in compressed data structures to the case of IP FIBs. *XBW-b* compresses a contemporary FIB to the entropy limit of just 100–300 kB and, at the same time, provides longest prefix match in asymptotically optimal time. Unfortunately, it turns out that the relatively immature hardware and software background for compressed string indexes greatly constrain the lookup and update performance of *XBW-b*. Therefore, we also present a practical FIB compression scheme, called the trie-folding algorithm.

Trie-folding is in essence a “compressed” reinvention of prefix trees, a commonly used FIB implementation in IP routers, and therefore readily deployable with minimal or no modification to router ASICs [53]. We show that trie-folding compresses to within a small constant factor of FIB entropy, supports lookup in strictly optimal time, and admits updates in nearly optimal time for FIBs of reasonable entropy (see later for precise definitions). The viability of trie-folding will be demonstrated on a Linux prototype and an FPGA implementation. By extensive tests on real and synthetic IP FIBs, we show that trie-folding supports tens of millions of IP lookups and hundreds of thousands updates per second, in less than 150–500 kB of memory.

C. Structure of the Paper

This paper is organized as follows. In Section II, we survey standard FIB representation schemes and cast compressibility metrics. In Section III, we describe *XBW-b*, while in Section IV we introduce trie-folding and establish storage size bounds. Section V is devoted to numerical evaluations and measurement results, Section VI surveys related literature, and finally Section VII concludes the paper.

II. PREFIX TREES AND SPACE BOUNDS

Consider the sample IP routing table in Fig. 1(a), storing address-prefix-to-next-hop associations in the form of an index into a neighbor table, which maintains neighbor specific information, like next-hop IP address, aliases, ARP info, etc. Associate a unique *label* from the alphabet Σ with each next-hop in the neighbor table. We shall usually treat labels as positive integers, complemented with a special *invalid label* $\perp \in \Sigma$ to mark blackhole routes. Let N denote the number of entries in the FIB, and let $\delta = |\Sigma|$ be the number of next-hops. An IP router does not keep an adjacency with every other router in the Internet, thus $\delta \ll N$. Specifically, we assume that δ is $O(\text{polylog } N)$ or $O(1)$, which is in line with reality [54], [55]. Finally, let W denote the width of the address space in bits (e.g., $W = 32$ for IPv4).

To actually forward a packet, we need to find the entry that matches the destination address in the packet on the greatest




Fig. 1. Representations of an IP forwarding table: (a) tabular form with address in binary format, prefix length, and next-hop address label; (b) prefix tree with state transitions marked; (c) ORTC-compressed prefix tree; (d) level-compressed multibit trie; and (e) leaf-pushed trie.

number of bits, starting from the MSB. For the address 0111, each of the entries $-/0$ (the default route), 0/1, 01/2, and 011/3 match. As the most specific entry is the last one, the lookup operation yields the next-hop label 1. This is then used as an index into the neighbor table, and the packet is forwarded on the interface facing that neighbor. This tabular representation is not really efficient, as a lookup or update operation requires looping through each entry, taking $O(N)$ time. The storage size is $(W + \lg \delta)N$ bits.¹

Binary prefix trees, or *tries* [37], support lookup and update much more efficiently; see Fig. 1(b). A trie is a labeled ordinal tree, in which every path from the root node to a leaf corresponds to an IP prefix and lookup is based on successive bits of the destination address: If the next bit is 0, proceed to the left subtree, otherwise proceed to the right, and if the corresponding child is missing, return the last label encountered along the way. Prefix trees generally improve the time to perform a lookup or update from linear to $O(W)$ steps.

A prefix tree representation is usually not unique, which opens the door to a wide range of optimization strategies to find more space-efficient forms. For instance, the prefix tree in Fig. 1(c) is *forwarding equivalent* with the one in Fig. 1(b) in that it orders the same label to every complete W -bit-long key, yet contains only three labeled nodes instead of seven (see the ORTC algorithm in [31] and [40]). Alternatively, *level-compression* [25], [26], [39] is a technique to remove excess levels from a binary trie to obtain a forwarding equivalent multibit trie that is substantially smaller; see Fig. 1(d).

A standard technique to obtain a *unique, normalized form* of a prefix tree is *leaf-pushing* [31], [39], [50]: In a first pre-order traversal, labels are pushed from the parents toward the children, and then in a second post-order traversal, each parent with identically labeled leaves is substituted with a leaf marked with the children’s label; see Fig. 1(e). The resultant trie is called a *leaf-labeled* trie since interior nodes no longer maintain labels, and it is also a *proper* binary trie with nice structure: Any node is either a leaf node, or it is an interior node with exactly two children. Updates to a leaf-pushed trie, however, may be expensive; modifying the default route, for instance, can result in practically all leaves being relabeled, taking $O(N)$ steps in the worst case.

A. Information-Theoretic Limit

How can we know for sure that a particular prefix tree representation, from the many, is indeed space-efficient? The first verifiable cornerstone of a space-efficient data structure is whether its size meets the *information-theoretic lower bound*,

corresponding to the minimum number of bits needed to uniquely identify any instance of the data. For example, there are exactly δ^n strings of length n on an alphabet of size δ , and to be able to distinguish between any two, we need at least $\lg(\delta^n) \approx n \lg \delta$ bits. In this example, even a naive string representation meets the bound, but in more complex cases, attaining it is much more difficult.

This argumentation generalizes from strings to leaf-labeled tries as follows (see also Ferragina *et al.* [18]).

Proposition 1: Let T be a proper, binary, leaf-labeled trie with n leaves on an alphabet of size δ . The information-theoretic lower bound to encode T is $2n + n \lg \delta$ bits.²

The bound is easily justified with a simple counting argument. The number of proper binary trees on n leaves is the $(n - 1)$ th Catalan number $C_{n-1} = (1/n) \binom{2n-2}{n-1}$, therefore we need at least $\lg C_{n-1} = 2n - \Theta(\log n)$ bits to encode the tree itself [56]; storing the label map defined on the n leaves of T requires an additional $n \lg \delta$ bits; and assuming that the two are independent, we need $2n + n \lg \delta$ bits overall.

A representation that encodes to within the constant factor of the information-theoretic lower bound (up to lower-order terms) and simultaneously supports queries in optimal time is called a *compact data structure*, while if the constant factor is 1, then it is also a *succinct data structure* [56].

B. Entropy Bounds

A succinct representation very often contains further redundancy in the form of regularity in the label mapping. For instance, in the sample trie of Fig. 1(e), there are three leaves with label 2, but only one with label 1 or 3. Thus, we could save space by representing label 2 on fewer bits, similarly to how Huffman coding does for strings. This correspondence leads to the following notion of *entropy* for leaf-labeled tries (on the traces of Ferragina *et al.* [18]).

Proposition 2: Let T be a proper, binary, leaf-labeled trie with n leaves on an alphabet Σ , let p_s denote the probability that some symbol $s \in \Sigma$ appears as a leaf label, and let H_0 denote the Shannon entropy of the probability distribution p_s , $s \in \Sigma$

$$H_0 = \sum_{s \in \Sigma} p_s \log_2 1/p_s. \quad (1)$$

Then, the zero-order entropy of T is $2n + nH_0$ bits.³


Intuitively speaking, the entropy of the tree structure corresponds to the information-theoretic limit of $2n$ bits as we do not assume any regularity in this regard. To this, the leaf labels add an extra nH_0 bits of entropy.

The entropy of a trie depends mainly on the size of the underlying tree and the distribution of labels on it. This transforms to FIBs quite naturally: The more prefixes go to the same next-hop and the more the FIB resembles “a default route with few exceptions,” the smaller the Shannon entropy of the next-hop distribution and the tighter the space limit. Accordingly, we shall define the notions *FIB information-theoretic lower bound* and *FIB entropy* as those of the underlying leaf-pushed prefix tree. Both space bounds are well defined as the normalized form is unique. Note, however, that in contrast to general trees, IP FIBs

²In the original manuscript [1], the information-theoretic lower bound is wrongly set to $n(4 + \lg \delta)$ bits instead of $n(2 + \lg \delta)$. The difference is due to erroneously taking the number of trees on n nodes, instead of n leaves.

³The claim is revised from [1], where the entropy wrongly appears as $4n + nH_0$. See previous note for an explanation.

¹The notation $\lg x$ is shorthand for $\lceil \log_2(x) \rceil$.

Fig. 2. Leaf-pushed trie and its XBW - b transform.

are of bounded height, so the real bounds should be somewhat smaller. Additionally, for the purposes of this paper, our space bounds involve binary leaf-labeled tries only. We relax this restriction in [52] using the generic trie entropy measure of [18].

III. ATTAINING ENTROPY BOUNDS

We present our first compressed FIB data structure, the *Burrows-Wheeler transform for binary leaf-labeled tries* (XBW - b). This data structure is a stripped down version of MBW , the *Multibit Burrows-Wheeler transform* from [52], and the XBW - l transform from [1], which in turn build on the succinct level-indexed binary trees of Jacobson [56] and the XBW transform due to Ferragina *et al.* [18]. In contrast to XBW - b that is binary only, the original MBW and XBW - l transforms support level-compressed tries as well, at the price of encoding to a slightly larger representation and missing the information-theoretical storage size bounds defined above.

The basis for the XBW - b transform is a normalized, binary, proper leaf-labeled trie. Let T be a binary tree on t nodes, let L be the set of leaves with $n = |L|$, and let l be a mapping $V \mapsto \Sigma$ specifying for a node v either that v does not have a label associated with it (i.e., $l(v) = \emptyset$) or the corresponding label $l(v) \in \Sigma$. If T is proper, binary, and leaf-labeled, then the following invariants hold.

- P1) Either $v \in L$, or v has 2 children.
- P2) $l(v) \neq \emptyset \Leftrightarrow v \in L$.
- P3) $t < 2n$ and so $t = O(n)$.

The main idea in XBW - b is serializing T into a bitstring S_I that encodes the tree structure and a string S_α on the alphabet Σ encoding the labels, and then using a sophisticated lossless string compressor to obtain the storage size bounds. Correspondingly, the XBW - b transform is defined as the tuple $\text{xbwb}(T) = (S_I, S_\alpha)$, where:

- S_I is a bitstring of size t with zero in position i if the i th node of T in level-order is an interior node, and 1 otherwise.
- S_α is a string of size n on the alphabet Σ encoding the leaf labels.

For our sample FIB, the leaf-pushed trie and the corresponding XBW - b transform are given in Fig. 2.

A. Construction and IP Lookup

In order to generate the XBW - b transform, one needs to fill up the strings S_I and S_α , starting from the root and traversing T in a breadth-first-search order. Meanwhile, two counters are maintained; i is used to enumerate the nodes and index into S_I , while j counts the leaves and indexes S_α . For every node, we decide whether it is interior, in which case the corresponding entry of S_I is set to zero, or a leaf. In the latter case, S_I is set to 1, and the leaf label is appended to S_α .

```

1:  $i \leftarrow 1; j \leftarrow 1$ 
2: BFS-TRAVERSE (node  $v$ , integer  $i$ , integer  $j$ )
3:   if  $v \notin L$  then  $S_I[i] \leftarrow 0$ 
4:     else  $S_I[i] \leftarrow 1; S_\alpha[j] \leftarrow l(v); j \leftarrow j + 1$ 
5:    $i \leftarrow i + 1$ 

```

The following statement is now obvious.

Lemma 1: Given a proper binary, leaf-labeled trie T on t nodes, $\text{xbwb}(T)$ can be built in $O(t)$ time.

The transform $\text{xbwb}(T)$ has some appealing properties. For instance, the children of some node, if the exist, are stored on consecutive indices in S_I and S_α . In fact, *all* nodes at the same level of T are mapped to consecutive indices.

The next step is to actually compress the strings. This promises easier than compressing T directly as $\text{xbwb}(T)$, being a sequential string representation, lacks the intricate structure of tries. An obvious choice would be to apply some standard string compressor (like the venerable `gzip(1)` tool), but this would not admit queries like “get all children of a node” without first decompressing the transform. Thus, we rather use a compressed string self-index [12], [13], [18], [56] to store $\text{xbwb}(T)$, which facilitates efficient navigation immediately on the compressed form.

The way string indexers usually realize navigability is to implement a certain set of simple primitives in constant $O(1)$ time in place. Given a string $S[1, t]$ on alphabet Σ , a symbol $s \in \Sigma$, and integer $q \in [1, t]$, these primitives are as follows:

- $\text{access}(S, q)$: return the symbol at position q in S ;
- $\text{rank}_s(S, q)$: return the number of times symbol s occurs in the prefix $S[1, q]$;
- $\text{select}_s(S, q)$: return the position of the q th occurrence of symbol s in S .

Curiously, these simple primitives admit strikingly complex queries to be implemented and supported in optimal time. In particular, the IP lookup routine on $\text{xbwb}(T)$ takes the following form.

```

1: lookup (address  $a$ )
2:    $q \leftarrow 0, i \leftarrow 1$ 
3:   while  $q < W$ 
4:     if  $\text{access}(S_I, i) = 1$  then
5:       return  $\text{access}(S_\alpha, \text{rank}_1(S_I, i))$ 
6:      $r \leftarrow \text{rank}_0(S_I, i)$ 
7:      $j \leftarrow \text{bits}(a, q, 1)$ 
8:      $i \leftarrow 2r + j; q \leftarrow q + 1$ 

```

The code first checks if the actual node, encoded at index i in $\text{xbwb}(T)$, is a leaf node. If it is, then $\text{rank}_1(S_I, i)$ tells how many leaves were found in the course of the BFS traversal *before* this one, and then the corresponding label is returned from S_α . If, on the other hand, the actual node is an interior node, then r tells how many interior nodes precede this one. Note that, as one easily checks, in a level-ordered tree traversal the children of the r th interior node are encoded from position $2r$ [56]. Next, we obtain the index j of the child to be visited next from the address to be looked up, we set the current index to $2r + j$, and then we carry on with the recursion.

B. Memory Size Bounds

First, we show that $XBW-b$ is a succinct FIB representation, in that it supports lookup in optimal $O(W)$ time and encodes to information-theoretic lower bound.

Lemma 2: Given a proper, binary, leaf-labeled trie T with n leaves on an alphabet of size δ , $\text{xbwb}(T)$ can be stored on $2n + n \lg \delta + o(n)$ bits so that `lookup` on $\text{xbwb}(T)$ terminates in $O(W)$ time.

Proof: Encode S_I on $2n + o(n)$ bits using the RRR succinct bitstring index [12], which supports select and rank in $O(1)$. In addition, even the trivial encoding of S_α needs only another $n \lg \delta$ bits and provides access in $O(1)$. Thus, every iteration of `lookup` takes constant time, hence the result. ■

Next, we show that $XBW-b$ can take advantage of regularity in leaf labels (if any) and encode *below* the information-theoretic bound to zero-order entropy.

Lemma 3: Let T be a proper, binary, leaf-labeled trie with n leaves on an alphabet of size $O(\text{polylog } n)$, and let H_0 denote the Shannon entropy of the leaf-label distribution. Then, $\text{xbwb}(T)$ can be stored on $2n + nH_0 + o(n)$ bits so that `lookup` on $\text{xbwb}(T)$ terminates in $O(W)$ time.

Proof: S_I can be encoded as above, and S_α can be stored on $nH_0 + o(n)$ bits using generalized wavelet trees so that access is $O(1)$, under the assumption that the alphabet size is $O(\text{polylog } n)$ [13]. ■

Interestingly, the above *zero-order* entropy bounds can be easily upgraded to *higher-order* entropy. A fundamental premise in data compression is that elements in a data set often depend on their neighbors, and the larger the context, the better the prediction of the element from its context and the more efficient the compressor. Higher-order string compressors can use the Burrows–Wheeler transform to exploit this contextual dependency, a reversible permutation that places symbols with similar context close to each other. This argumentation readily generalizes to leaf-labeled tries; simply, the *context of a node corresponds to its level in the tree* and because $XBW-b$ organizes nodes at the same level (i.e., of similar context) next to each other, it realizes the same effect for tries as the Burrows–Wheeler transform for strings (hence the name). Deciding whether or not such contextual dependency is present in real IP FIBs is well beyond the scope of this paper. Here, we only note that if it is, then $XBW-b$ can take advantage of this and compress an IP FIB to higher-order entropy using the techniques in [18] and [52].

In summary, the $XBW-b$ transform can be built quickly, supports lookup in asymptotically optimal time, and compresses to within entropy bounds. Updates, however, may be expensive. Even the underlying leaf-pushed trie takes $O(n)$ steps in the worst case to update, after which we could either rebuild the string indexes from scratch (again in $O(n)$) or use a dynamic compressed index that supports updates to the compressed form efficiently. For instance, [10] implements insertion and deletion in roughly $O(\log n)$ time, at the cost of slower rank and select. The other shortcoming of $XBW-b$ is that, even if it supports lookups in *theoretically* optimal time, it is just too slow for line speed IP lookup (see Section V-C). In Section IV, therefore, we discuss a somewhat more practical FIB compression scheme.

IV. PRACTICAL FIB COMPRESSION

The string indexes that underly $XBW-b$ are *pointerless*, encoding all information in compact bitmaps. This helps squeezing $XBW-b$ into higher-order entropy bounds, but also causes that we need to perform multiple rank and select operations just to, say, enter a child of a node. Even though these primitives run in $O(1)$, the constants still add up, building up delays too huge for line speed IP lookup. In contrast, a traditional *pointer machine*, like a prefix tree, can follow a child pointer in just a single indirection with only one random memory access overhead. The next compressed FIB data structure we introduce is, therefore, pointer-based.

The idea is to essentially reinvent the classic prefix tree, borrowing the basic mechanisms from the Lempel–Ziv (LZ78) string compression scheme [2]. LZ78 attains entropy by parsing the text into unique substrings, yielding a form that contains no repetitions. Tree threading is a generalization of this technique to *unlabeled trees*, merging isomorphic subtrees into a directed acyclic graph (DAG) [35], [50], [57], [58]. In this paper, we apply this idea to *labeled trees* by taking into account both the underlying tree structure *and* the labels when merging subtrees [59], [60]. If the trie is regular, then this will eliminate all recurrent substructures, producing a representation that contains no repetitions and hence, so the argument goes, admits entropy bounds like LZ78.

The below equivalence relation serves as the basis of our trie-merging technique.

Definition 1: Two leaf-labeled tries are identical if each of their subtrees are pairwise identical, and two leaves are identical if they hold the same label.

We call the algorithmic manifestation of this recursion the *trie-folding algorithm* and the resultant representation a *prefix DAG*. Fig. 3(a) depicts a sample prefix tree, Fig. 3(b) shows the corresponding leaf-pushed trie, and Fig. 3(c) gives the prefix DAG. For instance, the subtrees that belong to the prefix 0/1 and 11/2 are equivalent in the leaf-pushed trie and are thus merged into a single subtree that is now available in the prefix DAG along two paths from the root. This way, the prefix DAG is significantly smaller than the original prefix tree as it contains only half the nodes.

For the trie-folding algorithm, it is essential that the underlying trie be normalized; for instance, in our example it is easy to determine from the leaf-pushed trie that the two subtrees under the prefixes 00/2 and 10/2 are identical, but this is impossible to know from the original prefix tree. Thus, leaf-pushing is essential to realize good compression but, at the same time, makes updates prohibitive [50].

To avoid this problem, we apply a simple optimization. We separate the trie into two parts; “above” a certain level λ , called the *leaf-push barrier*, where subtrees are huge and so common subtrees are rare, we store the FIB as a standard binary prefix tree in which update is fast; and “below” λ , where common subtrees most probably turn up, we apply leaf-pushing to obtain good compression. Then, by a cautious setting of the leaf-push barrier, we simultaneously realize fast updates and entropy-bounded storage size.

The prefix DAGs for $\lambda = 1$ and $\lambda = 2$ are depicted in Fig. 3(e) and (f). The size is somewhat larger, but updating, say, the default route now only needs setting the root label without having to cycle through each leaf.

A. Construction and IP Lookup

We are given a binary trie T (not necessarily proper and leaf-pushed) of depth W , labeled from an alphabet Σ of size δ . Let V_T ($|V_T| = t$) be the set of nodes and L_T ($|L_T| = n$) be the set of leaves. Without loss of generality, we assume that T does not contain explicit blackhole routes. Then, the trie-folding algorithm transforms T into a prefix DAG $D(T)$, on nodes V_D and leaves L_D , with respect to the leaf-push barrier $\lambda \in [0, W]$. The algorithm itself is a simple variant of trie threading [57]: Assign a unique ID to each subtree that occurs below level λ and merge two tries if their IDs are equal (as of Definition 1). The algorithm actually works on a copy of T and always keeps an intact instance of T available. This instance, called the control FIB, can exist in the DRAM of the line card's control CPU as it is only consulted to manage the FIB. The prefix DAG itself is constructed in fast memory. We also need two ancillary data structures, the leaf table and the subtree index, which can also live in DRAM.

The *leaf table* will be used to coalesce leaves with identical labels into a common leaf node. Accordingly, for each $s \in \Sigma$, the leaf table $\text{lp}(s)$ stores a leaf node (no matter which one) with that label. Furthermore, the *subtrie index* \mathcal{S} will be used to identify and share common subtrees. \mathcal{S} is in fact a reference counted associative array, addressed with pairs of IDs $(i, j) \in \mathbb{N} \times \mathbb{N}$ as keys and storing for each key a node whose children are exactly the subtrees identified by i and j . \mathcal{S} supports the following primitives.

- $\text{put}(i, j, v)$: If a node with key (i, j) exists in \mathcal{S} , then increase its reference count and return it, otherwise generate a new ID in $v.\text{id}$, store v at key (i, j) with reference count 1, and return v .
- $\text{get}(i, j)$: Dereference the entry with key (i, j) and delete it if the reference count drops to zero.

In our code, we used a hash to implement \mathcal{S} , which supports the above primitives in amortized $O(1)$ time.

Now, suppose we are given a node v to be subjected to trie-folding and a leaf-push barrier λ . First, for each descendant u of v at depth λ , we normalize the subtree rooted at u using label $l(u)$ as a “default route,” and then we call the `compress` routine to actually merge identical leaves and subtrees below u , starting from the bottom and working upwards until we reach u . Consider the following pseudocode for the main `trie_fold` routine.

```

1: trie_fold(node  $v$ , integer  $\lambda$ )
2:   each  $\lambda$ -level child  $u$  of  $v$  do
3:     if  $l(u) = \emptyset$  then leaf_push( $u, \perp$ )
4:     else leaf_push( $u, l(u)$ )
5:     POSTORDER-TRAVERSE-AT- $u$  (compress)
6:      $l(\text{lp}(\perp)) \leftarrow \emptyset$ 
7:   compress (node  $w$ )
8:   if  $w \in L_D$  then  $w.\text{id} = l(w)$ ;  $u \leftarrow \text{lp}(w)$ 
9:   else  $u = \text{put}(w.\text{left}.\text{id}, w.\text{right}.\text{id}, w)$ 
10:  if  $u \neq w$  then repointer the parent of  $w$  to  $u$ ; delete( $w$ )

```

Here, $w.\text{left}$ is the left child and $w.\text{right}$ is the right child for w , and $w.\text{id}$ is the ID of w . As `trie_fold` visits each node at most twice and `compress` runs in $O(1)$ if `put` is $O(1)$, we arrive to the following conclusion.

Lemma 4: Given a binary trie T on t nodes, $D(T)$ can be constructed in $O(t)$ time.

Lookup on a prefix DAG goes exactly the same way as on a conventional prefix tree: Follow the path traced out by the successive bits of the lookup key and return the last label found. We need to take care of a subtlety in handling invalid labels, though. Namely, in our sample FIB of Fig. 3(a), the address 000 is associated with label 1 (the default route), which in the prefix DAG for $\lambda = 1$ [Fig. 3(e), derived from the trie on Fig. 3(d)] would become \perp if we were to let leaf nodes’ \perp labels override labels inherited from levels above λ . This problem is easy to overcome, though, by removing the label from the leaf $\text{lp}(\perp)$. By our assumption, the FIB contains no explicit blackhole routes, thus every traversal yielding the empty label on T terminates in $\text{lp}(\perp)$ on $D(T)$ and, by the above modification, also gives an empty label.

That being said, the last line of the `trie_fold` algorithm renders standard trie lookup correct on prefix DAGs. Since this is precisely the lookup algorithm implemented in many IP routers on the ASIC [53], we conclude that prefix DAGs can serve as compressed drop-in replacements for trie-based FIBs in many router products (similarly to, e.g., [40]).

Lemma 5: The `lookup` operation on $D(T)$ terminates in $O(W)$ time.

In this regard, trie-folding can be seen as a generalization of conventional FIB implementations: For $\lambda = 32$, we get good old prefix trees, and for smaller settings of λ , we obtain *increasingly smaller FIBs with exactly zero cost on lookup efficiency*. Correspondingly, there is no memory size versus lookup complexity “space–time” tradeoff in trie-folding.

B. Memory Size Bounds

The tries that underlie *XBW-b* are proper and leaf-labeled, and the nice structure makes it easy to reason about the size thereof. Unfortunately, the tries that prefix DAGs derive from are of arbitrary shape, and so it is difficult to infer space bounds in the same generic sense. We chose a different approach, therefore, in that we *view trie-folding as a generic string compression method, and we compare the size of the prefix DAG to that of an input string* given to the algorithm. The space bounds obtained this way transform to prefix trees naturally, as trie entropy itself is also defined in terms of string entropy (recall Proposition 2).

Instead of being given a trie, therefore, we are now given a string S of length n on an alphabet of size δ and zero-order entropy H_0 . Supposing that n equals some power of 2 (which we do for the moment), say, $n = 2^W$, we can think as if the symbols in S were written to the leaves of a complete binary tree of depth W as labels. Then, trie-folding will convert this tree into a prefix DAG $D(S)$, and we are curious as to how the size of $D(S)$ relates to the information-theoretic limit for storing S , that is, $n \lg \delta$ bits, and the zero order entropy nH_0 (see Fig. 4). Note that every FIB has such a “complete binary trie” representation, and vice versa.

The memory model for storing the prefix DAG is as follows. Above the leaf-push barrier λ , the children of a node are arranged on consecutive memory locations, and so each node holds a single node pointer of size to be determined later plus a label index of $\lg \delta$ bits. At and below level λ , nodes hold two pointers but no label, plus we also need an additional $\delta \lg \delta$ bits to store the coalesced leaves.

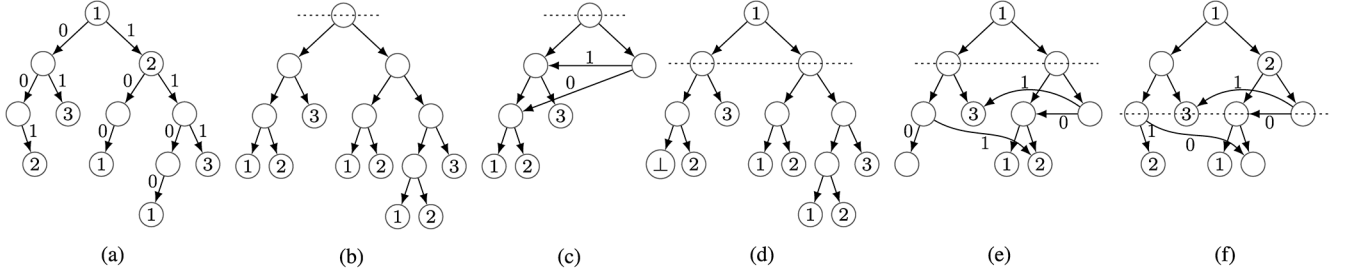


Fig. 3. (a) Binary trie for a sample FIB. (b) Same trie when leaf-pushing is applied from level $\lambda = 0$. (c) Prefix DAG for leaf-push barrier $\lambda = 0$. (d) Trie and (e) prefix DAG for $\lambda = 1$. (f) Prefix DAG for $\lambda = 2$. Dashed lines indicate the leaf-push barrier λ , and the invalid label \perp was removed from the corresponding leaf nodes in the prefix DAGs.




Fig. 4. Trie-folding as string compression: (a) a string, (b) the complete binary trie, and (c) the compressed DAG. The third character can be accessed by looking up the key $3 - 1 = 010_2$.

Now, we are in a position to state the first space bound. In particular, we show that $D(S)$ attains information-theoretic lower bound up to some small constant factor, and so it is a *compact data structure*. Our result improves the constant factor from 5 (as of [1]) to 4.

Theorem 1: Let S be a string of length $n = 2^W$ on an alphabet of size δ and set the leaf-push barrier as

$$\lambda = \left\lfloor \frac{1}{\ln 2} \mathcal{W}(n \ln \delta) \right\rfloor \quad (2)$$

where $\mathcal{W}()$ denotes the Lambert \mathcal{W} -function. Then, $D(S)$ can be encoded on at most $4 \lg(\delta)n + o(n)$ bits.

Note that the Lambert \mathcal{W} -function (or product logarithm) $\mathcal{W}(z)$ is defined as $z = \mathcal{W}(z)e^{\mathcal{W}(z)}$. The detailed proof, based on a counting argument, is deferred to the Appendix.

Next, we show that trie-folding compresses to within a constant factor of the zero-order entropy bound, subject to some reasonable assumptions on the alphabet. Furthermore, the constant term is improved from 7 (as of [1]) to 6.

Theorem 2: Let S be a string of length n and zero-order entropy H_0 , and set the leaf-push barrier as

$$\lambda = \left\lfloor \frac{1}{\ln 2} \mathcal{W}(nH_0 \ln 2) \right\rfloor. \quad (3)$$

Then, the expected size of $D(S)$ is at most $(6 + 2 \lg(1/H_0) + 2 \lg \delta)H_0n + o(n)$ bits.

Again, refer to the Appendix for the proof.

It turns out that the compression ratio depends on the specifics of the alphabet. For reasonable δ , say, $\delta = O(1)$ or $\delta = O(\text{polylog } n)$, the error $\lg \lg \delta$ is very small, and the bound gradually improves as H_0 increases, to the point that at maximum entropy $H_0 = \lg \delta$, we get precisely $6H_0n$. For extremely small entropy, however, the error $2 \lg(1/H_0)$ can become dominant as the overhead of the DAG outweighs the size of the very string in such cases.

C. Update

What remained to be done is to set the leaf-push barrier λ in order to properly balance between compression efficiency and

update complexity. Crucially, small storage can only be attained if the leaf-push barrier is chosen according to (3). Strikingly, we found that precisely this setting is the key to fast FIB updates as well.⁴

Herein, we only specify the update operation that changes an *existing* association for prefix a of prefix length p to the new label s or, within the string model, rewrites an entire block of symbols at the lowest level of the tree with a new one. Adding a new entry or deleting an existing one can be done in similar vein.

```

update (address  $a$ , integer  $p$ , label  $s$ , integer  $\lambda$ )
     $v \leftarrow D(T).root$ ;  $q \leftarrow 0$ 
    while  $q < p$ 
        if  $q \geq \lambda$  then  $v \leftarrow \text{decompress}(v)$ 
        if bits( $a$ ,  $q$ , 1) = 0 then  $v \leftarrow v.left$  else  $v \leftarrow v.right$ 
         $q \leftarrow q + 1$ 
        if  $p < \lambda$  then  $l(v) \leftarrow s$ ; return
         $w \leftarrow T.copy(v)$ ; repointer the parent of  $v$  to  $w$ ;  $l(w) \leftarrow s$ 
        POSTORDER-TRAVERSE-AT- $v(u: \text{get}(u.left.id, u.right.id))$ 
        trie_fold( $w$ , 0)
        for each parent  $u$  of  $w$ : level( $u$ )  $\geq \lambda$  do compress( $u$ )
    decompress (node  $v$ )
     $w \leftarrow \text{new\_node}$ ;  $w.id \leftarrow v.id$ 
    if  $v \in L_D$  then  $l(w) \leftarrow l(v)$ 
    else  $w.left \leftarrow v.left$ ;  $w.right \leftarrow v.right$ 
    get( $v.left.id, v.right.id$ )
    re-pointer the parent of  $v$  to  $w$ ; return  $w$ 
    
```

First, we walk down and decompress the DAG below λ along the path traced out by the successive bits of a until we reach level p . The `decompress` routine copies a node out from the DAG and removes the reference wherever necessary. At this point, if $p < \lambda$, then we simply update the label and we are ready. Otherwise, we replace the subtree with a new copy of the corresponding subtree from T , taking care of calling `get` on the descendants to remove dangling references, and we set the label on the root of the new copy to s . Then, we recompress the portions of the prefix DAG affected by the change by calling `trie_fold` and then calling `compress` on all the nodes along the upstream path towards to root.

Theorem 3: If the leaf-push barrier λ is set as (3), then update on $D(T)$ terminates in $O(W(1 + (1/H_0)))$ time.

Proof: If $p < \lambda$, then updating a single entry can be done in $O(W)$ time. If, on the other hand, $p \geq \lambda$, then update

⁴Note that (3) transforms into (2) at maximum entropy.

TABLE I

RESULTS FOR $XBW-b$ AND TRIE-FOLDING ON *access*, *core*, AND SYNTHETIC (*syn.*) FIBS: NAME, NUMBER OF PREFIXES N , AND NEXT-HOPS δ ; SHANNON ENTROPY OF THE NEXT-HOP DISTRIBUTION H_0 ; FIB INFORMATION-THEORETIC LIMIT I , ENTROPY E , AND $XBW-b$ AND PREFIX DAG SIZE (pDAG, $\lambda = 11$) IN KILOBYTES; COMPRESSION EFFICIENCY ν ; AND BITS/PREFIX EFFICIENCY FOR $XBW-b$ (η_{XBW-b}) AND TRIE-FOLDING (η_{pDAG})

	FIB	N	δ	H_0	I	E	$XBW-b$	pDAG	ν	η_{XBW-b}	η_{pDAG}
access	taz	410,513	4	1.00	94	56	63	178	3.17	1.12	3.47
	hbone	410,454	195	2.00	356	142	149	396	2.78	1.05	7.71
	access(d)	444,513	28	1.06	206	90	100	370	4.1	1.12	6.65
	access(v)	2,986	3	1.22	2.8	2.2	2.5	7.5	3.4	1.13	20.23
	mobile	21,783	16	1.08	0.8	0.4	1.1	3.6	8.71	2.36	1.35
core	as1221	440,060	3	1.54	130	115	111	331	2.86	2.03	6.02
	as4637	219,581	3	1.12	52	41	44	129	3.13	1.62	4.69
	as6447	445,016	36	3.91	375	277	277	748	2.7	5	13.45
	as6730	437,378	186	2.98	421	209	213	545	2.6	3.91	9.96
syn.	fib_600k	600,000	5	1.06	257	157	179	462	2.93	1.14	6.16
	fib_1m	1,000,000	5	1.06	427	261	297	782	2.99	1.14	6.26

visits at most $W + 2^{W-\lambda} \leq W + (W/H_0)$ nodes, using that $\lambda \geq W - \lg(W/H_0)$ whenever λ is as (3). ■

In summary, under mild assumptions on the label distribution, a prefix DAG realizes the Shannon entropy up to a small factor and allows indexing arbitrary elements and updates to any entry in roughly $O(\log n)$ time. As such, it is in fact a dynamic, entropy-compressed string self-index. As far as we are aware, this is the first pointer machine of this kind, as the rest of compressed string-indexes are pointerless. Regrettably, both the space bound and the update complexity weaken when the label distribution is extremely biased, i.e., when H_0 is very small. As we argue in Section V, though, this rarely causes problems in practice.

V. NUMERICAL EVALUATIONS

At this point, we have yet to demonstrate that the appealing theoretical properties of compressed FIBs indeed manifest as practical benefits. For this reason, we conducted a series of numerical evaluations with the goal to quantify the compressibility of real IP FIBs and see how our compressors fare. It was *not* our intention, however, to compare to other FIB storage schemes from the literature, let alone evince that ours is the fastest or the most efficient one. After all, information-theoretic space bounds are purposed precisely at making such comparisons unnecessary, serving as analytically justified ground truth. Instead, our motivation is merely to demonstrate that FIB compression allows to reduce memory tax without any compromise on the efficiency of longest prefix match or FIB updates.

For the evaluations, we coded up a full-fledged Linux prototype, where FIB compression and update run in user space and IP lookup is performed by a custom kernel module embedded in the kernel's IP stack. The code executed on a single core of a 2.50-GHz Intel Core i5 CPU, with 2x32 kB L1 data cache, 256 kB L2 cache, and 3 MB L3 cache.

Research on IP FIB data structures has for a long time been plagued by the unavailability of real data, especially from the Internet core. Alas, we could obtain only five FIB instances from real IP routers, each from the access: *taz* and *hbone* are from a university access, *access(d)* is from a default and *access(v)* from a virtual instance of a service provider's router, and *mobile* is from a mobile operator's access (see Table I). The first three are in the DFZ; the rest contain default routes. Apart from these, however, what is available publicly is RIB dumps from BGP collectors, like RouteViews or looking glass servers (named *as** in the data set). Unfortunately, these only

very crudely model real FIBs because collectors run the BGP best-path selection algorithm on their peers and these adjacencies differ greatly from real next-hops on production routers. We experimented with heuristics to restore the original next-hop information (e.g., set next-hop to first AS-hop), but the results were basically the same. Thus, these FIBs are included only for reference. All FIBs were collected at various times during 2012. We also used two randomly generated FIBs, one of 600 000 (*fib_600k*) and another of 1 million prefixes (*fib_1m*), to future-proof our results. These synthetic FIBs were generated by iterative random prefix splitting and setting next-hops according to a truncated Poisson distribution with parameter 3/5 ($H_0 = 1.06$, $\delta = 4$).

A. Update Complexity

First, we set out to determine a good setting for the leaf-push barrier λ . Recall that λ was introduced to balance between the compression efficiency and update complexity (also recall that no such compromise exists between compression and lookup). Our theoretical results provide the essential pointers to set λ [see (2) and (3)], but these are for compressing strings over complete binary trees. IP FIBs, however, are not complete.

We exercised the memory footprint versus update complexity tradeoff by varying λ between 0 and 32. The update time was measured over two update sequences: a random one with IP prefixes uniformly distributed on $[0, 2^{32} - 1]$ and prefix lengths on $[0, 32]$, and a BGP-inspired one corresponding to a real BGP router log taken from RouteViews. Here, we treated all BGP prefix announcements as generating an FIB update, with a next-hop selected randomly according to the next-hop distribution of the FIB. The results are mean values over 15 runs of 7500 updates, each run roughly corresponding to 15 min worth of BGP churn.

Herein, we only show the results for the *taz* FIB instance in Fig. 5. The results suggest that standard prefix trees (reproduced by the setting $\lambda = 32$), while pretty fast to update, occupy a lot of space. Fully compressed DAGs ($\lambda = 0$), in contrast, consume an order of magnitude less space but are expensive to modify. There is a domain, however, at around $5 \leq \lambda \leq 12$, where we win essentially all the space reduction and still handle about 100 000 updates per second (that is roughly 2.5 h of BGP update load). What is more, *the space-time tradeoff only exists for the synthetic, random update sequence, but not for BGP updates*. This is because BGP updates are heavily biased toward




Fig. 5. Update time versus memory footprint on `taz` for random and BGP update sequences.

longer prefixes (with a mean prefix length of 21.87), which implies that the size of leaf-pushed subtrees needed to be repacked per update is usually very small, and hence update complexity is relatively insensitive to λ . Based on these considerations, we set $\lambda = 11$ for the rest of the evaluations.

B. Storage Size

Storage size results are given in Table I. Notably, real FIBs that contain only a few next-hops compress down to about 60–150 kB with *XBW-b* at 1–2 bits/prefix efficiency, and only about 2–3 times more with trie-folding. This is chiefly attributed to the small next-hop entropy, indicating the presence of a dominant next-hop. Core FIBs, on the other hand, exhibit considerably larger next-hop entropy, with *XBW-b* transforms in the range of 100–300, and prefix DAGs in 330–700 kB. Recall, however, that these FIBs exhibit unrealistic next-hop distribution. Curiously, even the extremely large FIB of 1 million prefixes shrinks below 300 kB (800 kB with trie-folding). In contrast, small instances compress poorly, as it is usual in data compression. Finally, we observe that many *FIBs show high next-hop regularity* (especially the real ones), reflected in the fact that entropy bounds are 20%–40% smaller than the information-theoretic limit. *XBW-b very closely matches entropy bounds, with trie-folding off by only a small factor.*

We also studied compression ratios on synthetic FIBs, whose entropy was controlled by us. In particular, we regenerated the next-hops in `access` (`d`) according to Bernoulli distribution: A first next-hop was set with probability p and another with probability $1 - p$. Then, varying p in $[0, 1/2]$ we observed the FIB entropy, the size of the prefix DAG, and the compression efficiency ν , i.e., the factor between the two (see Fig. 6). We found that *the efficiency is around 3 and, in line with our theoretical analysis, degrades as the next-hop distribution becomes extremely biased.* This, however, never occurs in reality (see again Table I). We repeated the analysis in the string compression model: Here, the FIB was generated as a complete binary trie with a string of 2^{17} symbols written on the leaves, again chosen by a Bernoulli distribution, and this was then compressed with trie-folding (see Fig. 7, with *XBW-b* omitted). The observations are similar, with compression efficiency again varying around 3 and the spike at low entropy more prominent.

C. Lookup Complexity

Finally, we tested IP lookup performance on real software and hardware prototypes. Our software implementations run inside the Linux kernel’s IP forwarding engine. For this, we hijacked the kernel’s network stack to send IP lookup requests to our custom kernel module, working from a serialized blob generated by the FIB encoders. Our *XBW-b* lookup engine uses a kernel




Fig. 6. Size and compression efficiency ν over FIBs with Bernoulli distributed next-hops as the function of parameter p .




Fig. 7. Size and compression efficiency ν over strings with Bernoulli distributed symbols as the function of parameter p .

port of the RRR bitstring index [12] and the Huffman-shaped WaveletTree [13] from `libcds` [9]. Trie-folding was coded in pure C. We used the standard trick to collapse the first $\lambda = 11$ levels of the prefix DAGs in the serialized format [41], as this greatly eases implementation and improves lookup time with practically zero effect on updates. We also experimented with the Linux-kernel’s stock `fib_trie` data structure, an adaptive level- and path-compressed multibit trie-based FIB, as a reference implementation [25]. Lastly, we also realized the prefix DAG lookup algorithm in hardware, on a Xilinx Virtex-II Pro 50 FPGA with 4.5 MB of synchronous SRAM representing the state of the art almost 10 years ago. The hardware implementation uses the same serialized prefix DAG format as the software code. All tests were run on the `taz` instance.

For the software benchmarks we used the standard Linux network micro-benchmark tool `kbench` [62], which calls the FIB lookup function in a tight loop and measures the execution time with nanosecond precision. We modified `kbench` to take IP addresses from a uniform distribution on $[0, 2^{32} - 1]$ or, alternatively, from a packet trace in the “CAIDA Anonymized Internet Traces 2012” data set [61]. The route cache was disabled. We also measured the rate of CPU cache misses by monitoring the `cache-misses` CPU performance counter with the `perf(1)` tool. For the hardware benchmark, we mirrored `kbench` functionality on the FPGA, calling the lookup logic repeatedly on a list of IP addresses statically stored in the SRAM, and we measured the number of clock ticks needed to terminate the test cycle.

The results are given in Table II. On the software side, the most important observations are as follows. The prefix DAG, taking only about 180 kB of memory, is most of the time accessed from the cache, while `fib_trie` occupies an impressive 26 MB and so it does not fit into fast memory. Thus, even though the number of memory accesses needed to execute an IP lookup is smaller with `fib_trie`, as most of these go to slower memory, the prefix DAG supports about three times as many

TABLE II

LOOKUP BENCHMARK WITH *XBW-b*, PREFIX DAGS, *fib_trie*, AND THE FPGA IMPLEMENTATION ON *taz*: SIZE, AVERAGE, AND MAXIMUM DEPTH; AND MILLION LOOKUP PER SECOND, LOOKUP TIME IN CPU CYCLES, AND CACHE MISSES PER PACKET OVER RANDOM IP ADDRESSES (*rand.*) AND ADDRESSES TAKEN FROM THE TRACE [61] (*trace*)

		Linux			HW
		<i>XBW-b</i>	pDAG	<i>fib_trie</i>	FPGA
	size [Kbyte]	106	178	26,698	178
	average depth	–	3.7	2.42	–
	maximum depth	–	21	6	–
<i>rand.</i>	million lookup/sec	0.033	12.8	3.23	6.9
	CPU cycle/lookup	73940	194	771	7.1
	cache miss/packet	0.016	0.003	3.17	–
<i>trace</i>	million lookup/sec	0.037	13.8	5.68	6.9
	CPU cycle/lookup	67200	180	438	7.1
	cache miss/packet	0.016	0.003	0.29	–

lookups per second. Accordingly, *not just that FIB space reduction does not ruin lookup performance, but it even improves it*. In other words, there is no space–time tradeoff involved here. The address locality in real IP traces helps *fib_trie* performance to a great extent, as *fib_trie* can keep lookup paths to popular prefixes in cache. In contrast, the prefix DAG is pretty much insensitive to the distribution of lookup keys. Finally, we see that *XBW-b* is a distant third from the tested software lookup engines, suggesting that the constant in the lookup complexity is indeed prohibitive in practice and that our lookup code exercises some pretty pathologic code path in *libcds*.

The real potential of trie–folding is most apparent with our hardware implementation. The FPGA design executes a single IP lookup in just 7.1 clock cycles on average since the prefix DAG fits nicely into the SRAM running synchronously with the logic. This is enough to roughly 7 million IP lookups per second even on our rather ancient FPGA board. On a modern FPGA or ASIC, however, with clock rates in the gigahertz range, our results indicate that prefix DAGs could be scaled to hundreds of millions of lookups per second at a terabit line speed.

We also measured packet throughput using the *udpflood* macro-benchmark tool [62]. This tool injects UDP packets into the kernel destined to a dummy network device, which makes it possible to run benchmarks circumventing network device drivers completely. The results were similar as above, with prefix DAGs supporting consistently 2–3 times larger throughput than *fib_trie*.

VI. RELATED WORK

In line with the unprecedented growth of the routed Internet and the emerging scalability concerns thereof [44], [48], [49], finding efficient FIB representations has been a heavily researched question in the past and, judging from the substantial body of recent work [33], [36], [40], [41], still does not seem to have been solved completely.

Trie-based FIB schemes date back to the BSD kernel implementation of Patricia trees [37]. This representation consumes a massive 24 B per node, and a single IP lookup might cost 32 random memory accesses. Storage space and search time can be saved on by expanding nodes’ strides to obtain a multibit trie [26]; see, e.g., controlled prefix expansion [35], [39], level- and path-compressed tries [25], Lulea [29], Tree Bitmaps [32] and successors [27], [38], etc. Another approach is to shrink

the routing table itself by cleverly relabeling the tree to contain the minimum number of entries (see ORTC and derivatives [31], [40]). In our view, trie-folding is complementary to these schemes, as it can be used in combination with basically any trie-based FIB representation, realizing varying extents of storage space reduction.

Further FIB representations include hash-based schemes [27], [42], dynamic pipelining [34], CAMs [24], Bloom filters [30], binary search trees and search algorithms [28], [41], massively parallelized lookup engines [33], [41], FIB caching [36], and different combinations of these (see the textbook [23]). None of these comes with information-theoretic space bounds, although next-hop entropy itself appears in [40], but no analytical evaluation ensued. In contrary, *XBW-b* and trie-folding come with *theoretically proven* space limits, and thus *predictable memory footprint*. The latest reported FIB size bounds for >400 K prefixes range from 780 kB (DXR, [41]) to 1.2 MB (SMALTA, [40]). *XBW-b* improves this to just 100–300 kB, which easily fits into today’s SRAMs or can be realized right in chip logic with modern FPGAs.

Compressed data structures have been in the forefront of theoretical computer science research [3], [6]–[13], ever since Jacobson in his seminal work [56] defined succinct encodings of trees that support navigational queries in optimal time within information-theoretically limited space. Jacobson’s bitmap-based techniques later found important use in FIB aggregation [27], [32], [38]. With the extensive use of bitmaps, *XBW-b* can be seen as a radical rethinking of these schemes, inspired by the state of the art in succinct and compressed data structures.

The basic idea of folding a labeled tree into a DAG is not particularly new; in fact, this is the basis of many tree compacting schemes [57], space-efficient ordered binary decision diagrams and deterministic acyclic finite state automata [59], common subexpression elimination in optimizing compilers [60], and it has also been used in FIB aggregation [35], [50], [58] earlier. Perhaps the closest to trie-folding is Shape graphs [50], where common subtrees, without regard to the labels, are merged into a DAG. However, this necessitates storing a giant hash for the next-hops, making updates expensive, especially considering that the underlying trie is leaf-pushed. Trie-folding, in contrast, takes labels into account when merging and also allows cheap updates.

VII. CONCLUSION

With the rapid growth of the Web, social networks, mobile computing, data centers, and the Internet routing ecosystem as a whole, the networking field is in a sore need of compact and efficient data representations. Today’s networking practice, however, still relies on *ad hoc* and piecemeal data structures for basically all storage-sensitive and processing-intensive applications, of which the case of IP FIBs is just one salient example.

Our main goal in this paper was to advocate compressed data structures to the networking community, pointing out that space reduction does not necessarily hurt performance. Just the contrary: The smaller the space, the more data can be squeezed into fast memory, leading to faster processing. This lack of space–time tradeoff is already exploited to a great extent in information retrieval systems, business analytics, computational biology, and computational geometry, and we

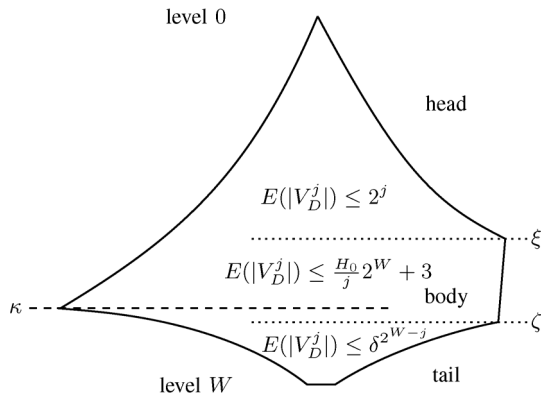


Fig. 8. Shape of the DAG as divided into three parts with bounds on the expected number of nodes at each level.

believe that it is just too appealing not to be embraced in networking as well. This paper is intended as a first step in that direction, demonstrating the basic information-theoretic and algorithmic techniques needed to attain entropy bounds, on the simple but relevant example of IP FIBs. Our techniques could then prove instructive in designing compressed data structures for other large-scale data-intensive networking applications, like OpenFlow and MPLS label tables, Ethernet self-learning MAC tables, BGP RIBs, access rules, log files, or peer-to-peer paths [63].

Accordingly, this paper can in no way be complete. For instance, we deliberately omitted IPv6 for brevity, even though storage burden for IPv6 is getting just as pressing as for IPv4 [38]. We see no reasons why our techniques could not be adapted to IPv6, but exploring this area in depth is for further study. Multibit prefix DAGs also offer an intriguing future research direction, for their potential to reduce storage space as well as improving lookup time from $O(W)$ to $O(\log W)$. On a more theoretical front, FIB entropy lends itself as a new tool in compact routing research, the study of the fundamental scalability of routing algorithms. We need to see why IP FIBs contain vast redundancy, track down its origins, and eliminate it to enforce zero-order entropy bounds right at the level of the routing architecture. To what extent this argumentation can then be extended to higher-order entropy is, for the moment, unclear at best.

APPENDIX

Proof of Theorem 1: As $D(S)$ is derived from a complete binary tree, the number of nodes V_D^j of $D(S)$ at level j is at most $|V_D^j| \leq 2^j$, and each node at level j corresponds to a 2^{W-j} -long substring of S so $|V_D^j| \leq \delta^{2^{W-j}}$. Let κ denote the intersection of the two bounds $2^\kappa = \delta^{2^{W-\kappa}}$, which gives

$$\kappa 2^\kappa = 2^W \log_2(\delta) = n \log_2(\delta). \quad (4)$$

Set the leaf push barrier at $\lambda = \lfloor \kappa \rfloor = \lfloor (1/\ln 2)\mathcal{W}(n \ln \delta) \rfloor$ where $\mathcal{W}()$ denotes the Lambert W -function. The left side of Fig. 8 is an illustration of the shape of DAG. Above level λ , we have

$$\sum_{j=0}^{\lambda} |V_D^j| = \sum_{j=0}^{\lambda} 2^j = 2^{\lambda+1} - 1 \leq 2 \cdot 2^\kappa$$

nodes; at level $\lambda + 1$, we have 2^κ nodes at maximum; at $\lambda + 2$, there are $|V_D^{\lambda+2}| \leq \delta^{2^{W-\lambda-2}} \leq \sqrt{2^\kappa}$ nodes; and finally, below

level $\lambda + 3$, we have an additional $\sqrt{2^\kappa}$ nodes at most as levels shrink as $|V_D^{j+1}| \leq \sqrt{|V_D^j|}$ downwards in $D(S)$. Finally, setting the pointer size at $\lceil \kappa \rceil$ bits and summing up the above yields that the size of $D(S)$ is at most

$$\begin{aligned} & \left(2 + 2 \left(1 + \frac{2}{\sqrt{2^\kappa}} \right) \right) \lceil \kappa \rceil 2^\kappa + (2 \cdot 2^\kappa + \delta) \log_2 \delta \\ & = 4n \log_2(\delta) + o(n) \end{aligned}$$

bits, using the fact that $\lceil \kappa \rceil 2^\kappa = n \log_2(\delta) + o(n)$ by (4) and the number of labels stored in the DAG is at most $2 \cdot 2^\kappa$ above the leaf-push barrier and further δ below it. ■

Proof of Theorem 2: Let $E(|V_D^j|)$ denote the expected number of nodes at level j of the DAG. We shall use the following bounds on $E(|V_D^j|)$ to prove the claim:

$$E(|V_D^j|) \leq \min \left\{ 2^j, \frac{H_0}{j} 2^W + 3, \delta^{2^{W-j}} \right\}. \quad (5)$$

Here, the first and the last bounds are from the proof of Theorem 1, while the second one is obtained below by treating the problem as a coupon collector's problem on the subtrees of $D(T)$ at level j . Suppose that we are given a set of coupons C , each coupon representing a string of length 2^{W-j} on the alphabet Σ of size δ and entropy H_0 , and we draw a coupon o with probability $p_o : o \in C$. Let $H_C = \sum_{o \in C} p_o \log_2(1/p_o) = H_0 2^{W-j}$, let V denote the set of coupons after $m = 2^j$ draws, and suppose $m \geq 3$.

Lemma 6: $E(|V|) \leq (m/\log_2(m))H_C + 3$.

Using this lemma, we have that the expected number of nodes at the j th level of $D(S)$ is at most $E(|V_D^j|) \leq (2^j/\log_2(2^j))H_0 2^{W-j} + 3 = (H_0/j)n + 3 = (H_0/j)2^W + 3$, which coincides with the second bound in (5). Note that here 2^j is an increasing function of j , while both $H_0 2^W/j + 3$ and $\delta^{2^{W-j}}$ are monotone decreasing functions.

Using these bounds, we divide the DAG into three parts (the ‘‘head,’’ ‘‘body,’’ and ‘‘tail’’) as illustrated at the right side of Fig. 8. Let ξ denote the intersection of the first two upper bounds, let ζ be that of the latter two, and let κ be the level where 2^j and $\delta^{2^{W-j}}$ meet. It is easy to see that the relation between these three values can only be $\xi \leq \kappa \leq \zeta$ or $\xi \geq \kappa \geq \zeta$. We discuss these two cases separately.

Case 1: $\xi \leq \kappa \leq \zeta$

The three parts of the DAG are as follows (again, see the right side of Fig. 8):

- head* for levels $0, \dots, \lfloor \xi \rfloor$;
- body* for levels $\lfloor \xi \rfloor + 1, \dots, \lfloor \zeta \rfloor - 1$;
- tail* for levels $\lfloor \zeta \rfloor, \dots, W$.

In the following, we give upper bounds on the expected number of the nodes in the head, the tail, and the body of the DAG. Set the leaf-push barrier at $\lambda = \lfloor \xi \rfloor$.

First, the expected number of nodes in the head is

$$\begin{aligned} \sum_{j=0}^{\lfloor \xi \rfloor} E(|V_D^j|) & \leq \sum_{j=0}^{\lfloor \xi \rfloor} 2^j = 2^{\lfloor \xi \rfloor + 1} - 1 \\ & < 2 \cdot 2^\xi = 2 \left(\frac{H_0}{\xi} 2^W + 3 \right) = 2M \quad (6) \end{aligned}$$

where $M = ((H_0/\xi)2^W + 3)$.

Second, for the size of the tail, we have

$$\begin{aligned}
 \sum_{j=\lceil \zeta \rceil}^W E\left(|V_D^j|\right) &\leq \sum_{j=\lceil \zeta \rceil}^W \delta^{2^{W-j}} = \sum_{j=1}^{W-\lceil \zeta \rceil} \delta^{2^j} \\
 &= \delta^{2^{W-\lceil \zeta \rceil}} + \sum_{j=1}^{W-\lceil \zeta \rceil-1} \delta^{2^j} \\
 &< \delta^{2^{W-\zeta}} + \sum_{i=1}^{2^{W-\lceil \zeta \rceil-1}} \delta^i \\
 &= \delta^{2^{W-\zeta}} + \delta^{2^{W-\lceil \zeta \rceil-1}+1} - 2 \\
 &< \delta^{2^{W-\zeta}} + \delta \sqrt{\delta^{2^{W-\zeta}}} \\
 &= \delta^{2^{W-\zeta}} \left(1 + \frac{\delta}{\sqrt{\delta^{2^{W-\zeta}}}}\right) \\
 &= \left(\frac{H_0}{\zeta} 2^W + 3\right) (1 + \epsilon_1) \\
 &< \left(\frac{H_0}{\xi} 2^W + 3\right) (1 + \epsilon_1) = M(1 + \epsilon_1)
 \end{aligned} \tag{7}$$

where

$$\epsilon_1 = \frac{\delta}{\sqrt{\delta^{2^{W-\zeta}}}} \leq \frac{\delta}{\sqrt{\delta^{\frac{W-\log_2 W/H_0}{\log_2 \delta}}}} = \frac{\delta \sqrt{\frac{W}{H_0}}}{2^{\frac{W}{2}}}$$

which tends to zero if W goes to infinity.

Third, for the number of nodes in the body, we write

$$\begin{aligned}
 \sum_{j=\lfloor \xi \rfloor + 1}^{\lceil \zeta \rceil - 1} E\left(|V_D^j|\right) &\leq \sum_{j=\lfloor \xi \rfloor + 1}^{\lceil \zeta \rceil - 1} \left(\frac{H_0}{j} 2^W + 3\right) \\
 &< \sum_{j=\lfloor \xi \rfloor + 1}^{\lceil \zeta \rceil - 1} \left(\frac{H_0}{\xi} 2^W + 3\right) \\
 &= (\lceil \zeta \rceil - 1 - \lfloor \xi \rfloor) M \\
 &\leq (\zeta - (\xi - 1)) M = (1 + \zeta - \xi) M.
 \end{aligned}$$

Lemma 7: The following bounds on ξ , ζ , and κ apply:

$$\xi \geq W - \log_2 W/H_0 \tag{8}$$

$$\zeta \leq W - \log_2(W - \log_2 W/H_0) + \log_2 \log_2(\delta) \tag{9}$$

$$\kappa \leq W - \log_2(W - \log_2 W) + \log_2 \log_2 \delta. \tag{10}$$

Using (8) and (9) for the body, we write

$$\begin{aligned}
 \sum_{j=\lfloor \xi \rfloor + 1}^{\lceil \zeta \rceil - 1} E\left(|V_D^j|\right) &\leq (1 + \zeta - \xi) M \\
 &\leq \left(1 + \log_2\left(\frac{W}{H_0}\right) - \log_2\left(\frac{W - \log_2\left(\frac{W}{H_0}\right)}{\log_2(\delta)}\right)\right) M \\
 &= (1 + \log \log_2(\delta) - \log_2 H_0 + \epsilon_2) M \tag{11}
 \end{aligned}$$

where

$$\epsilon_2 = \log_2 \frac{W}{W - \log_2\left(\frac{W}{H_0}\right)} \xrightarrow{W \rightarrow \infty} 0.$$

Choose the pointer size to $\lceil \kappa \rceil$ bits, using that the DAG contains at most $2^{\lceil \kappa \rceil}$ nodes at its broadest level. For the head, we need one pointer for each node, while for the rest we need two. Summing up with (7), (6), and (11), we get the following bound on the number of pointers:

$$\begin{aligned}
 (2 + 2(1 + \epsilon) + 2(1 - \log_2 H_0 + \log \log_2(\delta) + \epsilon_2)) M \\
 = \left(6 + 2 \log_2 \frac{\log_2 \delta}{H_0} + 2(\epsilon_1 + \epsilon_2)\right) M.
 \end{aligned}$$

We have to store labels above the barrier and at the bottom level, which is at most $2M + \delta$ labels, hence the average number of the bits is at most

$$\left(6 + 2 \log_2 \frac{\log_2 \delta}{H_0} + 2(\epsilon_1 + \epsilon_2)\right) M \lceil \kappa \rceil + (2M + \delta) \log_2 \delta.$$

Using (8) for ξ and (10) for κ , we have

$$\begin{aligned}
 M \lceil \kappa \rceil &= \lceil \kappa \rceil \left(\frac{H_0}{\xi} 2^W + 3\right) \leq \frac{\kappa + 1}{\xi} H_0 n + 3 \lceil \kappa \rceil \\
 &\leq \left(\frac{W - \log_2 \frac{W - \log_2 W}{\log_2 \delta} + 1}{W - \log_2 W/H_0}\right) H_0 n + 3 \lceil \kappa \rceil \\
 &= \left(1 + \frac{\log_2 W/H_0 - \log_2 \frac{W - \log_2 W}{\log_2 \delta} + 1}{W - \log_2 W/H_0}\right) H_0 n + 3 \lceil \kappa \rceil \\
 &= H_0 n + o(n).
 \end{aligned}$$

In summary, for the expected size of the DAG, we get $(6 - 2 \log_2 H_0 + 2 \log \log_2(\delta)) H_0 n + o(n)$ bits since $(2M + \delta) \log_2 \delta = o(n)$.

Case 2: $\zeta \leq \kappa \leq \xi$

In this case the DAG contains only the head and tail parts. According to Theorem 1, we get the upper bound $5n \log_2 \delta + o(n)$ on the number of bits. As $\kappa \leq \xi$, we have

$$\frac{n \log_2 \delta}{\kappa} = 2^\kappa \leq \frac{H_0}{\kappa} n + 3.$$

Hence, the upper bound on the number of required bits is $5n \log_2 \delta + o(n) \leq 5H_0 n + o(n) < (6 + 2 \log(\log_2(\delta)/H_0)) H_0 n + o(n)$. ■

Proof of Lemma 6: The probability of having coupon o in V is $\Pr(o \in V) = 1 - (1 - p_o)^m$ and so the expected cardinality of V is $E(|V|) = \sum_{o \in C} (1 - (1 - p_o)^m)$. The right-hand side of the statement of the lemma is $(m/\log_2(m)) \sum_{o \in C} p_o \log_2(1/p_o) + 3$. The claim holds if $\forall o \in C$

$$p_o < \frac{1}{e} \Rightarrow 1 - (1 - p_o)^m \leq \frac{m}{\log_2(m)} p_o \log_2 \frac{1}{p_o}. \tag{12}$$

First, assume $m \geq 1/p_o$. As the right-hand side is a monotone increasing function of m when $e \leq 1/p_o \leq m$

$$1 - (1 - p_o)^m \leq 1 = \frac{\frac{1}{p_o}}{\log_2\left(\frac{1}{p_o}\right)} p_o \log_2 \frac{1}{p_o} \leq \frac{m}{\log_2(m)} p_o \log_2 \frac{1}{p_o}.$$

Otherwise, if $m < 1/p_o$, then let $x = \log_{1/p_o} m$. Note that $0 < x < 1$. After substituting $m = 1/p_o^x$, we have

$$\begin{aligned} 1 - (1 - p_o)^{\frac{1}{p_o^x}} &\leq \frac{\frac{1}{p_o^x}}{\log_2 \left(\frac{1}{p_o^x} \right)} p_o \log_2 \left(\frac{1}{p_o} \right) \\ &= \frac{\frac{1}{p_o^x}}{x \log_2 \left(\frac{1}{p_o} \right)} p_o \log_2 \left(\frac{1}{p_o} \right) \\ &= \frac{\frac{1}{p_o^x}}{x} p_o = \frac{1}{x p_o^{x-1}}. \end{aligned}$$

which can be reordered as

$$(1 - p_o)^{\frac{1}{p_o^x}} \geq 1 - \frac{1}{x p_o^{x-1}}.$$

Taking the $p_o^{x-1} > 0$ power of both sides, we get

$$(1 - p_o)^{\frac{1}{p_o}} \geq \left(1 - \frac{1}{x p_o^{x-1}} \right)^{p_o^{x-1}}.$$

Using that $x < 1$ and so $1/x > 1$, we see that the above holds if

$$(1 - p_o)^{\frac{1}{p_o}} \geq \left(1 - \frac{1}{p_o^{x-1}} \right)^{p_o^{x-1}}.$$

Note that $(1 - p_o)^{1/p_o}$ is monotone decreasing function, thus the inequality holds if $p_o \leq 1/p_o^{x-1}$, but this is true because $p_o^x \leq 1$. This proves (12) under the assumption $p_o < 1/e$. Note also that there are at most $3 > 1/e$ coupons for which (12) cannot be applied. ■

Proof of Lemma 7: To prove (8), for level $l = W - \log_2(W/H_0)$, we have

$$2^l = 2^{W - \log_2 W / H_0} \leq 2^{W - \log_2 W / H_0} = 2^W \frac{H_0}{W} < \frac{H_0}{l} 2^W + 3.$$

For (9), at level $l = W - \log_2((W - \log_2(W/H_0))/\log_2 \delta)$, we have

$$\begin{aligned} \delta^{2^{W-l}} &= \delta^{\frac{W - \log_2 \left(\frac{W}{H_0} \right)}{\log_2(\delta)}} = \delta^{\frac{\log_2 \left(\frac{H_0}{W} 2^W \right)}{\log_2(\delta)}} \\ &= \frac{H_0}{W} 2^W < \frac{H_0}{l} 2^W + 3. \end{aligned}$$

Finally, for (10) at level $l = W - \log_2((W - \log_2 W)/\log_2 \delta)$, we write

$$\delta^{2^{W-l}} = \delta^{\frac{W - \log_2 W}{\log_2 \delta}} = 2^{W - \log_2 W} \leq 2^l. \quad \blacksquare$$

ACKNOWLEDGMENT

The authors wish to thank B. Mihálka, Z. Csernátóy, G. Barna, L. Rónyai, A. Gulyás, G. Enyedi, A. Császár, G. Pongrácz, F. Claude, and S. Gorinsky for their invaluable assistance. Special thanks go to J. Lu for pointing out the mistake in [1] regarding FIB entropy.

REFERENCES

- [1] G. Rétvári, J. Tapolcai, A. Kőrösi, A. Majdán, and Z. Heszberger, "Compressing IP forwarding tables: Towards entropy bounds and beyond," in *Proc. ACM SIGCOMM*, 2013, pp. 111–122.
- [2] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY, USA: Wiley-Interscience, 1991.
- [3] N. Ziviani, E. S. de Moura, G. Navarro, and R. Baeza-Yates, "Compression: A key for next-generation text retrieval systems," *IEEE Comput.*, vol. 33, no. 11, pp. 37–44, Nov. 2000.
- [4] J. Woods, "PPP deflate protocol," RFC 1979, 1996.
- [5] J. Bonwick and B. Moore, "ZFS—The last word in file systems," Sun Microsystems, 2004.
- [6] W.-K. Hon, R. Shah, and J. S. Vitter, "Compression, indexing, retrieval for massive string data," in *Proc. CPM*, 2010, pp. 260–274.
- [7] G. Navarro and V. Mäkinen, "Compressed full-text indexes," *Comput. Surv.*, vol. 39, no. 1, p. 2, 2007.
- [8] P. Ferragina and G. Manzini, "Opportunistic data structures with applications," in *Proc. IEEE FOCS*, 2000, pp. 390–398.
- [9] G. Navarro and F. Claude, "libcds: Compact data structures library," 2004 [Online]. Available: <http://libcds.sourceforge.net>
- [10] V. Mäkinen and G. Navarro, "Dynamic entropy compressed sequences and full-text indexes," *Trans. Algor.*, vol. 4, no. 3, pp. 32:1–32:38, 2008.
- [11] E. de Moura, G. Navarro, N. Ziviani, and R. Baeza-Yates, "Fast and flexible word searching on compressed text," *Trans. Inf. Syst.*, vol. 18, no. 2, pp. 113–139, 2000.
- [12] R. Raman, V. Raman, and S. S. Rao, "Succinct indexable dictionaries with applications to encoding k -ary trees and multisets," in *Proc. ACM-SIAM SODA*, 2002, pp. 233–242.
- [13] P. Ferragina, G. Manzini, V. Mäkinen, and G. Navarro, "Compressed representations of sequences and full-text indexes," *Trans. Algor.*, vol. 3, no. 2, p. 20, 2007.
- [14] S. Vigna and P. Boldi, "MG4J: Managing gigabytes for Java," 2007 [Online]. Available: <http://mg4j.di.unimi.it>
- [15] LuceneTransform, "Transparent compression for Apache Lucene," 2012 [Online]. Available: <http://code.google.com/p/lucenetransform>
- [16] I. H. Witten, A. Moffat, and T. C. Bell, *Managing Gigabytes: Compressing and Indexing Documents and Images*. San Mateo, CA, USA: Morgan Kaufmann, 1999.
- [17] P. M. Tolani and J. R. Haritsa, "XGRIND: A query-friendly XML compressor," in *Proc. ICDE*, 2002, pp. 225–234.
- [18] P. Ferragina, F. Luccio, G. Manzini, and S. Muthukrishnan, "Compressing and indexing labeled trees, with applications," *J. ACM*, vol. 57, no. 1, pp. 1–33, 2009.
- [19] WebGraph, "A framework for graph compression," [Online]. Available: <http://webgraph.di.unimi.it>
- [20] J. Rossignac, "Edgebreaker: Connectivity compression for triangle meshes," *IEEE Trans. Visual Comput. Graphics*, vol. 5, no. 1, pp. 47–61, Jan.–Mar. 1999.
- [21] H. Sun, O. Ozturk, and H. Ferhatosmanoglu, "CoMRI: A compressed multi-resolution index structure for sequence similarity queries," in *Proc. IEEE CSB*, 2003, p. 553.
- [22] G. Navarro and J. Tarhio, "LZgrep: A Boyer-Moore string matching tool for Ziv-Lempel compressed text," *Softw. Pract. Exper.*, vol. 35, no. 12, pp. 1107–1130, 2005.
- [23] W. Wu, *Packet Forwarding Technologies*. New York, NY, USA: Auerbach, 2008.
- [24] A. McAuley and P. Francis, "Fast routing table lookup using CAMs," in *Proc. IEEE INFOCOM*, 1993, pp. 1382–1391.
- [25] S. Nilsson and G. Karlsson, "IP-address lookup using LC-tries," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 6, pp. 1083–1092, Jun. 1999.
- [26] G. Cheung and S. McCanne, "Optimal routing table design for IP address lookups under memory constraints," in *Proc. IEEE INFOCOM*, 1999, pp. 1437–1444.
- [27] M. Bando, Y.-L. Lin, and H. J. Chao, "FlashTrie: Beyond 100-Gb/s IP route lookup using hash-based prefix-compressed trie," *IEEE/ACM Trans. Netw.*, vol. 20, no. 4, pp. 1262–1275, Aug. 2012.
- [28] P. Gupta, B. Prabhakar, and S. P. Boyd, "Near optimal routing lookups with bounded worst case performance," in *Proc. IEEE INFOCOM*, 2000, pp. 1184–1192.
- [29] M. Degermark, A. Brodnik, S. Carlsson, and S. Pink, "Small forwarding tables for fast routing lookups," in *Proc. ACM SIGCOMM*, 1997, pp. 3–14.
- [30] S. Dharmapurikar, P. Krishnamurthy, and D. E. Taylor, "Longest prefix matching using Bloom filters," in *Proc. ACM SIGCOMM*, 2003, pp. 201–212.
- [31] R. Draves, C. King, S. Venkatachary, and B. Zill, "Constructing optimal IP routing tables," in *Proc. IEEE INFOCOM*, 1999.
- [32] W. Eatherton, G. Varghese, and Z. Dittia, "Tree bitmap: Hardware/software IP lookups with incremental updates," *Comput. Commun. Rev.*, vol. 34, no. 2, pp. 97–122, 2004.

- [33] S. Han, K. Jang, K. Park, and S. Moon, "PacketShader: A GPU-accelerated software router," in *Proc. ACM SIGCOMM*, 2010, pp. 195–206.
- [34] J. Hasan and T. N. Vijaykumar, "Dynamic pipelining: Making IP-lookup truly scalable," in *Proc. ACM SIGCOMM*, 2005, pp. 205–216.
- [35] I. Ioannidis and A. Grama, "Level compressed DAGs for lookup tables," *Comput. Netw.*, vol. 49, no. 2, pp. 147–160, 2005.
- [36] Y. Liu, S. O. Amin, and L. Wang, "Efficient FIB caching using minimal non-overlapping prefixes," *Comput. Commun. Rev.*, vol. 43, no. 1, pp. 14–21, Jan. 2012.
- [37] K. Sklower, "A tree-based packet routing table for Berkeley UNIX," Berkeley, CA, USA, Tech. Rep., 1991.
- [38] H. Song, J. Turner, and J. Lockwood, "Shape shifting tries for faster IP route lookup," in *Proc. IEEE ICNP*, 2005, pp. 358–367.
- [39] V. Srinivasan and G. Varghese, "Faster IP lookups using controlled prefix expansion," *Perform. Eval. Rev.*, vol. 26, no. 1, pp. 1–10, 1998.
- [40] Z. A. Uzmi *et al.*, "SMALTA: Practical and near-optimal FIB aggregation," in *Proc. ACM CoNEXT*, 2011, pp. 1–12.
- [41] M. Zec, L. Rizzo, and M. Mikuc, "DXR: Towards a billion routing lookups per second in software," *Comput. Commun. Rev.*, vol. 42, no. 5, pp. 29–36, 2012.
- [42] M. Waldvogel, G. Varghese, J. Turner, and B. Plattner, "Scalable high speed IP routing lookups," in *Proc. ACM SIGCOMM*, 1997, pp. 25–36.
- [43] A. Elmokashfi, A. Kvalbein, and C. Dovrolis, "BGP churn evolution: A perspective from the core," *IEEE/ACM Trans. Netw.*, vol. 20, no. 2, pp. 571–584, Apr. 2012.
- [44] G. Huston, "BGP routing table analysis reports," [Online]. Available: <http://bgp.potaroo.net/>
- [45] R. Bolla and R. Bruschi, "RFC 2544 performance evaluation and internal measurements for a Linux based open router," in *Proc. IEEE HPSR*, 2006, p. 6.
- [46] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," Internet Engineering Task Force, RFC 4984, 2007.
- [47] K. Fall, G. Iannaccone, S. Ratnasamy, and P. B. Godfrey, "Routing tables: Is smaller really much better?," in *Proc. ACM HotNets-VIII*, 2009.
- [48] X. Zhao, D. J. Pacella, and J. Schiller, "Routing scalability: An operator's view," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 8, pp. 1262–1270, Oct. 2010.
- [49] V. Khare *et al.*, "Evolution towards global routing scalability," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 8, pp. 1363–1375, Oct. 2010.
- [50] H. Song, M. S. Kodialam, F. Hao, and T. V. Lakshman, "Scalable IP lookups using shape graphs," in *Proc. IEEE ICNP*, 2009, pp. 73–82.
- [51] P. Francois, C. Filsfils, J. Evans, and O. Bonaventure, "Achieving sub-second IGP convergence in large IP networks," *Comput. Commun. Rev.*, vol. 35, no. 3, pp. 35–44, 2005.
- [52] G. Rétvári *et al.*, "Compressing IP forwarding tables for fun and profit," in *Proc. ACM HotNets-XI*, 2012, pp. 1–6.
- [53] EZChip, Campbell, CA, USA, "NP-4: 100-gigabit network processor for carrier ethernet applications," 2011 [Online]. Available: http://www.ezchip.com/Images/pdf/NP-4_Short_Brief_online.pdf
- [54] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, "In search of path diversity in ISP networks," in *Proc. ACM IMC*, 2003, pp. 313–318.
- [55] J. Choi, J. H. Park, P. Chun Cheng, D. Kim, and L. Zhang, "Understanding BGP next-hop diversity," in *Proc. IEEE INFOCOM Workshops*, 2011, pp. 846–851.
- [56] G. Jacobson, "Space-efficient static trees and graphs," in *Proc. IEEE FOCS*, 1989, pp. 549–554.
- [57] J. Katajainen and E. Mäkinen, "Tree compression and optimization with applications," *Int. J. Found. Comput. Sci.*, vol. 1, no. 4, pp. 425–447, 1990.
- [58] S. Stergiou and J. Jain, "Optimizing routing tables on systems-on-chip with content-addressable memories," in *Proc. Syst.-on-Chip*, 2008, pp. 1–6.
- [59] R. E. Bryant, "Symbolic boolean manipulation with ordered binary-decision diagrams," *Comput. Surv.*, vol. 24, no. 3, pp. 293–318, 1992.
- [60] J. Cocke, "Global common subexpression elimination," *SIGPLAN Not.*, vol. 5, no. 7, pp. 20–24, 1970.
- [61] P. Hick, K. C. Claffy, and D. Andersen, "CAIDA anonymized Internet traces," 2014 [Online]. Available: <http://www.caida.org/data/passive>
- [62] D. S. Miller, "net_test_tools," 2011 [Online]. Available: https://kernel.googlesource.com/pub/scm/linux/kernel/git/davem/net_test_tools
- [63] H. Madhyastha, E. Katz-Bassett, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane Nano: Path prediction for peer-to-peer applications," in *Proc. USENIX*, 2009, pp. 137–152.



Gábor Rétvári (S'03–M'05) received the M.Sc. and Ph.D. degrees in electrical engineering from the Budapest University of Technology and Economics (BME), Budapest, Hungary, in 1999 and 2007, respectively.

He is now a Senior Research Fellow with the High Speed Networks Laboratory, Department of Telecommunications and Media Informatics, BME. His research interests include any aspects of packet routing in data networks, particularly the networking applications of computational geometry, information theory, and network flows. He maintains several open-source scientific tools written in Perl, C, and Haskell.



János Tapolcai (M'05) received the M.Sc. degree in technical informatics and Ph.D. degree in computer science from the Budapest University of Technology and Economics (BME), Budapest, Hungary, in 2000 and 2005, respectively, and the D.Sc. degree in engineering science from Hungarian Academy of Sciences (MTA), Budapest, Hungary, in 2013.

Currently, he is an Associate Professor with the High-Speed Networks Laboratory, Department of Telecommunications and Media Informatics, BME.

Dr. Tapolcai is a winner of the MTA Lendület Program and Google Faculty Award.



Attila Kőrösi received the M.Sc. degree in mathematics from the Budapest University of Technology and Economics, Budapest, Hungary, in 2007.

He is an Assistant Research Fellow with the MTA-BME Information Systems Research Group, Budapest, Hungary. He has experiments in Wolfram Mathematica, C, JAVA, and Linux. He has published several papers on P2P-based multimedia services and their stochastic models. His research interest is mathematical modeling of routing protocols, FIB compression, and data centers.



András Majdán received the M.Sc. degree in technical informatics (majoring in network and services) from the Budapest University of Technology and Economics, Budapest, Hungary, in 2012, and is currently pursuing the Ph.D. degree in telecommunications and media informatics at the same university.

His research involves the compression of IP forwarding tables and particle filter-based mobile positioning.



Zalán Heszberger received the M.Sc. and Ph.D. degrees in electrical engineering from the Budapest University of Technology and Economics (BME), Budapest, Hungary, in 1997 and 2007, respectively.

Currently, he is an Associate Professor with the Department of Telecommunications and Media Informatics, BME. His main research interests are future Internet technologies and complex networks. Currently, he is working on clean-slate design Internet routing and network management algorithms.