

Theory and Practise of Monotone Minimal Perfect Hashing

Djamal Belazzougui*

Paolo Boldi†

Rasmus Pagh‡

Sebastiano Vigna§

Abstract

Minimal perfect hash functions have been shown to be useful to compress data in several data management tasks. In particular, *order-preserving* minimal perfect hash functions [12] have been used to retrieve the position of a key in a given list of keys: however, the ability to preserve any given order leads to an unavoidable $\Omega(n \log n)$ lower bound on the number of bits required to store the function. Recently, it was observed [1] that very frequently the keys to be hashed are sorted in their intrinsic (i.e., lexicographical) order. This is typically the case of dictionaries of search engines, list of URLs of web graphs, etc. We refer to this restricted version of the problem as *monotone minimal perfect hashing*. We analyse experimentally the data structures proposed in [1], and along our way we propose some new methods that, albeit asymptotically equivalent or worse, perform very well in practise, and provide a balance between access speed, ease of construction, and space usage.

1 Introduction

A minimal perfect hash function maps bijectively a set S of n keys into the set $\{0, 1, \dots, n-1\}$. The construction of such functions has been widely studied in the last years, leading to fundamental theoretical results such as [13, 14, 19].

From an application-oriented viewpoint, *order-preserving* minimal perfect hash functions have been used to retrieve the position of a key in a given list of keys [12, 26]. In [1] we note that all existing techniques for this task assume that keys can be provided in any order, incurring an unavoidable $\Omega(n \log n)$ lower bound on the number of bits required to store the function. However, very frequently the keys to be hashed are sorted in their intrinsic (i.e., lexicographical) order. This is typically the case of dictionaries of search engines, list of URLs of web graphs, etc. Thus, it is interesting to study *monotone minimal perfect hashing*—the problem of mapping each key of a lexicographically sorted set to its ordinal position.

In this paper our main goal is that of minimising the function description while maintaining quick (ideally,

constant-time) access. At SODA 2009 [1], we will present two solutions for the case where elements of S are taken from a universe of u elements. The first solution (based on longest common prefixes) provides $O((\log u)/w)$ access time, where w is the size of a machine word (so it is constant time if the string length is linear in the machine-word size), but requires $O(\log \log u)$ bits per element. The second solution (based on a *relative z-fast trie*) requires just $O(\log \log \log u)$ bits per element, but access requires $O((\log u)/w + \log \log u)$ steps.

In this paper, we present some new structures for this problem, and compare them experimentally with the above-mentioned ones. The purpose is twofold: first of all, we want to understand the constants hidden in the asymptotic estimates of [1]; second, we want to establish whether in a real-world scenario the new solutions proposed here have some practical advantage over the theoretically better ones.

To this purpose, we provide precise, big-Oh-free estimates of the number of bits occupied by each structure, which turn out to match very closely the number of bits required in the actual implementations. Moreover, we implement and engineer in detail all solutions in Java, and run them against large and real data.

In Section 2 we define precisely our problem, and in Section 3 we set up the tools that will be used in the rest of the paper. Then, in Sections 4, 5, 6, 7 and 8 we present data structures that provide different tradeoffs between space and time. Throughout the paper, we use the example of Figure 1 to illustrate the algorithms. Finally, in Section 10 we present experiments based on Java implementations of our data structures.

The code used for our experiments is distributed as part of the Sux4J project (<http://sux4j.dsi.unimi.it/>) under GNU Lesser General Public License. The lists of URLs used in the experimental section are available as part of the data sets distributed by the Laboratory for Web Algorithmics (<http://law.dsi.unimi.it/>), so as to make our experiments fully reproducible.

2 Definitions and notation

Sets and integers. We use von Neumann's definition and notation for natural numbers: $n = \{0, 1, \dots, n-1\}$. We thus freely write $f : m \rightarrow n$ for a function from the first m natural numbers to the first n natural numbers. We do the same with real numbers, with a slight abuse of notation,

*École Nationale Supérieure d'Informatique, Algiers, Algeria

†Università degli Studi di Milano, Italy

‡IT University of Copenhagen, Denmark

§Università degli Studi di Milano, Italy

s_0	0001001000000	s_3	0010011000000	s_6	0010011010100	s_9	0010011110110
s_1	0010010101100	s_4	0010011001000	s_7	0010011010101	s_{10}	0100100010000
s_2	0010010101110	s_5	0010011010010	s_8	0010011010110		

Figure 1: A toy example: $S = \{s_0, \dots, s_{10}\}$ is divided into three buckets of size three (except for the last one that contains just two elements), whose delimiters $D = \{s_2, s_5, s_8\}$ appear in boldface.

understanding a ceiling operator.

In the following, we will always assume that a universe u of integers, called *keys*, is fixed; this set may in many applications be infinite, but unless otherwise specified we will suppose that it is finite. The set u has a *natural order* which corresponds to the string lexicographical order of the $\log u$ -bit left-zero-padded binary representation. We assume, for sake of simplicity, that all strings have the same length $\log u$. We describe at the end of the paper the few modifications that are needed to state our results in terms of the *average* string length of a set of variable-length strings.

Given $S \subseteq u$ with $|S| = n$, and given m , an *m-bucket hash function* for S is any function $h : S \rightarrow m$. We say that:

- h is *perfect* iff it is injective;
- h is *minimal perfect* iff it is injective and $n = m$;
- h is *monotone* iff $x \leq y$ implies $h(x) \leq h(y)$ for all $x, y \in S$;
- h is *order-preserving* with respect to some total order \leq on U iff $x \leq y$ implies $h(x) \leq h(y)$ for all $x, y \in S$.

We would like to stress the distinction between *monotone* and *order-preserving* functions, which we introduce because the structures proposed in the literature as “order-preserving” [12] actually make it possible to *impose* any order on the keys. On the contrary, we are interested in the *existing*, standard lexicographical order on keys viewed as strings. The distinction is not moot because the lower bound $\Omega(n \log n)$ for order-preserving hashing does not hold for monotone hashing.

Notice that since monotone hash functions are a special case of order-preserving hash functions (applied to the natural order), any structure for the latter can be used to implement the former, but not vice versa.

Approximations. In this paper we purposely avoid asymptotic notation; our interest is in providing fairly precise estimates of the number of bits used by each structure. Nonetheless, we must allow some approximation if we want to control the size of our expressions. We will tacitly assume the following:

$$\log(\varepsilon + \log n) \approx \log \log n \quad \text{for small } \varepsilon$$

$$\log(\ln n) \approx \log \log n$$

$$\log n - \log \log n \approx \log n \quad \text{when appearing as a subexpression.}$$

Moreover, $o(n)$ components will be tacitly omitted.

3 Tools

The data structures described in this paper are based on a combination of techniques from two different threads of research: minimal perfect hashing based on random hypergraphs, and succinct data structures.

3.1 Rank and select. We will make extensive use of the two basic blocks of several succinct data structures—rank and select. Given a bit array (or bit string) $\mathbf{b} \in \{0, 1\}^n$, whose positions are numbered starting from 0, $\text{rank}_{\mathbf{b}}(p)$ is the number of ones up to position p , exclusive ($0 \leq p \leq n$), whereas $\text{select}_{\mathbf{b}}(r)$ is the position of the r -th one in \mathbf{b} , with bits numbered starting from 0 ($0 \leq r < \text{rank}_{\mathbf{b}}(n)$). These operations can be performed in constant time on a string of n bits using additional $o(n)$ bits [22, 8]. When \mathbf{b} is obvious from the context we shall omit the subscript.

3.2 Storing functions. In the rest of the paper we will frequently need to associate values to the key set S , that is, to store a function $f : S \rightarrow 2^r$ for some constant r . An obvious possibility is to store a minimal perfect hash function on S and use the resulting value to index a table of rn bits. Much better theoretical solutions were made available recently [6, 9, 29]: essentially, it is possible to evaluate a function in constant time storing just $rn + o(n)$ bits. Since we are interested in practical applications, however, we will use an extension of a technique developed by Majewski, Wormald, Havas and Czech [26] that has a slightly larger space usage, but has the advantage of being extremely fast, as it requires just the evaluation of three hash functions¹ plus three accesses to memory.

The technique developed in [26] was used to compute order-preserving hash functions in γrn bits, where $\gamma = 1.23$. Actually, the very same approach allows one to assign *any value* to the keys—emitting a distinct value in n for each element of S is just one of the possibilities. Thus, we will extend (and improve) the technique to store arbitrary functions in just $\gamma n + rn$ bits.

We recall briefly the technique of [26]. We start by building a random 3-hypergraph with γn nodes and n hyperedges—one per element of S —defined by three random² hash functions $h_1, h_2, h_3 : S \rightarrow \gamma n$. The choice of

¹Actually, we use Jenkins hashing [23], which provides three 64-bit hash values with a single evaluation.

²In this paper we make the *full randomness* assumption—our hash func-

$\gamma = 1.23$ makes the probability that the resulting graph is acyclic positive (see [5, 26] for details).

The acyclicity check computes a (sort-of) topological order of the hyperedges with the property that by examining the hyperedges in that order, at least one vertex of each hyperedge, the *hinge*, will have never appeared previously. We now assign values a_i to vertices, with the aim of obtaining that

$$f(x) = (a_{h_1(x)} + a_{h_2(x)} + a_{h_3(x)}) \bmod 2^r.$$

This is always possible, because by examining the vertices in the order produced by the acyclicity check we can always choose the value for the hinge (if there are more unassigned values, we set them to zero).

Storing the function in this way would require γrn bits. We call such a structure an *MWHC function* (from Majewski, Wormald, Havas and Czech). We note, however, that when r is large we can use an additional bit array s to mark those vertices that have a non-zero value, and record in an array b only the (at most n) nonzero values. To compute a_i , we first look at s_i : if it is zero, $v_i = 0$; otherwise, we compute $\text{rank}_s(i)$ and use the resulting value to index the array b .

The resulting structure, which we call a *compacted MWHC function*, uses $\gamma n + rn$ bits: this is advantageous as long as $\gamma + r < \gamma r$, which happens when $r > 5$.³

Three remarks are in order:

- even the best imaginable solution obtained by coupling a minimal perfect hash function (requiring at least $n \log e \approx 1.44n$ bits [13]) and an array of rn bits is never advantageous;
- for an order-preserving hash function, $\log(n!) = n \log n - O(n)$ bits is an obvious lower bound (as we can store the keys following any order), so a compacted MWHC function provides an optimal solution: thus, we will not discuss order-preserving functions further.

Another, complementary approach to the storage of static functions uses a minimal perfect hash function to index a *compressed* bit array (see the next section for some examples of suitable techniques). To obtain a minimal perfect hash function, we can adopt again the above hypergraph technique and use two bits per vertex to code the *index of the hash function outputting the hinge*. This effectively provides a perfect hash function $S \rightarrow \gamma n$ into the vertex space (by mapping each key to its hinge, a value in $\{1, 2, 3\}$). Thus, the perfect hash function can be turned into a *minimal* perfect hash function by ranking, as it is immediate to devise a space

$o(n)$, constant-time ranking structure that counts nonzero pairs of bits, so we obtain minimal perfect hashing in $2\gamma n$ bits (the idea of coding the hinge position appeared for the first time in Bloomier filters [7]; using ranking to obtain a minimal perfect hash function was suggested in [3]). This approach is advantageous if the bit array can be compressed with a ratio better than $1 - \gamma/r$.

Two-step MWHC functions. To gain a few more bits when the distribution of output values is significantly skewed, we propose *two-step MWHC functions*. We fix an $s < r$: then, the $2^s - 1$ most frequent output values are stored in a (possibly compacted) s -bit MWHC function, whereas the value $2^s - 1$ is used as an escape to access a secondary MWHC function storing all remaining values. Of course, we need to store explicitly the mapping from $2^s - 1$ to the set of most frequent values. The value s is chosen so to optimise the space usage.

A large-scale approach. MWHC functions require a large amount of memory to be built, as they require random access to the 3-hypergraph to perform a visit. To make their construction suitable for large-size key sets we reuse some techniques from [4]: we divide keys into *chunks* using a hash function, and build a separate MWHC function for each chunk. We must now store for each chunk the offset in the array a where the data relative to the chunk is written, but using a chunk size $\omega(\log n)$ (say, $\log n \log \log n$) the space is negligible. The careful analysis in [4] shows that this approach can be made to work even at a theoretical level by carefully reusing the random bits when building the MWHC functions of each chunk.

3.3 Elias–Fano representation of monotone functions.

We will frequently need to store either a list of arbitrary integers, or a list of nondecreasing natural numbers. In both cases, we would like to consume as little space as possible. To this purpose, we will use the *Elias–Fano representation of monotone functions* [10, 11]. Such a data structure stores a monotone function $f : n \rightarrow 2^s$, that is, a list of nondecreasing natural numbers $0 \leq x_0 \leq x_1 \leq \dots \leq x_{n-1} < 2^s$, provides constant-time access⁴ to each x_i , and uses $2 + s - \log n$ bits per element when $n < 2^s$, and $1 + 2^s/n$ bits when $2^s \leq n$.

Here, for completeness we briefly recall from [11] the representation: we store explicitly the $s - \log n$ lower bits of each x_i in a bit array. The value v of the upper $\log n$ bits of x_i are written in a bit array b of $2n$ bits by setting the bit of position $i + v$. It is easy to see that now v can be recovered as $\text{select}_b(i) - i$. Since the lower bits use

³tions are fully random. Albeit controversial, this is a common practical assumption that makes it possible to use the results about random hypergraphs.

³This evaluation does not take into account that ranking structures are asymptotically $o(n)$, but on real data they occupy a significant fraction of the original data size. The actual threshold depends on that fraction.

⁴Actually, in the original Elias paper access is not constant, as it relies on a selection structure that is not constant-time. On one side, replacing the selection structure with a modern, constant-time structure provides constant-time access to the x_i s; on the other side, practical implementations use the original Elias inventory-based linear scan, as it turns out to be faster.

$s - \log n$ bits per element, and the bit array requires $2n$ bits, we obtain the stated space bounds (in the case $2^s \leq n$, the first array is empty, and the second array requires just $n + 2^s$ bits). Selection in b can be performed with one of the many selection structures in the literature (see, e.g., [8, 16, 17, 24, 30]; we note that if $n \ll 2^s$, the impact of the selection structure on the space usage is irrelevant).

Finally, if we have to store a list of natural numbers x_0, x_1, \dots, x_{n-1} , we can just juxtapose the binary representation of $x_0 + 1, x_1 + 1, \dots, x_{n-1} + 1$ (without the most significant bit) and use an Elias–Fano monotone sequence to store the starting point of each binary representation. The resulting structure provides significant gains if the distribution is skewed towards small values, as the overall space usage is $2n + n \log(\sum_i \lfloor \log(x_i + 1) \rfloor / n) + \sum_i \lfloor \log(x_i + 1) \rfloor$. We will speak in this case of an *Elias–Fano compressed list*.

3.4 Bucketing. We now discuss here briefly a general approach to minimal perfect monotone hashes that we will use in this paper and that will be referred to as *bucketing*. The same idea has been widely used for non-monotone perfect hashing, and its extension proves to be fruitful.

Suppose you want to build a minimal perfect monotone hash function for a set S ; you start with:

- a monotone hash function $f : S \rightarrow m$ mapping S to a space of m buckets, called the *distributor*;
- for each $i \in m$, a minimal perfect monotone hash function g_i on $f^{-1}(i)$;
- a function $\ell : m \rightarrow n$ such that, for each $i \in m$,

$$\ell(i) = \sum_{j < i} |f^{-1}(j)|.$$

Then, the function $h : S \rightarrow n$ defined by

$$h(x) = \ell(f(x)) + g_{f(x)}(x)$$

is a minimal perfect monotone hash function for S . The idea behind bucketing is that the distributor will consume little space (as we do not require minimality or perfection), and that the functions hashing each element in its bucket will consume little space if the bucket size is small enough.

4 Bucketing with longest common prefixes

The first solution we propose is taken from [1] and is based on *longest common prefixes*. This solution has the advantage of requiring just the evaluation of a fixed number of hash functions; on the other side, it has in practise the highest memory usage among the algorithms we discuss.

Let b be a positive integer, and divide the set S into buckets B_i of size b preserving order. We now apply bucketing as described in Section 3.4, but to store the function $f : S \rightarrow m$ (with $m = \lceil n/b \rceil$, as before), we proceed as follows:

s_0	2	00	0
s_1	2	00100110	1
s_2	2	00100110101	2
s_3	8	0	3
s_4	8		
s_5	8		
s_6	11		
s_7	11		
s_8	11		
s_9	1		
s_{10}	1		

Figure 2: Bucketing with least common prefix for the set S of Figure 1: (a) f_0 maps each element x of S to the length of the least common prefix of the bucket to which x belongs; (b) f_1 maps each least common prefix to the bucket index.

- we store a compacted MWHC function $f_0 : S \rightarrow \log(u/b)$ that assigns, to each $x \in S$, the length $\ell_{f(x)}$ of the longest common prefix of the bucket containing x (note that the length of the common prefix of a set of b strings of length $\log u$ cannot exceed $\log u - \log b$);
- we store a compacted (or a two-step) MWHC function $f_1 : \{p_0, p_1, \dots, p_{m-1}\} \rightarrow m$ mapping p_i to i .

To compute $f(x)$ for a given $x \in S$, one first applies f_0 obtaining the length $\ell_{f(x)}$; from this one can compute $p_{f(x)}$, whence, using f_1 , one obtains $f(x)$. Figure 2 displays the functions f_0 and f_1 for the example of Figure 1.

The function f_0 requires $(\gamma + \log \log(u/b))n$ bits, whereas f_1 requires $(\gamma + \log(n/b))n/b$ bits; the g_i s would require $(\gamma + \log b)n$ bits, as before, but we can pack the information returned by f_0 and the g_i s into a single function. Altogether, we obtain:

$$\left(\gamma + \log \log \frac{u}{b} + \log b\right)n + \left(\gamma + \log \frac{n}{b}\right)\frac{n}{b}.$$

Approximating $\log \log(u/b)$ with $\log \log u$, the above function is minimised by $b = W(\xi n)$ where W is Lambert's W function and $\xi = e^{2\gamma} \approx 6.4$, so

$$b \approx \ln(\xi n) - \ln \ln(\xi n) \approx 1 + \gamma \ln 2 + \ln n - \ln \ln n,$$

giving about

$$\left(\gamma + \log e + \log \log n + \log \log \frac{u}{\log n}\right)n$$

bits. A good upper bound on the space usage is $(2.7 + 2 \log \log u)n$ bits.

4.1 Variable-size bucketing. In case n is close to u , we can still use longest prefixes, but we must modify our

structure so as to use $O(n \log \log(u/n))$ bits. To do so, we divide the elements of S in buckets on the basis of their first $\log n$ bits. Clearly, the size of each bucket will now be variable, but the position of the first element of each bucket is a monotone function that can be recorded in $2n$ bits using the Elias–Fano representation. This construction provides the monotone hashing part of the bucketing scheme.

We are now left to define the minimal perfect monotone hashing inside each bucket. Let $b(x) = 1 + \gamma \ln 2 + \ln x - \ln \ln x$ be the bucket size used in the previous section. If a bucket has size smaller than $b(u/n)$, we simply use a MWHC function (which requires at most $\gamma + \log b(u/n)$ bits per element). Otherwise, we exploit the fact that each bucket can be seen as a subset of a smaller universe of size u/n by stripping the first $\log n$ bits of each string. Thus, the bound given in the previous section, when applied to a bucket of size s , becomes

$$\left(2\gamma + \log e + \log \log s + \log \log \frac{u}{n \log s}\right) s,$$

where the additional γs bits come from the need of separating the storage for the f_0 and the g_i s: in this way, both the $(\log b(u/n))$ -sized data for large and small buckets can be stored in the same bit array with no additional costs.

We conclude that the entire data structure requires at most

$$\left(2\gamma + \log e + 2 + \min \left\{ \log \log n, \log \log \frac{u}{n} \right\} + \log \log \frac{u}{n}\right) n$$

bits; evaluation is obviously still constant time if $\log u$ is smaller than a constant number of machine words. Note that the latter formula only becomes advantageous when n is very large ($n \approx \sqrt{u}$ or larger).

5 Bucketing with partial compacted tries

We now turn to a data structure requiring access time linear in the length of the key (e.g., $O(\log u)$). As in the previous section, the idea is always that of dividing S into equally sized buckets, and then compute the bucket offset using a compacted MWHC function. However, this time for each bucket B_i we consider its *delimiter* k_i , which is the last string in B_i (with respect to the natural \leq -order). Let D denote the set of delimiters of all buckets, except for the last one: we will locate the bucket of a key $x \in S$ by counting the number of elements of D that are smaller than x . This is actually a standard operation, called *rank*:

$$\text{rank}_T(x) = |\{t \in T \mid t < x\}| \text{ for } T \subseteq u \text{ and } x \in u.$$

There are many data structures in the literature that can be used to compute ranks. An obvious, naive possibility is using a compacted trie [25] containing the strings in D (see Figure 3). A much more sophisticated approach is described




Figure 3: The standard compacted trie built from the set D of Figure 1. This data structure can be used to rank arbitrary elements of the universe with respect to D : when the trie is visited with an element not in D , the visit may terminate at some arbitrary node, determining that the given element is to the left (i.e., smaller than) or to the right (i.e., larger than) all the leaves that descend from that node. The picture shows, for each element of S , the node where the visit would end.

in [18], whose *fully indexable dictionaries* make it possible to store the distributor f in $n/b \log(ub/n)$ bits (plus lower-order terms); each g_i requires $(\gamma + \log b)n/b$ bits. So we need altogether

$$\frac{n}{b} \log \frac{ub}{n} + (\gamma + \log b)n \text{ bits.}$$

This quantity is minimised letting $b = -W(-ne/u)$, so when $n \ll u$

$$b \approx \ln \frac{u}{ne} + \ln \ln \frac{u}{ne} \approx \ln \frac{u}{n} + \ln \ln \frac{u}{n} + 1$$

and we need

$$\left(\gamma + \log e + 2 \log \log \frac{u}{n}\right) n \text{ bits.}$$

The time required to access the data structure is now dominated by the computation of the rank (see [18]; the cost should be actually multiplied by $\log u$ as we do not assume that the $\log u$ is equal to the word size).

We propose, however, an interesting alternative based on the observation that both tries and the above dictionaries are able to rank *any element of u* on D . This is actually not necessary, as we need just ranking the elements of S .

5.1 Partial compacted tries. When using a compacted trie for the computation of rank, one has to compare at each step the sequence contained in the currently visited node (say p) with the same-length prefix x of the element that is sought. If $x < p$ or $x > p$, the visit ends at this point (the sequence was, respectively, smaller than or larger than the elements of D corresponding to the leaves that descend from the current node). If $x = p$, we must continue our visit on the left or right subtree, depending on the $|p|$ -th bit of the

string (this is what happens, in particular, for the elements of D).

In this section we introduce a new data structure, the *partial compacted trie*, that reduces the space usage of a compacted trie if you know in advance that you want to rank only strings out of a certain set $S \supseteq D$ (as opposed to ranking all strings in u).

To understand the idea behind this, suppose that you build the standard compacted trie of D , and that the root is labelled by p : when the trie is used to rank an element of S that has p as prefix, the visit continues on the left or right subtrie; this happens, in particular, for all elements of D (because p is the least common prefix of D) and, more generally, for all the elements of S that are between $\min D$ and $\max D$ (because they also have prefix p). On the other hand, some of the remaining elements of S (those that do not start with p and are smaller than $\min D$ or larger than $\max D$) cause the visit to end at the root. Now, in many cases it is not necessary to look at the *whole* prefix of length $|p|$ to understand that the visit must end; for example, suppose that $p = 01010$ but all the elements smaller than $\min D$ start with 00 and all elements larger than $\max D$ start with 11 : then the first two bits of p are enough to determine if the visit must stop, or if we must proceed to some subtrie. We might store $01???$ to mean that if the string starts with something *smaller* (larger, resp.) than 01 , then it falls on the left (right, resp.) of the whole D , otherwise, we can ignore the following three bits (whose values are anyway fixed for the remaining elements of S), and proceed in the usual way looking at the sixth bit.

This intuitive idea is formalised as follows: a *partial compacted trie* (PaCo trie, for short) is a binary tree in which every node contains not a binary string but rather a pattern formed by a binary string followed by zero or more “don’t know” symbols (?), for instance, “00101???”. Given a PaCo trie, and an $x \in u$, the visit of the trie with x starts from the root and works as follows:

- if the node we are visiting is labelled by the pattern $w^?^k$, we compare the first $|w|$ symbols of x (possibly right-padded with zeroes), say x' , with w ;
- if x' is smaller than w , or if $x' = w$ and the current node is a leaf, we end the visit and return the number of leaves to the left of the current node;
- if x' is larger than w , we end the visit and return the number of leaves to the left of the current node *plus* the number of leaves in the subtrie rooted at the current node;
- if $x' = w$ and the current node is not a leaf, let b be the $(|w| + k)$ -th bit of x , and y be the suffix following it: we recursively visit the left or right subtrie (depending on whether $b = 0$ or $b = 1$) with y .

Differently from a standard compacted trie, the construction of a PaCo trie depends both on S and D ; it is similar to the construction of a standard trie, but instead of taking the whole least common prefix p of the elements of D , we just take the smallest part that is enough to disambiguate it from all the elements of S that do not start with p (and that are, necessarily, all smaller than $\min D$ or larger than $\max D$). In particular, it is enough to find the largest element smaller than $\min D$ that does not start with p (say s') and the smallest element larger than $\max D$ that does not start with p (say s''); these two elements alone are sufficient to determine how short we can take the prefix of p : if the prefix we take is enough to understand that s' falls on the left, then the same will *a fortiori* happen for smaller elements of S , and similarly for s'' .

Formally, we define the PaCo trie associated to a set $S = \{s_0 < \dots < s_{n-1}\} \subseteq u$ with respect to $D = \{s_{i_0} < \dots < s_{i_{k-1}}\}$ as follows (the reader may find it easier to understand this construction by looking at Figure 4):

- let p be the least common prefix of D ;
- let j' be the largest integer in $\{0, \dots, i_1 - 1\}$ such that p is not prefix of $s_{j'}$, and ℓ' be the length of the longest common prefix between p and $s_{j'}$; in the case all strings in s_0, \dots, s_{i_1-1} start with p , we let $\ell' = |p| - 1$ and $j' = 0$;
- let j'' be the smallest integer in $\{i_k + 1, \dots, n - 1\}$ such that p is not prefix of $s_{j''}$, and ℓ'' be the length of the longest common prefix between p and $s_{j''}$; in the case all strings in $s_{i_k+1}, \dots, s_{n-1}$ start with p , we let $\ell'' = |p| - 1$ and $j'' = n$;
- let now $\ell = 1 + \max(\ell', \ell'')$, and q be the prefix of length ℓ of p ;
- if $k = 1$, the PaCo trie of S w.r.t. D is the one-node PaCo trie labelled with $q^{|p|-\ell}$;
- if $k > 1$, let j be the smallest index such that s_j starts with $p1$, and t be the smallest index such that $j \leq i_t$: the PaCo trie of S w.r.t. D has root labelled with $q^{|p|-\ell}$, and the left and right subtrees are defined as follows⁵:
 - the left subtrie is the PaCo trie of $\{s_{j'+1} - p, \dots, s_{j-1} - p\}$ with respect to $\{s_{i_0} - p, \dots, s_{i_{t-1}} - p\}$;
 - the right subtrie is the PaCo trie of $\{s_j - p, \dots, s_{j''-1} - p\}$ with respect to $\{s_{i_t} - p, \dots, s_{i_{k-1}} - p\}$.

⁵Here, we are writing $s - p$ for the string obtained omitting the prefix p from s .




Figure 4: Constructing the first level of the PaCo trie for S with respect to D (see Figure 1): the central box corresponds to the longest common prefix $p = 001001$; here $\ell' = 2$ and $\ell'' = 1$, hence $\ell = 3$ giving rise to $q = 001$.

The PaCo trie for the example of Figure 1 is shown in Figure 5.

This definition is justified by the following theorem:

THEOREM 5.1. *Let $D \subseteq S \subseteq u$ and T be the PaCo trie of S with respect to D . Then, for every $x \in S$, the visit of T with x returns $\{y \in D \mid y < x\}$.*

Proof. We will use the same notation as above, and show that the visit of a PaCo trie with an $x \in S$ follows the same route as a visit to the full compacted trie built over D . Let us suppose that $x = s_m$. If $j' < m < j''$, then p is a prefix of x , hence also q is a prefix of x , and we proceed with the visit going down to the left or right subtree, as we would do with the full compacted trie. If $m = j'$ (the case $m = j''$ is analogous) is treated as follows: q is strictly longer than the longest common prefix between p and $s_{j'}$ (because $|q| > \max\{\ell', \ell''\} \geq \ell'$); so the prefix of x of length $|q|$, say q' , is different from q , and it must necessarily be lexicographically smaller than it. We therefore end the visit, as we would do in the full trie (x is smaller than all keys). If $m < j'$ (the case $m > j''$ is analogous) follows *a fortiori*, because the prefix of x of length $|q|$ will be even smaller than (or equal to) q' .

In our general framework, we will use a PaCo trie as a distributor; more precisely, given $S = \{s_0, \dots, s_{n-1}\}$ consider, for some positive integer b , the set $D = \{s_0, s_b, s_{2b}, \dots\}$: the PaCo trie of S with respect to D is a distributor for the function $f : S \rightarrow \lceil n/b \rceil$ that maps s_m to $\lceil m/b \rceil$. In this application, all buckets have size b (except possibly for the last one, which may be smaller).

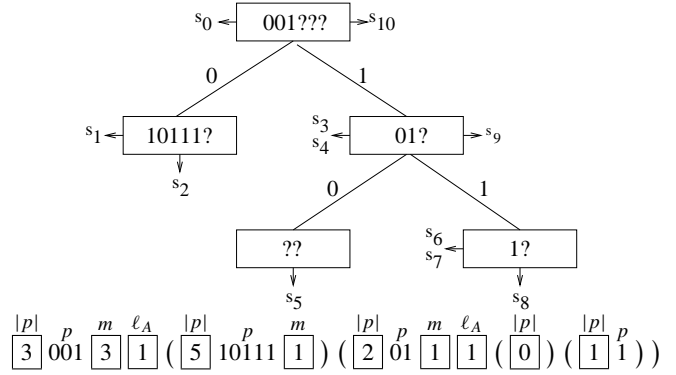


Figure 5: The PaCo trie for S with respect to D (see Figure 1): like for the trie in Figure 3, we can use this data structure to rank with respect to D , but *only* for elements of S , and not for arbitrary elements of the universe. At the bottom, we show the recursive bit stream representation: all framed numbers are written in δ coding. The skip component is omitted but we use parentheses to isolate the parts of the stream corresponding to each subtree.

5.2 Implementation issues. Experimental evidence suggests that PaCo tries usually save 30 – 50% of the space occupied for paths compared to a standard trie. To get most benefits from this saving, we propose to store PaCo tries in a *recursive bit stream format*. More precisely, the representation $\llbracket T \rrbracket$ of a trie T whose root contains a bit string p followed by m don't-care symbols, left subtree A and right subtree B is given by the concatenation

$$s \mid p \mid p \mid m \mid \ell_A \mid \llbracket A \rrbracket \mid \llbracket B \rrbracket,$$

where s is the length in bits of $\llbracket A \rrbracket$, ℓ_A is the number of leaves of A , and all numbers are represented in δ coding. Leaves have $s = 0$, and they do not record the information that follows p . Figure 5 shows the encoding of our example trie.

This representation (which is trivially decodable) makes it possible to navigate the PaCo trie in at most $\log u$ steps, as the left subtree of a node is available as soon as the node data is decoded, and it is possible to jump to the right subtree by skipping s bits. Moreover, at each navigation step we can compute in constant time the number of leaves at the left or under the current node using ℓ_A .⁶

By sizing buckets appropriately, the space usage of this representation is linear in n , and indeed we will see that using a PaCo trie as a distributor provides in practise the best space/time tradeoff for long keys.

We observe, however, that due to the difficulty of estimating the bit gain of a PaCo trie w.r.t. a dictionary or a

⁶We remark that in principle ℓ_A can be reconstructed by visiting the trie. However, adding it to the structure makes it possible to return the correct value immediately after locating the exit node.

standard trie, it is very likely that the computation of the optimal bucket size is far off the best value. Since trying exhaustively all bucket sizes is out of the question, we propose the following heuristic approach: we first estimate the bucket size using the formula above. Then, we compute the PaCo trie and assume that halving the bucket size (thus doubling the number of delimiters) will also approximately double the size of the PaCo trie. Using this simple model, we compute a new estimation of the bucket size and build the corresponding structure. In practise, this approach provides a very good approximation.

The algorithm given in the previous section requires accessing the strings of S and D randomly. However, there is a two-pass sequential algorithm that keeps in internal memory a representation of the trie built on D , only.

6 Succinct hollow tries

The *hollow trie* associated to S is a compacted trie [25] in which all paths of internal nodes have been replaced with their length, and all paths on leaves have been discarded. More in detail, given S , we can define the hollow trie inductively as follows:

- if $|S| \leq 1$, the associated hollow trie is the empty binary tree;
- otherwise, if p is the longest common prefix of strings in S , the associated hollow trie is a binary tree whose root is labelled by $|p|$, and whose left and right subtrees are the hollow tries associated to the sets $\{x \in \{0, 1\}^* \mid p \cdot x \in S\}$ for $i = 0, 1$, respectively.

The hollow trie for the example of Figure 1 is shown in Figure 6.

Note that a hollow trie is very similar to the *blind trie* that underlies a Patricia trie [27]: however, in a blind trie we keep track of the lengths of all compacted paths, whereas in a hollow trie the lengths on the leaves are discarded. Indeed, the blind trie of a single string is given by a node containing the first character of the string and the number of remaining characters, whereas the hollow trie of a single string is the empty binary tree.

We store hollow tries by computing the corresponding forest (in the first-child/next-sibling representation), adding a node at the top, and using Jacobson's representation [22] to store the string of balanced parentheses associated to the tree. Thus, we use two bits per node plus $\log \log u$ bits for each label of an internal node (actually, in practise using a variable-length bit array for the labels provides significant savings). We recall that Jacobson's representation divides parentheses in blocks of $\log n$ elements and stores explicitly the closing parentheses corresponding to a subset of open parentheses called *pioneers*. In our case, we use 64-bit blocks; we store the list of positions of open pioneers using




Figure 6: The hollow trie for the set S of Figure 1, and the associated forest (in this case, a tree); at a node labelled by i , look at the i -th bit (numbered from 0), follow the corresponding arc, and discard all bits up to the i -th (inclusive). At the bottom, we show the corresponding representation by balanced parentheses; the bold pair corresponds to the fictitious (round) node.

an Elias–Fano monotone sequence, and the distance of the corresponding closed parenthesis using an Elias–Fano compressed list. In practise, this brings down the space needed to less than one bit per node. Much more sophisticated schemes (e.g., [15, 28]) achieve in theory $o(n)$ bits usage, but they require either to store both open and closed pioneers (we need just open pioneers—see below) or to use the standard trick of using several levels of blocking, which in our case leads to a very small space gain and a significant slowdown. Figure 6 shows our example trie, the corresponding forest, and the associated balanced parentheses.

All in all, just $(2 + \log \log u) n$ bits are sufficient, plus the bits required to store the structure for balanced parentheses. Traversing such a trie on the succinct representation is an easy process: if we are examining the open parenthesis at position p , moving on the left is just moving on the open parenthesis at position $p + 1$ (if the parenthesis at $p + 1$ is closed, the left binary subtree is empty), whereas moving on the right is just matching the current parenthesis, getting position q , and moving to position $q + 1$ (if the parenthesis at $q + 1$ is closed, the right binary subtree is empty). For every element $x \in S$ the resulting leaf will be exactly the leaf associated to x ; of course, since we have discarded all paths, the result for strings not in S will be unpredictable.

We note that while walking down the trie we can easily compute the number of strings in S that are lexicographically smaller than w and the number of open parentheses the

precede the current one. When we move to the left subtree, the first number does not change and the second one is incremented by one; when we move to the right subtree, we must add to both numbers the number of leaves in the left subtree: this number is exactly $(q - p + 1)/2$ —the number of open parentheses between p and q , using the same notation of the previous paragraph. This observation makes it possible to get the skip associated to each node, and also to return the correct value as soon as the right leaf has been found.

Note that the process above requires time $O(\log u)$, albeit for not-so-pathological sets it will be more like $O(\log n)$. However, the constant factors associated to the navigation of the succinct representation are very high. In theory and in practise, a hollow trie occupies less space than the previous two structures, but we will see that its main usefulness lies in the possibility of using it as the base of a distributor.

6.1 Variable-size bucketing. Similarly to the case of longest common prefixes, for dense sets we can get to a space usage $O(\log \log(u/n))$ by dividing the elements of S in buckets on the basis of their first $\log n$ bits. We can record the start of each bucket using a $2n$ -bits Elias–Fano representation, and concatenate the bit representations of the hollow tries of each bucket. Since we store each trie in $2b$ bits we can recover the trie of a bucket using only the index of the first element. This gives a space usage of

$$\left(4 + \log \log \frac{u}{n}\right) n \text{ bits.}$$

It seems that this modification is not going to be practical unless n is very close to u ($n \geq u^{4/5}$).

6.2 Implementation issues. The speed of rank/select operation is of course crucial in implementing a hollow trie. We use the currently fastest practical implementations available on 64-bit processors [30].

Another important issue is the storage of the labels. As in the case of a compacted trie, it is difficult to estimate the actual size of the labels of a hollow trie, so we store the labels in an Elias–Fano list. An interesting path for future research is to explore alternative, more efficient and cache-effective succinct representations for hollow tries (e.g., using a bit stream format similar to that used for PaCo tries).

7 Bucketing with hollow tries

We finally turn to the slowest data structure, which make it possible to use just $O(\log \log \log u)$ bits per element: in practise, unless the string length exceed a billion, any data set can be monotonically hashed using less than one byte per element in time $O(\log u)$.

The basic idea is that of using a succinct hollow trie

built on the delimiters as a distributor. Since skips require $\log \log u$ bits each, we can use the same size for our buckets; then, the occupation of the trie will be linear in the number of keys, and the bucket offset will require just $\log \log \log u$ bits.

It is however very easy to check that a hollow trie is not sufficient to map correctly each key into its bucket: in the example of Figure 1, if we built a hollow trie on the delimiters, the string s_{10} would be misclassified (more precisely, mapped to the first bucket), as the first test would be on the seventh bit, which is zero in s_{10} .

The idea we use to solve this problem is to mimic the behaviour of a trie-based distributor: all we need to know is, for each node and for each key, which behaviour (exit on the left, exit on the right, follow the trie) must be followed. This information, however, can be coded in very little space. To see how this can be done, consider what happens when visiting a trie-based distributor using a key x .

At each node, we know that x starts with a prefix that depends only on the node, and that there is a substring s of x , following the prefix, that must be compared with the compacted path stored at the node. The behaviour of x at the node depends *only* on the comparison between s and the compacted path. The problem with using a hollow trie as distributor is that we know just the *length* of the compacted path, and thus also s , but we do not know the compacted path.

The solution is to store explicitly the function (depending on the node) that maps s to its appropriate behaviour (exit on the left, exit on the right, follow the trie). The fundamental observation is that the number of overall keys of such maps (upon which the space usage depends) is very low, as we have one exit substring per key (n keys) and one follow substring per node (n/b keys, where b is the bucket size), as you follow the trie *only* when s coincides with the compacted path associated to the node. Figure 7 shows the distributor obtained for our example.

This, we have to map $n + n/b$ strings to one of three values. We note, however, that there is no “follow” behaviour on leaves (from a distributor viewpoint, “follow” on a leaf coincides with “exit on the left”). All in all, the space requirement is $n/b(2 + \log \log u)$ for the trie, $2\gamma n$ bits to store the behaviour on internal nodes, and $\gamma n/b$ bits to store the behaviour on the leaves using MWHC functions⁷; finally, we need $2n\gamma \log \log u$ bits to store the offset of each key into

⁷Actually, these are upper bounds, as it might happen that some strings exit on a leaf.




Figure 7: The distributor based on a hollow trie built from the set D of Figure 1. This data structure can be used to rank elements of S with respect to D (the reader is invited to compare with Figure 3). The picture shows, for each node of the hollow trie, the associated function.

his bucket⁸, resulting in

$$\frac{n}{b}(2 + \log \log u) + 2\gamma n + \gamma \frac{n}{b} + \gamma n \log b$$

bits. This is minimised when

$$b = \frac{\ln 2}{\gamma}(\log \log u + 2) + \ln 2,$$

resulting in

$$\left(\frac{1}{\ln 2} + 2 + \gamma \log \log \log u\right)n$$

overall bits, that is, approximately $3.44 + 1.23 \log \log \log u$ bits per element. The time required to hash a string is $O(\log u)$, as in the case of the hollow trie; at each node we have to compute a MWHC function, but since the overall length of all the strings involved in such computation is bounded by $\log u$, this has no asymptotic effect on the computation time.

7.1 Implementation issues. Storing n/b functions explicitly would increase significantly space usage. Rather, we store just two functions mapping pairs node/string to the associated action.

The distributor we discussed can be easily built using a three-pass linear procedure similar to that of the PaCo trie:

⁸We remark that we are actually wasting $2 - \log 3$ bits per element in coding the behaviour of internal nodes, but this has a very minor influence on the final result. Moreover, we are using non-compacted MWHC functions as in this case they are advantageous unless u is preposterously large.

in the first phase we build a compacted trie on the delimiters. Then, we compute for each key the action at each node, and store it in a temporary file, taking care of never saving twice the “follow” action at an internal node. Then, we read the temporary file and build the necessary MWHC functions. Finally, we compute the MWHC function mapping each key to its offset.

8 Bucketing with relative z-fast tries

In [1] we have proposed a data structure for the monotone minimal perfect hashing problem which has space usage $O(\log \log \log u)$ and access time $O((\log u)/w + \log \log u)$, where w is the size of a machine word. The structure, which uses as distributor a *relative z-fast trie*, is very complex and we invite the interested reader to consult [1] for the details. We have implemented the structure using the techniques described in this paper: the formula giving the space usage, however, is wide and ugly. We simply report that the optimal bucket size is $b \approx 16 + 7 \ln u \ln u + \ln \ln \ln u$, which is the value used in our implementation.

9 Average length

In this section we remark the few modifications that are necessary to make the space usage of our structures proportional to $\log \ell$, where ℓ is the average length of a set of prefix-free, variable-length strings S . The main ingredient is again the Elias–Fano representation. We remark, however, that on the files used for our experiments (see Section 10) the (small) additional constant costs in bits per key due to the additional structures makes the final hashing structure actually more expensive than the original one in terms of bits per key. However, generally speaking a bound in terms of the average length makes the method robust against the situation where a few keys in S are much longer than the average.

In the case of bucketing based on longest common prefixes, we simply note that since the longest common prefix of a bucket is shorter than the shortest string in the bucket, the sum of lengths of all prefixes does not exceed $(n/b)\ell$, so we can replace the function storing the prefixes length with a minimal perfect hash function (2γ bits per prefix) followed by an Elias-Fano list $(\log((n/b)\ell)/(n/b)) = \log \ell$ bits per prefix). The analysis of Section 4 can be carried on using ℓ in place of $\log u$, obtaining the desired result.

In the case of bucketing based on PaCo tries, we choose as delimiters the *shortest string* of each bucket, and then again the overall number of bits in the trie paths cannot exceed $(n/b)\ell$ (in lieu of the obvious bound $(n/b) \log u$). In this case, we also have to store an additional bit per key, which tells whether the string falls at the right or at the left of its delimiter.

Finally, in a hollow trie the sum of labels cannot exceed the overall string length ℓn , so storing the $n - 1$ labels actually requires $2 + \log \ell$ bits per key. The result can be

obviously carried over to the structure using the hollow trie as a distributor, using the same idea of PaCo tries.

10 Experiments

In this section we discuss a set of experiments conducted using the three data structure we introduced, and, for comparison, an order-preserving hash function computed by an MWHC compacted function. We used Java for all implementations; the tests were run using the Sun JVM 1.6 on a 64-bit Opteron processor running at 2814 MHz with 1 MiB of first-level cache. We remark that all our implementations are 64-bit clean (e.g., the number of strings, overall length of the files, etc. are limited by 2^{64}) and none requires loading the entire set of data to be processed into internal memory.

For our tests, we used a number of files that we describe in detail:

- `trec-title.terms` (9.1 MiB, ≈ 1 million strings): the terms appearing in titles of the TREC GOV2 collection (UTF-8);
- `trec-text.terms` (419 MiB, ≈ 35 million strings): the terms appearing in the text of the TREC GOV2 collection (UTF-8);
- `webbase-2001.urls` (6.7 GiB, ≈ 118 million strings): the URLs of a general crawl performed by the WebBase crawler [2] in 2001 (ISO-8859-1);
- `uk-2007-05.urls` (12 GiB, ≈ 106 million strings): the URLs of a 101 Mpages crawl of .uk performed by UbiCrawler [20] in May 2007 (ISO-8859-1);
- `largerandom.bin` (859 MiB, 100 million strings): random 64-bit strings;
- `smallrandom.bin` (477 MiB, 100 million strings): random 32-bit strings.

The first two files represent typical text dictionaries. We provide two different URL lists because they are very different in nature (the most recent one has significantly longer URLs). Finally, the random string files are useful to test the behaviour of our structures in the presence of random strings. The latter, in particular, is extremely dense (n is close to u).

Table 1 reports the results of our experiments. Beside the plain encoding of each file, we also tried a more sophisticated approach based on Hu–Tucker codes [21]. Hu–Tucker codes are *optimal lexicographical codes*—they compress optimally a source reflecting, however, the order between the symbols of the source in the lexicographical order of the codewords (this entails a loss of space w.r.t. to entropy, which is however bounded by 2 bits). It is interesting to experiment with Hu–Tucker codes because the increase of compression

(or, better, a lack thereof) can be used as a measure of the effectiveness of our data structures (in case of binary numbers, Hu–Tucker codes were computed on bytes).

We accessed one million randomly selected strings from each set. The tests were repeated thirteen times, the first three results discarded (to let the Java Virtual Machine warm-up and optimise dynamically the code) and the remaining ones averaged.

- From the results of our experiments, PaCo-based monotone hash function has the best tradeoff between space and time, but in case a high speed is required, LCP-based monotone hash functions have a much better performance.
- PaCo-based monotone hashing comes close to the space of hollow tries on long string. While we cannot justify this fact theoretically, in practise this shows that PaCo tries exploit very well the redundancy in our data.
- There no significant advantage in using Hu–Tucker coding. The only relevant effect is on LCP-based monotone hashing, as the size of the hash function is very dependent on the string length, and Hu–Tucker coding does reduce the string length; however, the same gain can be obtained by a two-step LCP-based function.
- The smallest structures come from hollow tries—either used directly, or as a distributor, but accessing them is usually very expensive. The case of short strings, however, is different: here hollow tries provide a very significant compression gain, with a small loss in access time. This is due to the significantly smaller number of accesses.
- Structures based on relative z-fast tries, albeit the best in theory, are always beaten by PaCo-based structures. The gap, however, is small, and might suggest that some additional engineering might make relative z-fast tries the best also in practise.

We remark that in our current implementation we focused on compression ratio and access time: our construction procedures are not particularly optimised and in particular building PaCo tries is quite expensive.

11 Conclusions

We have presented experimental data about some old and new data structures that provide monotone minimal perfect hashing in a very low number of bits. Our results show that our data structures are practical and cover a wide range of space/time tradeoffs.

Improvements to function storage (e.g., using $rn + o(n)$ methods) would yield immediately improvements to the data

File	trec-title 139 b/key			trec-text 202 b/key			webbase-2001 481 b/key			uk-2007-05 898 b/key			64-bit random			32-bit random		
	b/key	μ s/key	c.t.	b/key	μ s/key	c.t.	b/key	μ s/key	c.t.	b/key	μ s/key	c.t.	b/key	μ s/key	c.t.	b/key	μ s/key	c.t.
Plain string encoding																		
MWHC	22.47	0.54	3.9	27.47	0.78	3.2	28.46	1.18	4.8	28.46	1.32	5.7	28.47	0.74	3.7	28.46	0.64	2.3
LCP	16.62	0.74	4.1	22.17	1.11	13.0	25.20	1.48	9.7	22.20	1.93	10.2	13.18	1.21	6.1	13.18	1.17	3.5
LCP (2-step)	12.87	0.96	6.8	14.47	1.46	15.6	17.76	2.08	12.3	18.83	2.66	11.8	10.70	1.18	8.0	10.67	1.16	5.3
PaCo	8.43	1.66	7.8	8.62	2.60	25.1	9.85	4.26	19.9	10.53	4.62	16.7	7.04	1.82	7.1	6.78	1.70	3.6
Hollow	6.74	5.16	.3	6.91	8.73	.1	7.21	15.44	.1	7.57	15.31	.1	4.41	7.07	.1	4.37	6.60	.1
HTDist	6.38	10.42	9.0	6.36	16.66	23.4	6.51	24.87	18.1	6.49	29.07	15.4	5.99	11.28	9.8	5.60	13.41	9.7
RelZFast	9.40	4.50	12.8	9.94	6.24	25.1	10.22	8.88	23.2	9.91	10.65	24.5	8.57	6.15	10.0	8.58	6.03	9.3
Hu–Tucker encoding																		
LCP	14.62	1.29	5.9	22.18	1.83	8.0	24.20	4.40	17.9	22.20	7.12	26.7	13.18	1.73	6.9	13.18	1.55	3.9
LCP (2-step)	12.55	1.94	8.1	13.82	2.53	10.3	17.59	7.78	20.9	18.57	13.09	29.1	10.70	2.25	8.9	10.67	1.82	6.1
PaCo	7.30	1.93	8.1	7.53	2.63	25.0	9.47	6.35	41.9	9.89	9.62	60.9	7.00	2.29	10.3	7.31	2.06	5.3
Hollow	5.09	4.74	.3	5.23	7.69	.1	6.70	16.13	.1	7.11	20.78	.1	4.41	7.46	.1	4.37	6.89	.1
HTDist	5.96	8.78	12.6	6.16	12.39	19.8	6.44	27.41	31.8	6.43	31.02	42.5	5.99	11.55	13.1	5.60	13.45	10.7
RelZFast	9.26	4.49	13.9	9.80	6.48	20.1	10.14	10.59	38.1	9.87	14.76	53.4	8.58	6.58	11.1	8.58	6.15	9.4

Table 1: Experimental results. For each of the test files, we show the number of bits per key. Then, for each combination of test file and hash function we show the number of bits per key, the number of microseconds per key of a successful probe (unsuccessful probes are actually faster), and the time spent constructing the data structure (“c.t.”) expressed again in microseconds per key. The first table uses the encoding of each file as specified in Section 10, whereas the second table shows structures built using optimal lexicographical Hu–Tucker encoding.

structures presented here. However, current methods appear to be too slow for practical implementations.

The speed of our structures is presently mainly constrained by memory latency and unaligned access. Future work will concentrate on reducing this burden without increasing space usage.

References

- [1] D. Belazzougui, P. Boldi, R. Pagh, and S. Vigna. Monotone minimal perfect hashing: Searching a sorted table with $O(1)$ accesses. In *Proceedings of the 20th Annual ACM-SIAM Symposium On Discrete Mathematics (SODA)*, New York, 2009. ACM Press.
- [2] P. Boldi, B. Codenotti, M. Santini, and S. Vigna. Ubcrawler: A scalable fully distributed web crawler. *Software: Practice & Experience*, 34(8):711–726, 2004.
- [3] F. C. Botelho, R. Pagh, and N. Ziviani. Simple and space-efficient minimal perfect hash functions. In *WADS*, pages 139–150, 2007.
- [4] F. C. Botelho and N. Ziviani. External perfect hashing for very large key sets. In *CIKM*, pages 653–662, 2007.
- [5] J. Cain and N. C. Wormald. Encores on cores. *Electronic Journal of Combinatorics*, 13(1), 2006.
- [6] D. Charles and K. Chellapilla. Bloomier filters: A second look. In *Proc. ESA 2008*, 2008.
- [7] B. Chazelle, J. Kilian, R. Rubinfeld, and A. Tal. The Bloomier filter: an efficient data structure for static support lookup tables. In J. I. Munro, editor, *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11–14, 2004*, pages 30–39. SIAM, 2004.
- [8] D. R. Clark and J. I. Munro. Efficient suffix trees on secondary storage (extended abstract). In *SODA*, pages 383–391, 1996.
- [9] M. Dietzfelbinger and R. Pagh. Succinct data structures for retrieval and approximate membership (extended abstract). In L. Aceto, I. Damgård, L. A. Goldberg, M. M. Halldórsson, A. Ingólfssdóttir, and I. Walukiewicz, editors, *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7–11, 2008, Proceedings, Part I: Track A: Algorithms, Automata, Complexity, and Games*, volume 5125 of *Lecture Notes in Computer Science*, pages 385–396. Springer, 2008.
- [10] P. Elias. Efficient storage and retrieval by content and address of static files. *J. Assoc. Comput. Mach.*, 21(2):246–260, 1974.
- [11] R. M. Fano. On the number of bits required to implement an associative memory. Memorandum 61, Computer Structures Group, Project MAC, MIT, Cambridge, Mass., n.d., 1971.
- [12] E. A. Fox, Q. F. Chen, A. M. Daoud, and L. S. Heath. Order-preserving minimal perfect hash functions and information retrieval. *ACM Trans. Inf. Sys.*, 9(3):281–308, 1991.
- [13] M. L. Fredman and J. Komlós. On the size of separating systems and families of perfect hash functions. *SIAM J. Algebraic Discrete Methods*, 5(1):61–68, 1984.
- [14] M. L. Fredman, J. Komlós, and E. Szemerédi. Storing a sparse table with $O(1)$ worst case access time. *J. Assoc.*

- Comput. Mach.*, 31(3):538–544, July 1984.
- [15] R. F. Geary, N. Rahman, R. Raman, and V. Raman. A simple optimal representation for balanced parentheses. *Theor. Comput. Sci.*, 368(3):231–246, 2006.
 - [16] A. Golynski. Optimal lower bounds for rank and select indexes. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part I*, volume 4051 of *Lecture Notes in Computer Science*, pages 370–381. Springer, 2006.
 - [17] R. González, S. Grabowski, V. Mäkinen, and G. Navarro. Practical implementation of rank and select queries. In *Poster Proceedings Volume of 4th Workshop on Efficient and Experimental Algorithms (WEA'05)*, pages 27–38. CTI Press and Ellinika Grammata, 2005.
 - [18] A. Gupta, W.-K. Hon, R. Shah, and J. S. Vitter. Compressed data structures: Dictionaries and data-aware measures. *Theoret. Comput. Sci.*, 387(3):313–331, 2007.
 - [19] T. Hagerup and T. Tholey. Efficient minimal perfect hashing in nearly minimal space. In *Proceedings of the 18th Symposium on Theoretical Aspects of Computer Science (STACS '01)*, volume 2010 of *Lecture Notes in Computer Science*, pages 317–326. Springer–Verlag, 2001.
 - [20] J. Hirai, S. Raghavan, H. Garcia-Molina, and A. Paepcke. WebBase: a repository of Web pages. *Computer Networks*, 33(1–6):277–293, 2000.
 - [21] T. C. Hu and A. C. Tucker. Optimal Computer Search Trees and Variable-Length Alphabetical Codes. *SIAM J. Applied Math.*, 21(4):514–532, 1971.
 - [22] G. Jacobson. Space-efficient static trees and graphs. In *In Proc 30th Annual Symposium on Foundations of Computer Science*, pages 549–554, 1989.
 - [23] B. Jenkins. Algorithm alley: Hash functions. *Dr. Dobbs's Journal of Software Tools*, 22(9):107–109, 115–116, Sept. 1997.
 - [24] D. K. Kim, J. C. Na, J. E. Kim, and K. Park. Efficient implementation of rank and select functions for succinct representation. In S. E. Nikolettseas, editor, *Proc. of the Experimental and Efficient Algorithms, 4th International Workshop*, volume 3503 of *Lecture Notes in Computer Science*, pages 315–327. Springer, 2005.
 - [25] D. E. Knuth. *Sorting and Searching*, volume 3 of *The Art of Computer Programming*. Addison-Wesley, second edition, 1997.
 - [26] B. S. Majewski, N. C. Wormald, G. Havas, and Z. J. Czech. A family of perfect hashing methods. *Comput. J.*, 39(6):547–554, 1996.
 - [27] D. R. Morrison. PATRICIA—practical algorithm to retrieve information coded in alphanumeric. *J. Assoc. Comput. Mach.*, 15(4):514–534, 1968.
 - [28] J. I. Munro and V. Raman. Succinct representation of balanced parentheses and static trees. *SIAM J. Comput.*, 31(3):762–776, 2001.
 - [29] E. Porat. An optimal bloom filter replacement based on matrix solving, 2008. arXiv:0804.1845v1.
 - [30] S. Vigna. Broadword implementation of rank/select queries. In *Proc. WEA 2008: 7th International Workshop on Experimental Algorithms*, *Lecture Notes in Computer Science*. Springer–Verlag, 2008.