RecSplit: Minimal Perfect Hashing
via Recursive Splitting

Emmanuel Esposito[*]   Thomas Mueller Graf[†]   Sebastiano Vigna[‡]

Abstract

A *minimal perfect hash function* bijectively maps a key set S out of a universe U into the first $|S|$ natural numbers. Minimal perfect hash functions are used, for example, to map irregularly-shaped keys, such as strings, in a compact space so that metadata can then be simply stored in an array. While it is known that just 1.44 bits per key are necessary to store a minimal perfect hash function, no published technique can go below 2 bits per key in practice. We propose a new technique for storing minimal perfect hash functions with expected linear construction time and expected constant lookup time that makes it possible to build for the first time, for example, structures which need 1.56 bits per key, that is, within 8.3% of the lower bound, in less than 2 ms per key. We show that instances of our construction are able to *simultaneously* beat the construction time, space usage and lookup time of the state-of-the-art data structure reaching 2 bits per key. Moreover, we provide parameter choices giving structures which are competitive with alternative, larger-size data structures in terms of space and lookup time. The construction of our data structures can be easily parallelized or mapped on distributed computational units (e.g., within the MapReduce framework), and structures larger than the available RAM can be directly built in mass storage.

1 Introduction

Minimal perfect hash functions (MPHFs) are static data structures storing a bijection of a given set S of keys, $|S| = n$, into the first n natural numbers. While such a bijection can easily be stored using hash tables, MPHFs are allowed to return *any* value if the queried key is not in the original set S; this relaxation enables to break the information-theoretical lower bound of storing the set S. Indeed, MPHFs constructions achieve $O(n)$ bits of space, regardless of the size of the keys. This property makes MPHFs powerful techniques when handling, for instance, large sets of strings, and they are important building blocks of space-efficient data structures.

The usual requirement on MPHFs is that construction can be completed in linear time, and that lookup takes constant time; often these requirements are relaxed to *expected* linear/constant time.

The space lower bound of a MPHF is about $\lg e \approx 1.44$ bits per key [FKS84], and independent from S. Some theoretical constructions reach the lower bound, but they work only for preposterously large n [HT01]. In practice, the CHD construction [BBD09], detailed below, is currently able to reach about 2 bits per key, but at such a space we show that it is significantly slower than our approach for both construction and lookup. Techniques based on random linear systems, such as derivatives of the venerable construction by Majewski, Wormald, Havas, and Czech [MWHC96], are reasonably fast in construction and quite fast in lookups, but they each only achieve about 2.24 bits per key [GOV16]. Finally, techniques based on fingerprints [MSSZ14] are extremely fast in construction and lookup, but only when the space occupancy is extremely large (e.g., above 4 bits per key).

In this paper we present a new technique for storing minimal perfect hash functions based on recursive splitting and brute-force searching, with expected linear construction time and constant expected lookup time. Our technique is able to break the 2 bits barrier, and at the same time can improve the construction time, space usage and lookup time of CHD.

Our approach to build an MPHF is conceptually simple, and on a high level similar to existing approaches: we partition a set into buckets, and then process each bucket independently. Unlike other approaches, in our case the entries of a bucket logically form a tree which defines an independent MPHF. We analyze how exactly those trees should "look like", for a brute-force search to be efficient, for lookup to be fast, and to use little space. We also show how to implicitly encode such trees, and how to employ succinct data structures that allow efficient lookup and storage even for very small MPHFs.

One of the main virtues of our approach is that it makes it possible to create practical MPHFs using less than 9% more space than the information-theoretical lower bound, compared to the state of the art, which uses 40% more space. But our approach is also very flexible and amenable, such that it is useful for a wide range of parameters: it not only yields the most space-saving practical MPHFs, but it is also competitive regarding construction time and lookup over a wide range of parameter values. Existing algorithms, on the other hand, are competitive only in a narrow area.

All the code used in this paper is available from the authors under the GNU Lesser General Public License, version 3 or later, as a part of the Sux project (`http://sux.di.unimi.it/`).

---

[*]Università degli Studi di Milano, Italy
[†]Independent researcher, Switzerland
[‡]Università degli Studi di Milano, Italy

2 Notation

We use von Neumann's definition and notation for the set of natural numbers, so $n = \{0, 1, \ldots, n-1\}$. We use $\gg$ and $\ll$ for left and right bit shifting. Following Knuth [Knu11], we use λx for the index of the highest (leftmost) bit set (λx is undefined when $x \leq 0$).

3 Background and related work

3.1 Early work

Sprugnoli [Spr77] defines perfect hash functions on small sets by testing exhaustively constants in simple expressions involving integer divisions and remainders; in some small-sized cases, he is able to find a MPHF. He suggests that for large key sets one can first hash the keys into a fixed number of buckets (called therein *segments*), and then operate independently. Jaeschke [Jae81] refines his approach to obtain minimal perfect hash functions on small sets by exhaustive search. Both exhaustive search and hashing into buckets are essential elements of our construction.

3.2 Random Linear Systems

In their seminal paper [MWHC96], Majewski, Wormald, Havas, and Czech (MWHC hereinafter) introduced the first compact construction for static functions using the connection between linear systems and hypergraphs. To store a function $f : S \to t$, they generate a random system with $n = |S|$ equations in m variables of the form

$$w_{h_0(x)} + w_{h_1(x)} + \cdots + w_{h_{k-1}(x)} = f(x) \mod t \quad x \in S.$$

Here $h_i : U \to m$ are k fully random hash functions, and the w_i's are variables assuming values in $\mathbf{Z}/t\mathbf{Z}$. Due to bounds on the acyclicity of random graphs, if the ratio between the number of variables and the number of equations is above a certain threshold γ_k, the system can be almost always triangulated in linear time by peeling the associated hypergraph, and thus solved; in other words, we have both a probabilistic guarantee that the system is solvable, and that the solution can be found in linear time.

The data structure is then a solution of the system (i.e., the values of the variables w_i): storing the solution makes it possible to compute $f(x)$ in constant time. The space usage is approximately $\gamma_k b$ bits per key. The constant γ_k depends on the degree k, and attains its minimum at $k = 3$ ($\gamma_3 \approx 1.23$).

Chazelle, Kilian, Rubinfeld and Tal [CKRT04], unaware of the MWHC construction, proposed it independently, but also exploited the fact that a peelable hypergraph is also orientable: since as a side effect of the peeling process each hyperedge can be assigned a unique vertex (or, equivalently, each equation can be assigned a unique variable), each key can be assigned injectively an integer in $\lceil \gamma_k n \rceil$. Then, we just need to modify the MWHC construction so that instead of storing $f(x)$, we store which of the k vertices of the hyperedge is the assigned

one, and this can be done in approximately $\gamma_k \lceil \log k \rceil$ bits per key. At retrieval time, the value we store makes it possible to recover the unique integer assigned to the key.

To make the perfect hash function minimal, that is, a function to n, rather than $\lceil \gamma_r n \rceil$, a *ranking* data structure can be added. A bit vector of size $\lceil \gamma_r n \rceil$ records which vertices have been assigned to a hyperedge, and then with additional $o(n)$ bits the number of ones preceding a given one (i.e., its rank) can be computed in constant time (see [Vig08] for practical implementations).

Again, the best k is 3, which yields theoretically a $2.46n + o(n)$ bits data structure [BPZ13], using 2 bits per variable: since the value 3 never appears in the solution, it can be used instead of zero to mark vertices associated with hyperedges. In this way, vertices assigned to hyperedges have always nonzero values and are the only ones with nonzero values because of the way the solution is computed by the peeling process. In the end, it is easy to adapt ranking structures so as to rank nonzero pairs of bits, rather than ones, eliminating the need for the additional bit vector (see [BPZ13] for details). Finally, Genuzio, Ottaviano and Vigna have shown that it is possible to use bounds on *solvability*, rather than *acyclicity*, for the same purpose, obtaining a $2.24n + o(n)$ bits data structure [GOV16].

3.3 CHD

Belazzougui, Botelho, and Dietzfelbinger [BBD09] introduced a completely different construction, called CHD (compressed hash-and-displace), which makes it possible, in theory, to reach the information-theoretical lower bound of ≈ 1.44 bits per key, at the price of increasing the expected construction time. Keys are first mapped into small buckets of expected size λ (e.g., $\lambda = 4$), and then buckets are mapped into the codomain without collisions, starting with the largest ones. To map the buckets into the codomain they examine for each bucket an enumeration of fully random independent hash functions $\varphi_i^k : U \to k$, $i = 0, 1, 2, \ldots$: once a function mapping the current bucket in the codomain without collisions is found, its index is stored in the data structure. While the theoretical analysis suggests that by increasing the bucket size it is possible to achieve space usage as close to the lower bound as desired, in practice it is unfeasible to go below 2 bits per key.

3.4 Fingerprinting

Recently, Müller *et al.* [MSSZ14] introduced a completely new technique for minimal perfect hashing. A series of bit arrays of decreasing size, called *levels*, is used to record information about collisions between keys. More precisely, all positions in the first level to which a single key is mapped by a random hash function are marked with a one. Then, one takes all keys which participated in collisions, and tries to map them into the second level, using the same strategy, and so on. As a result, one obtains a perfect hashing of the key set: to retrieve the output associated with a key, one maps the key in turn to each level until a one is found, and then the (overall) position of

the one yields a distinct number for each key. A constant-time ranking structure over the concatenation of the bit arrays is then used, as in the case of hypergraph-based techniques, to turn the perfect hash function into a minimal perfect hash function.

Fingerprint-based minimal perfect hash functions have the advantage of being very tunable for speed (both for construction and lookup) by suitably setting the length of each level, and thus the space usage: at 3 bits per elements, for example, the authors report that only 1.56 levels are accessed on average, resulting in a very low number of cache misses. However, the best result for space is 2.58 bits per key, which is not competitive, for example, with CHD or the techniques described in this paper. Moreover, at that size construction and lookup are very slow.

4   RecSplit

RecSplit is based on the idea that for very small sets it is possible to find a MPHF simply by brute-force searching for a *bijection* with suitable codomain. However, we extend the idea to brute-force searching of *splittings* which bring large sets to an amenable size.

To build a RecSplit instance for a set of n keys, we first distribute keys randomly into buckets of average size b using a random hash function $g : S \to \lceil n/b \rceil$. Then, each bucket is recursively *split* into smaller pieces until we reach a target *leaf size* ℓ on which it is feasible to directly search for an MPHF. As we will see, the parameters ℓ and b provide different space and time tradeoffs. We will build an MPHF for each bucket independently, opening the way to parallel or distributed construction.

We search for both splittings and bijections considering an enumeration of fully random independent hash functions $\varphi_i^k :$ $U \to k$, $i = 0, 1, 2, \ldots$, whose indices will be stored in the data structure.

Given a set of keys X (e.g., a whole bucket) of size m, a split is defined by a list of parts $k_0, k_1, \ldots, k_{s-1}$ such that $\sum_{i=0}^{s-1} k_i = m$. We compute the splitting index by searching for the first function φ_i^m such that

$$\left| \left(\varphi_i^m \right)^{-1} \left(\left[\sum_{i=0}^{t-1} k_i \, .. \, \sum_{i=0}^{t} k_i \right) \right) \cap X \right| = k_t, \quad 0 \leq t < s-1.$$

In other words, we map the elements of X in m, and find the first function such that the elements mapped to a value smaller than k_0 are exactly k_0, the elements mapped to values greater than or equal to k_0 but smaller than $k_0 + k_1$ are exactly k_1, and so on.

Then, we proceed recursively on the s parts until the current size of a part is less than or equal to ℓ. At that point, given a part X we search for the first function $\varphi_i^{|X|}$ that is bijective on X. Thus, we obtain a rooted *splitting tree* with a *splitting index* associated to each internal node and a *bijection index* associated with each leaf (see Figure 1). We will represent each such index using an optimal Golomb-Rice instantaneous code [Sal07], as such indices have a geometric distribution.

While the size of each part is in principle arbitrary, we will see that it is convenient to work with sizes that are multiples of ℓ, and to have all parts of the same size except possibly the last one. In this way, in every bucket *all leaves have exactly size ℓ*, except possibly the last one, and for each tree node *all subtrees have the same shape*, except possibly the last one.

A *splitting strategy* is made of a leaf size ℓ and a desired *splitting unit* u (which must be a multiple of ℓ) for each $m > \ell$. These pieces of information univocally define the shape of the splitting tree for each $m > \ell$: the fanout of a node associated with m keys will be $\lceil m/u \rceil$, with all parts of size u except possibly the last one, which may have size $m \bmod u$.

Given a splitting tree for a set X, we can compute a MPHF on X: given $x \in X$, we follow the tree from the root up to a leaf using the splitting indices. Every time we move to a child, we know exactly how many keys we are leaving to our left, as that number only depends on the splitting strategy. Thus, when we get to a leaf, we know that, say, c keys are mapped to the left of the leaf. We then apply the bijection associated with the leaf to x, and add c to the resulting value. Multiple buckets are handled by keeping track of the prefix sums of the number of keys in each bucket (i.e., of how many keys are assigned to previous buckets) and modifying the result accordingly.

5   Searching for splittings and bijections

Many different splitting strategies are possible. While a completely analytical analysis appears to be formidable, in this section we prove a number of results which will make it much easier to devise a good strategy.

A first observation is that bucket sizes are distributed as a binomially distributed random variable S with parameters $p = b/n$ and n. The moments of S can be written as [Kno08]

$$\mathbf{E}\left[S^d \right] = \sum_{i=0}^{d} \left\{ \begin{matrix} d \\ i \end{matrix} \right\} i! \binom{n}{i} \left(\frac{b}{n} \right)^i$$

$$= \sum_{i=0}^{d} \left\{ \begin{matrix} d \\ i \end{matrix} \right\} \frac{n!}{n^i (n-i)!} b^i \leq \sum_{i=0}^{d} \left\{ \begin{matrix} d \\ i \end{matrix} \right\} b^i.$$

In particular, every algorithm that is polynomially bounded (even just in expectation) for each bucket will be linear in expectation when run on all keys, and any lookup algorithm doing polynomial work on a bucket will be constant-time in expectation.

5.1   Searching for splittings

Assuming that our family of functions is fully random, the probability of finding a split for a set of size m in a left part of
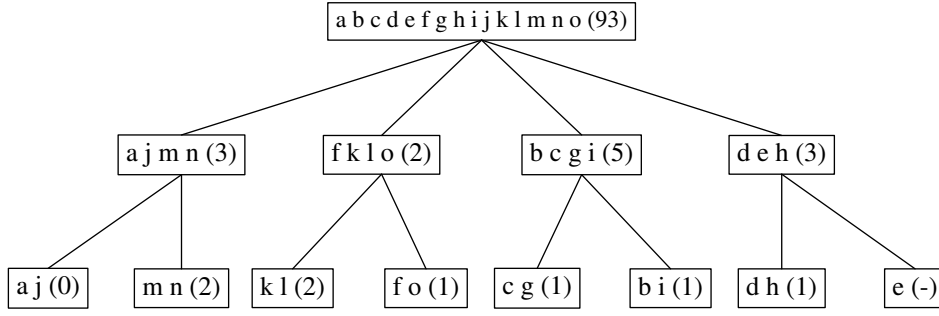
Figure 1: An example of a splitting tree for 15 element with $\ell = 2$. Each node is decorated with the associated keys and the splitting or bijection index. Note that the leaf of size one has no associated index.

size k and a right part of size $m - k$ is

$$\frac{1}{m^m}\binom{m}{k}k^k(m-k)^{m-k} \sim$$

$$\frac{1}{m^m}\frac{\left(\frac{m}{e}\right)^m\sqrt{2\pi m}}{\left(\frac{k}{e}\right)^k\sqrt{2\pi k}\left(\frac{m-k}{e}\right)^{m-k}\sqrt{2\pi(m-k)}}k^k(m-k)^{m-k}$$

$$= \sqrt{\frac{m}{2\pi k(m-k)}}.$$

using Stirling's approximation. Then, the average number of trials to find a splitting hash function will be asymptotic to

$$\sqrt{\frac{2\pi k(m-k)}{m}}$$

which is maximized when $k = \lceil m/2 \rceil$, that is, for balanced splits.

We remark that our choice of codomain m and threshold k is optimal. Indeed, if we consider the general case of a codomain r and a threshold t, the probability

$$(1) \qquad \frac{1}{r^m}\binom{m}{k}t^k(r-t)^{m-k}$$

depends only on the ratio r/t, as

$$\frac{1}{(\alpha r)^m}\binom{m}{k}(\alpha t)^k(\alpha r - \alpha t)^{m-k} = \frac{1}{r^m}\binom{m}{k}t^k(r-t)^{m-k}.$$

Now, if we explicitly write $t = \alpha r$ and minimize (1) in α, we obtain $\alpha = k/m$. Thus, any choice of r and t with $t/r = k/m$ provides the highest probability of success.

We will also use splittings with more than two parts. In general, if we want to split m into s parts as $m = k_0 + k_1 + \cdots + k_{s-1}$, the number of splitting functions is given by the number of possible ordered partitions of m into sets of cardinality $k_0, k_1, \ldots, k_{s-1}$ (i.e., the associated multinomial coefficient), multiplied by the number of possible functions within each set.

Dividing by the overall number of functions m^m we obtain

$$(2) \qquad \frac{1}{m^m}\binom{m}{k_0, k_1, \ldots, k_{s-1}}\prod_{i=0}^{s-1}(k_i)^{k_i}$$

$$= \frac{1}{m^m}\frac{m!}{\prod_{i=0}^{s-1}k_i!}\prod_{i=0}^{s-1}(k_i)^{k_i} = \frac{m!}{m^m}\prod_{i=0}^{s-1}\frac{(k_i)^{k_i}}{k_i!}$$

$$\sim \sqrt{\frac{m}{(2\pi)^{s-1}\prod_{i=0}^{s-1}k_i}},$$

where the last approximation is once again obtained using Stirling's. In the end, the average number of trials is asymptotically

$$(3) \qquad \sqrt{\frac{(2\pi)^{s-1}\prod_{i=0}^{s-1}k_i}{m}}.$$

By Jensen's inequality the worst case happens when all parts are equal, that is,

$$\sqrt{\frac{(2\pi)^{s-1}\left(\frac{m}{s}\right)^s}{m}},$$

which grows polynomially in m for fixed s, but with exponent s, which means that we cannot expect to be able to choose a large value of s.

We now want to give a bound of the expected amount of work done on a single bucket. To do so, we make a simplifying assumption, that is, that splitting always happens in s parts: we fix the leaf size ℓ and we use a bucket size of $s^h\ell$ (if the bucket has not a size of this form, we use the closest approximation from above). Let X_j^i be a set of random variables, where i represent the distance from the root, and $0 \le j \le s^i - 1$. The variables are independent and for each i they have a geometric distribution defined by (2); moreover, we upper bound the number of function evaluations using the number of elements

on which the splitting is computed:

$$\mathbf{E}\left[\sum_{i=0}^{h-1}\sum_{j=0}^{s^i-1} X_j^i s^{h-i}\ell\right] = s^h \ell \sum_{i=0}^{h-1} \mathbf{E}\left[X_0^i\right]$$

$$\sim s^h \ell \sum_{i=0}^{h-1} \sqrt{\frac{(2\pi)^{s-1}\left(s^{h-i-1}\ell\right)^s}{s^{h-i}\ell}}$$

$$= s^h \ell \sqrt{(2\pi\ell)^{s-1}} \sum_{i=0}^{h-1} \sqrt{\frac{\left(s^{h-i-1}\right)^s}{s^{h-i}}}$$

$$\le s^h \ell \sqrt{(2\pi\ell)^{s-1}} \sum_{i=0}^{h-1}\left(\sqrt{s^{s-1}}\right)^i$$

$$= s^h \ell \sqrt{(2\pi\ell)^{s-1}} \frac{1 - \left(\sqrt{s^{s-1}}\right)^h}{1 - \sqrt{s^{s-1}}}$$

$$\le \sqrt{(2\pi)^{s-1}\left(s^h\ell\right)^{s+1}}.$$

As we already remarked, since the expected work on a bucket is polynomial in the bucket size, the expected overall work for finding splittings will be linear in the number of keys. Our computation does not apply to a generic splitting strategy, but the strategies we will use in practice will be covered by our analysis.

5.2 Searching for bijections

Analogously, we can estimate the probability of finding a MPHF on a leaf of size m. The probability of hitting a bijection is $m!/m^m$, as there are $m!$ bijections and m^m overall possible functions from the leaf to itself, so the average number of trials is $m^m/m! \approx e^m/\sqrt{2\pi m}$ using Stirling's approximation. Note that since leaf sizes are bounded by a constant, the overall work for finding bijections is linear in the number of keys.

5.3 Invariance

A useful feature of our technique is that it satisfies an *invariance* property: the probability of finding a bijection on a set of m elements is equal to the probability of finding a splitting and finding a bijection on each part. More formally,

Theorem 1 *Given a set of m elements, the probability of finding a minimal perfect hash on the set is equal to the probability of finding a split in s parts k_0, $k_1, \ldots, k_{s-1}$, multiplied by the probabilities of finding a bijection on each part.*

Proof. Immediate, as

$$\frac{m!}{m^m} \prod_{i=0}^{s-1} \frac{(k_i)^{k_i}}{k_i!} \prod_{i=0}^{s-1} \frac{k_i!}{(k_i)^{k_i}} = \frac{m!}{m^m}. \quad \square$$

As an inductive consequence, for *every* tree defining splittings and bijections on a set of size m, the product of the probabilities on all nodes is constant, and equal to $m!/m^m$.

5.4 A splitting strategy

We start from the observation that if we had no limit on the construction time, searching for a bijection would provide a space-optimal structure in expectation. This happens because the optimal parameter $r(p)$ (i.e., the length of the fixed part) of a Golomb-Rice code [Sal07] for a geometrically distributed source with parameter p, i.e., $k \ge 0$ appears with probability $(1-p)^k p$, is given by [Kie04]

$$(4) \qquad r(p) = \max\left\{0, \left\lceil \lg\left(-\frac{\lg\varphi}{\lg(1-p)}\right)\right\rceil\right\},$$

where $\varphi = \left(\sqrt{5}+1\right)/2$ is the golden ratio. For this choice, the expected length of the unary part is

$$\frac{1}{1 - (1-p)^{2^{r(p)}}},$$

which is always between φ and $1 + \varphi$. In particular, in expectation the length of the codeword for trials with probability p when $p \to 0$ is

$$1 + \varphi + \left\lceil \lg\left(-\frac{\lg\varphi}{\lg(1-p)}\right)\right\rceil = 1 + \varphi + \left\lceil \lg\frac{\ln\varphi}{p + O(p^2)}\right\rceil$$

$$= 1 + \varphi + \lceil \lg\ln\varphi - \lg p - \lg(1 + O(p))\rceil$$

$$= c + \lg\frac{1}{p} + O(p),$$

with $1.56 \approx 1 + \varphi + \lg\ln\varphi \le c < 2 + \varphi + \lg\ln\varphi \approx 2.56$. Essentially, modulo less than three bits (which we pay for not using full Golomb codes, for rounding and for instantaneousness) the expected length of the codeword is the logarithm of its expected value. If we apply this formula to the case of a bijection on m elements, we obtain

$$c + \lg\frac{m^m}{m!} + O\left(\frac{m!}{m^m}\right) = m\lg e + O(\lg m)$$

so for large enough m we can get arbitrarily close to the lower bound.[1] More interestingly, the invariance property tells us that if we have probabilities p_0, $p_1, \ldots, p_{t-1}$ at the nodes of a splitting tree for m keys, then $\prod_{i=0}^{t-1} p_i = m!/m^m$. So now the cost of storing the tree is in expectation

$$\sum_{i=0}^{t-1}\left(c + \lg\frac{1}{p_i} + O(p_i)\right) = ct + \lg\prod_{i=0}^{t-1}\frac{1}{p_i} + O\left(\sum_{i=0}^{t-1} p_i\right)$$

$$= ct + m\lg e + O(\lg m).$$

In other words, we lose about c bits for each split and bijection, but otherwise the number of bits will tend to the optimal value when m grows (and t is fixed, which in particular means that the leaf size has to grow). Increasing the number of splits will thus speed up construction and slow down lookups, but just

---

[1]For feasible leaf sizes, the per key expected cost of the unary prefix for bijections is $\gtrless 0.1$.

slightly increase the space usage, as long as the number of keys going through the split is sufficiently large. The main space loss associated with c, however, is that due to the leaves, as it is amortized over the smallest number of keys.

Armed with the information gathered so far, we are now going to describe our splitting strategy. Our main drive is that of making the choice of the leaf size the main factor in bounding the time required to build the data structure, and we will assume $\ell \leq 24$ as the search on larger leaves is simply too slow; besides, the Golomb-Rice parameter in this range is smaller than 32 and can be packed into just five bits. We expect that larger leaves will lead to less evaluations and thus to structures with faster lookups, too.

From (3) it is clear that if we want to flatten the splitting tree, we should use larger values of s in the lower levels. We thus define the following criterion: starting from the bottom, we want to aggregate leaves so that the work that is necessary to compute the splitting is the same work as computing the leaf bijections, combined. We assume that each trial for a splitting will need m function evaluations, and that each trial for a bijection will require $\sqrt{\pi m/2}$ function evaluations [Ram12, FGKP95]; both estimations are approximations. If we look for the integer minimizing

$$
(5) \qquad \left| s\sqrt{\frac{\pi\ell}{2}}\frac{\ell^\ell}{\ell!} - s\ell\frac{(s\ell)^{s\ell}}{(s\ell)!}\left(\frac{\ell!}{\ell^\ell}\right)^s \right|
$$

for $\ell \leq 24$ we obtain $s = \max\{\,2, \lceil 0.35\ell + 0.5\rceil\,\}$. Thus, we will greedily try to aggregate this number of leaves.

If we go up another level, we can ask the same: that is, when the fanout is such that the work done is the same as the work done on the two lower levels. This leads to minimizing

$$
\left| st\sqrt{\frac{\pi\ell}{2}}\frac{\ell^\ell}{\ell!} - ts\ell\frac{(ts\ell)^{ts\ell}}{(ts\ell)!}\left(\frac{(s\ell)!}{(s\ell)^\ell}\right)^t \right|
$$

assuming s is the (exact) solution of (5). Once again, we can numerically compute the best integer solutions, obtaining $t = \lceil 0.21\ell + 0.9\rceil$ for $7 \leq \ell \leq 24$, and 2 for $\ell < 7$.

One can continue with further aggregation levels using the same criterion, but the impact on the space used by the data structure becomes negligible, and each aggregation level requires a test in the lookup code. After the second aggregation level we thus fix the fanout at 2, and use as unit (and left part) $\lceil \lfloor m/2\rfloor/st\ell\rceil \cdot st\ell$, with s and t as above.

6 Data representation

The following data must be stored to be able to evaluate a RecSplit MPHF efficiently:

- The initial bucket-assignment function g.

- For each bucket, a representation of its splitting tree, including the indices stored at each node.

- Since we build MPHFs independently for each bucket, we need to store the *prefix sums* of the number of keys in each bucket (i.e., for each bucket, the number of keys in all previous buckets).

- Finally, we will concatenate the representations of all buckets in a single bit array: we will thus need to store the *offset* (i.e., the starting position) of each bucket in the array.

6.1 The bucket-assignment function g

First, we replace every key $x \in S$ with a unique signature $s(x)$ of $\Theta(n)$ bits (in our implementation, 128). The bucket assignment is generated using *fixed-point multiplication*: we interpret the value $u(x)$ of the upper t bits (e.g., $t = 64$) of $s(x)$ as a real number $\alpha(x) = u(x)/2^t$ in the interval $[0\mathinner{\ldotp\ldotp}1)$ represented in fixed-point arithmetic, and we assign to x the bucket $\lfloor\alpha(x)\cdot\lceil n/b\rceil\rfloor$. If we interpret $\alpha(x)$ as a real random uniform value in the unit interval, this operation correspond to an *inversion* [Dev86] returning a random uniform discrete value in $\lceil n/b\rceil$, which is exactly what we need.[2] Since we use fixed-point arithmetic, this amounts to computing $\lfloor u(x)\cdot\lceil n/b\rceil/2^t\rfloor$, which can be performed with a multiplication and a shift. We will use the same technique when searching for splittings and bijections.

6.2 Splitting trees

Splitting trees will form the bulk of the space occupied by a RecSplit instance. It is thus essential to devise a parsimonious representation that is at the same time quickly accessible.

To attain these goals, we will not store the shape of the tree (e.g., pointers to children). Instead, we only write the indices associated with each node in preorder using an optimal Golomb-Rice code, in a bit array. We remind that once the splitting strategy has been fixed, the Golomb-Rice parameter of each index is known, as it depends only on the number of keys associated with the node.[3]

Since the number of children of every node is defined by the splitting strategy, we do not need to store them explicitly. However, when navigating the tree top-to-bottom, whenever we do not move to the first child of a node, we need to recursively skip the subtrees associated to the previous children.

In principle, to skip a subtree, we only need to know the Golomb-Rice parameter associated with the root: we then skip the first code, and recurse into each child. This strategy would be, however, very slow.

We thus write all the constant-length part of the codes, followed by all their unary parts (see Figure 2) in the bit array. This choice makes skipping the constant-length part of a subtree very fast, as we just have to move a pointer: the amount of

---

[2]Indeed, fixed-point inversion has been since the early days the technique of choice for turning the bits returned by a pseudorandom number generator into a uniform discrete value in a finite range (see, e.g., [Knu86]).

[3]Empirically, using optimal Golomb codes reduces the size of the structure by less than 1%. The fact that Golomb-Rice codes have a fixed part of constant length will have an important part in the implementation.

movement can be precomputed in a table (for a fixed splitting strategy, it depends only on the number of keys associated with the node). Using the same idea, we can retrieve the overall length of the fixed part, that is, the starting point of the unary codes, without additional information.

Skipping a number of unary codes is actually a *selection* operation (e.g., find the k-th one in a bit array), for which very fast *broadword programming* algorithms exist [Vig08, GP14].[4] We need to know the number of ones to skip, which corresponds to the number of skipped nodes: also this number can be precomputed and stored in a table.

6.3 Prefix sums and offsets

Finally, the prefix sums of bucket sizes s_i and the offset of each bucket o_i are stored using a customized Elias-Fano representation, which is a succinct representation of monotone sequences [Eli74, Fan71].

First, we compact our data by exploiting the dependence between these two values: we store s_i and $o_i - \beta s_i$, where β is the number of bits per key that are necessary to store the splitting trees. Then, as it is customary, for both sequences we compute the minimum delta δ between successive elements and use it to rescale the sequences by subtracting $i\delta$ from the i-th element: this operation reduces the range of the sequences and, correspondingly, the bits per element. If the minimum delta between items of the modified list $o_i - \beta s_i$ is negative, the rescaling will *enlarge* the elements of the list so that it is again monotone.

The Elias-Fano representation uses space proportional to the logarithm of the average gap between two elements, which in our case is b or αb. Thus, the space used to store prefix sums can be reduced arbitrarily by enlarging the target bucket size b (at the price of slower construction and lookup).

7 Implementation details

7.1 Logarithms

In several computations we need to estimate the closest integer to $\lg x$, that is, $\lfloor \lg x + 1/2 \rfloor$. For this, we use the approximate formula

$$\lambda(x + (x \gg 1)) = \lfloor \lg(x + \lfloor x/2 \rfloor) \rfloor \approx \lfloor \lg(x + x/2) \rfloor$$
$$= \lfloor \lg x + \lg 3 - 1 \rfloor \approx \lfloor \lg x + 0.58 \rfloor,$$

which contains only integer operations.

7.2 Choosing the parameter for Golomb-Rice codes

The optimal parameter $r(p)$ of a Golomb-Rice code for a geometrically distributed source with parameter p is given by (4).

The possible values of p depend only on m, as the splitting strategy specifies univocally the number and sizes of the parts in which to perform the split, so we can precompute the possible values of $r(p)$ and store the results in a table for the most likely bucket sizes.[5]

In the (practically negligible) case of large buckets we developed the following integer approximation for the optimal parameter when splitting m elements in s parts $k_0, k_1, \ldots, k_{s-1}$:

$$\left(((s-1) \cdot 5 \gg 1) + \sum_{i=0}^{s-1} \lambda\big(k_i + \big(k_i \gg 1\big)\big) - \lambda(m) \right) \gg 1.$$

Analogously, we store the optimal Golomb-Rice codes for bijections in a table, as they will be needed only for a few dozen leaf sizes.

7.3 Avoiding correlation

For our estimations of the number of trials to be applicable, every search for a split or for a bijection must be independent. We will compute both splittings and bijections on the $\Theta(n)$-bit signatures used for the bucket assignments, but discard the upper t bits, so we will be working, in each bucket, with a set of random signatures. When recursively descending during the splitting procedure, we take care of never reusing functions that have been already searched through. Thus, along a path independence is guaranteed by the independence of the functions φ_i^k, whereas on different parts it is guaranteed by the fact that the involved keys are actually a random set of signatures with an empty intersection.

7.4 Customizing Elias–Fano

We briefly recall the details of the Elias–Fano representation. We assume to have a monotonically non-decreasing sequence of $n > 0$ natural numbers

$$0 \le x_0 \le x_1 \le \cdots \le x_{n-2} \le x_{n-1} \le u,$$

where $u > 0$ is any upper bound on the last value. We will represent such a sequence in two bit arrays as follows:

- the lower $\ell = \max\{ 0, \lfloor \lg(u/n) \rfloor \}$ bits of each x_i are stored explicitly and contiguously in the *lower-bits array*;

- the upper bits are stored in the *upper-bits array* as a sequence of unary-coded gaps.

One then puts a *selection* data structure on the high bits, so that the position p of the k-th one can be found in constant time. At that point, the upper bits of the k-element are just $p - k$, and the lower bits can be retrieved directly.

First, we notice that since the lower bits have fixed width, if you need to work, as in our case, with *two* Elias–Fano representations of the same length and you always need to access the

---

[4]On the Intel cores after Haswell the PDEP instruction makes it possible to perform selection in a word using just three instructions.

[5]A single table of 32-bit integers indexed by bucket size is sufficient to memorize the Golomb-Rice parameter of a node, the skipping information of its subtree and the number of nodes in the subtree when $\ell \ge 4$ and $b \le 2000$.
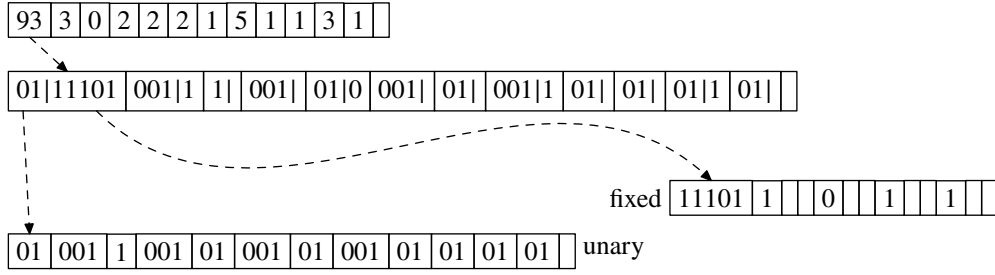
Figure 2: An example of coding for the tree in Figure 1. The indices are first laid out in preorder. Then each index is represented using a suitable Golomb-Rice code (the unary part is separated by a bar). Then the unary and fixed parts are stored separately.

elements with the same position in the two lists, the lower bits can be interleaved, saving one cache miss.

Then, we notice that the gaps between elements of the list are extremely regular, as they are bucket sizes, or the number of bits used to encode a bucket. Because of this property, the selection data structure can be significantly simplified to a two-level inventory: we record the position of the ones of index multiple of 2^{14} in a 64-bit integer, and then use an array of 16-bit integers to record the position of the intermediate ones whose index is a multiple of 2^q for some q (in our code, $q = 8$). For the remaining ones we perform a local linear search using broadword programming, starting from the closest one in the inventory.

By interleaving this information from the two lists, we often save a further cache miss.

8 Experiments

In this section, we present the results of our experiments, which were performed on an Intel® Core™ i7-7770 CPU @3.60GHz (Kaby Lake), with 64 GiB of RAM, Linux 4.17.19 and Java 12. For the C code, we used the GNU C compiler 8.1.1.

Time is measured by wall clock. For the lookup timings we report, the relative standard deviation is below 5%. Construction times (which include reading the input, generating the data structure and serializing it) have a bit more variability due to I/O. We fix the CPU clock to avoid variations due to throttling from the Turbo Boost controller.

We compare RecSplit with the three current state-of-the-art MPHF implementations:

- GOV [GOV16][6] is a structure based on random hypergraphs that provides maps at 2.24 bits per key with fast access.

- CHD [BBD09][7] has been discussed in Section 3.3. It is by far the slowest map we tested, but it is the only one that can reach about 2 bits per key. We report data for $\lambda = 5$ and $\lambda = 6$. We have not been able to build maps with $\lambda = 7$ beyond a few thousand keys.

- BBHash [LRCP17][8] is an implementation of fingerprinting-based minimal perfect hashing (see Section 3.4) that aims at being very fast in construction and lookup, but uses a large amount of space. We tested the version using the smallest amount of space ($\gamma = 1$), which however needs more than 3 bits per key to be stored.

In our experiment, we build and evaluate maps on 128-bit random keys: this way, we significantly increase the resolution of our results, as we cut off the time that is necessary to compute hashes of long keys. Since all implementations we consider hash as a first step their keys into random short keys, our choice of keys has no impact on the behavior of the structures.

RecSplit has different behavior depending on the parameters ℓ and b: increasing ℓ leads to smaller data structures *and* faster lookups, at the price of a greater construction time. Increasing b can further decrease space (as the Elias–Fano lists are better amortized), at the price of a slightly larger construction and lookup times. Figure 3 illustrates the dependence of space from these two parameters.

Among the many possible variations, we isolate:

- $\ell = 8$, $b = 100$ breaks the 2 bits/key barrier of CHD, providing better space, as well as significantly faster construction and lookup.

- $\ell = 12$, $b = 9$ uses less space than GOV and BBHash; it is faster or comparable in lookup time, but requires a longer time to build in a single threaded, non-distributed setting.

- $\ell = 5$, $b = 5$ uses less space than BBHash, and it is faster to build at large sizes.

- $\ell = 16$, $b = 2000$ is at the boundary of feasibility in construction, as it requires almost two milliseconds per key: nonetheless, it reaches 1.56 bits per key, that is, 8.3% from the lower bound. Its lookup time is significantly larger than that of GOV or BBHash, but it is still faster than CHD.

In Table 1 and 2 we report space usage in bits per key, and the construction and lookup time in nanoseconds per key. In

---

MPHF	b/key	Constr.	Lookup
CHD ($\lambda = 5$)	2.07	1627	91
CHD ($\lambda = 6$)	2.01	6440	92
$\ell = 8, b = 100$	1.79	1027	67
GOV	2.26	4210	58
$\ell = 12, b = 9$	2.23	7245	39
BBHash ($\gamma = 1$)	3.07	96	33
$\ell = 5, b = 5$	2.95	139	47
$\ell = 16, b = 2000$	1.56	1733801	87

Table 1: Space usage, construction time and lookup time, in nanoseconds per key, for a million keys.

MPHF	b/key	Constr.	Lookup
CHD ($\lambda = 5$)	2.07	6310	447
CHD ($\lambda = 6$)	2.01	25550	444
$\ell = 8, b = 100$	1.80	1063	227
GOV	2.25	1007	210
$\ell = 12, b = 9$	2.23	7331	189
BBHash ($\gamma = 1$)	3.06	290	172
$\ell = 5, b = 5$	2.95	181	241

Table 2: Space usage, construction time and lookup time, in nanoseconds per key, for a billion keys.
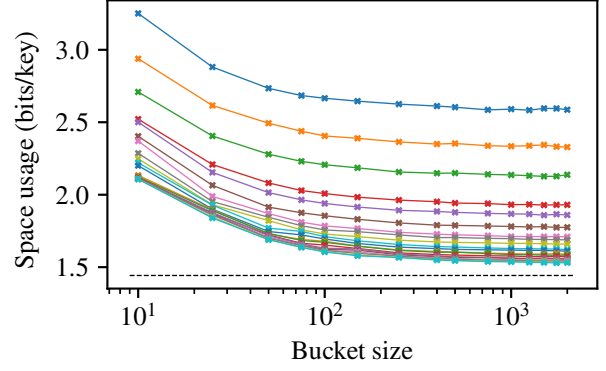

Figure 3: Space usage depending on bucket size and leaf size. Lines from top to bottom correspond to values of ℓ from 2 to 21. The dashed line shows the lower bound $\lg e \approx 1.44$.
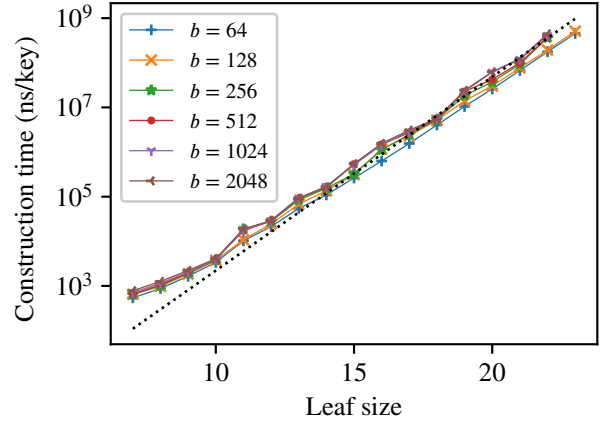

Figure 4: Construction time per bucket size, depending on the leaf size. The dashed line is $e^x/10$.

each group, we compare an alternative (CHD, GOV, BBHash) with RecSplit. We remark that for all structures we consider, the number of bits per key is essentially independent from the size of the key set.

An immediate observation is that in the case of Table 1 all structures fit the cache, whereas in the case of Table 2 they do not. As consequence, the lookup times are an order of magnitude larger, even though all structures perform lookups in constant time.

A more global view depending on the number of keys is given by Figure 5 and 6. Note that all code is in C or C++, except for the construction of GOV, which is available only in Java.

Immediately evident is that for each of the state-of-the-art maps we compare with, *there is an instantiation of RecSplit that is comparable in space and/or lookup speed, and improves at least one of the two*. Figure 7 shows the *Pareto frontier* (the set of coordinate-wise minimal points) of the space of pairs giving space usage and lookup time: below 3 bits per element, these
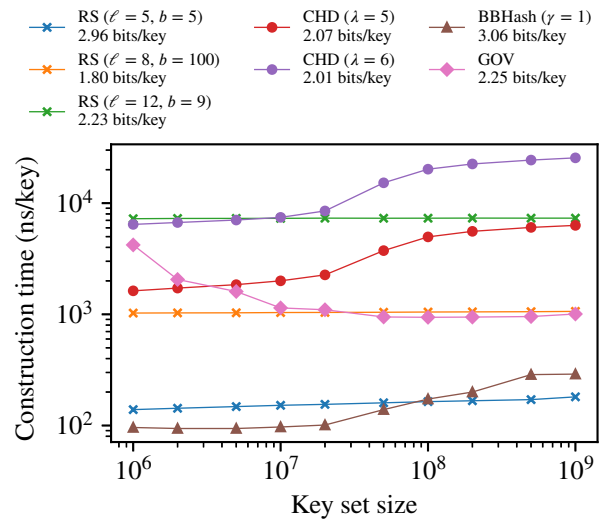

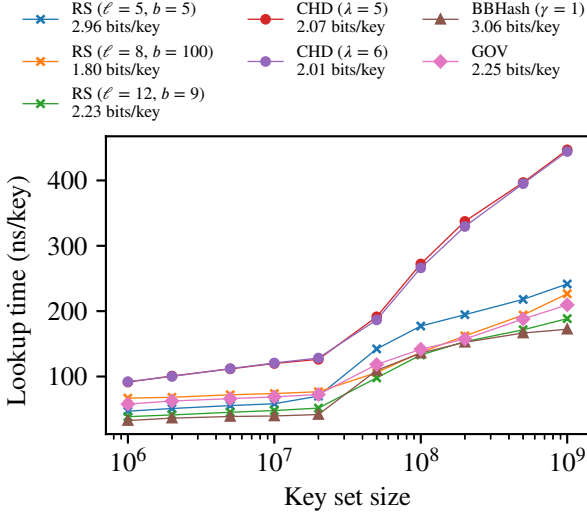Figure 5: Construction time depending on the key set size.

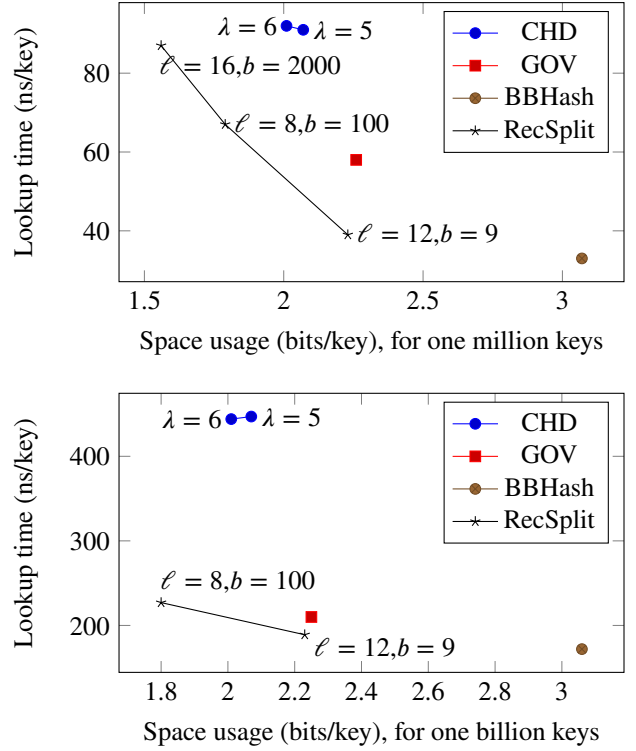Figure 6: Lookup time depending on the key set size.

Figure 7: Scatter plots of bits per key vs. lookup time. Points on the Pareto frontier are in the lower left region. The upper half is based on Table 1, whereas the lower on Table 2.

elements are all RecSplit instances. BBHash has a slightly better lookup speed ($\approx 10\%$ faster) at the price of a $\approx 40\%$ larger space.

If we consider the three-dimensional Pareto frontier, the situation is different because instances of RecSplit with fast lookup require usually more construction time. However, RecSplit improves on all three parameters at the same time with respect to CHD, which is presently the state of the art for small space. Moreover, both CHD and BBHash require that as much RAM as the final size of the data structure is available. On the contrary, GOV and RecSplit in principle allow to build the structure on-disk, possibly in a distributed fashion, using a very small amount of RAM, as both structures distribute keys in small buckets which can be processed independently.

Another interesting advantage of RecSplit is that, as we already mentioned, increasing ℓ leads to slower construction, but at the same time decreases space and makes lookup faster, whereas for BBHash, building a smaller map means slower construction and slower lookups; for CHD, a smaller map means slower construction, and no effect on lookups.

If we consider the points associated with RecSplit instances with $\ell \geq 4$ and $64 \leq b \leq 2048$ in isolation, they are essentially all on the three-dimensional Pareto frontier. This means that all these parameter choices provide different tradeoffs between our three measures.

From the figures, it is evident that the in-memory construction of CHD and BBHash induces a significantly nonlinear behavior due to cache and Translation Lookahead Buffer misses as the number of keys increase. This does not happen for GOV and RecSplit.

Finally, in Figure 4 we show the increase in construction time depending on the leaf size for a choice of exponentially spaced bucket sizes. As it is evident, increasing the leaf size by one approximately multiplies the construction time by e, as expected from the discussion in Section 5.4, whereas the bucket

size has less impact.

9 Conclusions

We have presented RecSplit, a new static data structure storing a minimal perfect hash function with expected linear-time construction and expected constant-time lookup. RecSplit is the first data structure able to break the 2 bits/key barrier in practice, and instances very close to the lower bound can be built feasibly, albeit slowly. Construction time can be reduced by using a distributed computational setting such as MapReduce [DG08], as fixed-point inversion is monotone, so buckets can be computed by a linear scan of the (offline) sorted signatures, and then the splitting tree of each bucket can be built independently.

By enlarging the size of the codomain of bijections and splittings, the techniques described in this paper can also be used to build extremely compact perfect (but not minimal) hash functions. We leave this extension for future work.

References

[BBD09]   Djamal Belazzougui, Fabiano C. Botelho, and Martin Dietzfelbinger. Hash, displace, and compress. In Amos Fiat and Peter Sanders, editors, *Algorithms - ESA 2009, 17th Annual European Symposium, Copenhagen, Denmark, September 7-9, 2009. Proceedings*, pages 682–693, 2009.

[BPZ13] Fabiano C. Botelho, Rasmus Pagh, and Nivio Ziviani. Practical perfect hashing in nearly optimal space. *Information Systems*, 38(1):108–131, 2013.

[CKRT04] Bernard Chazelle, Joe Kilian, Ronitt Rubinfeld, and Ayellet Tal. The Bloomier filter: an efficient data structure for static support lookup tables. In J. Ian Munro, editor, *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004*, pages 30–39. SIAM, 2004.

[Dev86] Luc Devroye. *Non-uniform random variate generation*. Springer Verlag, 1986.

[DG08] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.

[Eli74] Peter Elias. Efficient storage and retrieval by content and address of static files. *J. Assoc. Comput. Mach.*, 21(2):246–260, 1974.

[Fan71] Robert 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.

[FGKP95] Philippe Flajolet, Peter J. Grabner, Peter Kirschenhofer, and Helmut Prodinger. On Ramanujan's Q-function. *J. Comp. Appl. Math.*, 58(1):103–116, 1995.

[FKS84] Michael L. Fredman, János Komlós, and Endre Szemerédi. Storing a sparse table with $O(1)$ worst case access time. *J. Assoc. Comput. Mach.*, 31(3):538–544, July 1984.

[GOV16] Marco Genuzio, Giuseppe Ottaviano, and Sebastiano Vigna. Fast scalable construction of (minimal perfect hash) functions. In Andrew V. Goldberg and Alexander S. Kulikov, editors, *Experimental Algorithms: 15th International Symposium, SEA 2016, St. Petersburg, Russia, June 5-8, 2016, Proceedings*, number 9685 in Lecture Notes in Computer Science, pages 339–352. Springer, 2016.

[GP14] Simon Gog and Matthias Petri. Optimized succinct data structures for massive data. *Software: Practice and Experience*, 44(11):1287–1314, 2014.

[HT01] Torben Hagerup and Torsten Tholey. Efficient minimal perfect hashing in nearly minimal space. In Afonso Ferreira and Horst Reichel, editors, *STACS 2001, 18th Annual Symposium on Theoretical Aspects of Computer Science, Dresden, Germany, February 15-17, 2001, Proceedings*, pages 317–326, 2001.

[Jae81] Gerhard Jaeschke. Reciprocal hashing: A method for generating minimal perfect hashing functions. *Comm. ACM*, 24(12):829–833, 1981.

[Kie04] Anthony Kiely. Selecting the Golomb parameter in Rice coding. Technical Report 42-159, Jet Propulsion Laboratory, California Institute of Technology, 2004.

[Kno08] Andreas Knoblauch. Closed-form expressions for the moments of the binomial probability distribution. *SIAM Journal on Applied Mathematics*, 69(1):197–204, 2008.

[Knu86] Donald E. Knuth. *The Metafont book*. Addison-Wesley, 1986.

[Knu11] Donald E. Knuth. *The Art of Computer Programming: Volume 4, Combinatorial algorithms. Part 1*, volume 4A. Addison-Wesley, 2011.

[LRCP17] Antoine Limasset, Guillaume Rizk, Rayan Chikhi, and Pierre Peterlongo. Fast and Scalable Minimal Perfect Hashing for Massive Key Sets. In Costas S. Iliopoulos, Solon P. Pissis, Simon J. Puglisi, and Rajeev Raman, editors, *16th International Symposium on Experimental Algorithms (SEA 2017)*, volume 75 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:16. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.

[MSSZ14] Ingo Müller, Peter Sanders, Robert Schulze, and Wei Zhou. Retrieval and perfect hashing using fingerprinting. In Joachim Gudmundsson and Jyrki Katajainen, editors, *SEA 2014: Experimental Algorithms*, pages 138–149. Springer International Publishing, 2014.

[MWHC96] Bohdan S. Majewski, Nicholas C. Wormald, George Havas, and Zbigniew J. Czech. A family of perfect hashing methods. *Comput. J.*, 39(6):547–554, 1996.

[Ram12] Srinivasa Ramanujan. On question 294. *J. Indian Math. Soc.*, 4:151–152, 1912.

[Sal07] David Salomon. *Data Compression*. Springer–Verlag London, 2007.

[Spr77] Renzo Sprugnoli. Perfect hashing functions: a single probe retrieving method for static sets. *Comm. ACM*, 20(11):841–850, 1977.

[Vig08] Sebastiano Vigna. Broadword implementation of rank/select queries. In Catherine C. McGeoch, editor, *Experimental Algorithms. 7th International Workshop, WEA 2008*, number 5038 in Lecture Notes in Computer Science, pages 154–168. Springer–Verlag, 2008.