



PDF Download
3797036.pdf
14 March 2026
Total Citations: 0
Total Downloads: 113

Latest updates: <https://dl.acm.org/doi/10.1145/3797036>

SURVEY

Modern Minimal Perfect Hashing: A Survey

HANS-PETER LEHMANN, Karlsruhe Institute of Technology, Karlsruhe, Baden-Wurttemberg, Germany

THOMAS MUELLER

RASMUS PAGH, University of Copenhagen, Copenhagen, Hovedstaden, Denmark

GIULIO ERMANNO PIBIRI, Ca' Foscari University of Venice, Venice, VE, Italy

PETER SANDERS, Karlsruhe Institute of Technology, Karlsruhe, Baden-Wurttemberg, Germany

SEBASTIANO VIGNA, University of Milan, Milan, MI, Italy

View all

Open Access Support provided by:

Karlsruhe Institute of Technology

University of Milan

University of Copenhagen

Ca' Foscari University of Venice

Published: 11 March 2026
Online AM: 11 February 2026
Accepted: 04 February 2026
Revised: 18 January 2026
Received: 06 June 2025

Citation in BibTeX format

Modern Minimal Perfect Hashing: A Survey

HANS-PETER LEHMANN, Karlsruhe Institute of Technology, Karlsruhe, Germany

THOMAS MUELLER, Independent researcher, Basel, Switzerland

RASMUS PAGH, BARC, University of Copenhagen, Copenhagen, Denmark

GIULIO ERMANNIO PIBIRI, Ca' Foscari University of Venice, Venice, Italy

PETER SANDERS, Karlsruhe Institute of Technology, Karlsruhe, Germany

SEBASTIANO VIGNA, Università degli Studi di Milano, Milan, Italy

STEFAN WALZER, Karlsruhe Institute of Technology, Karlsruhe, Germany

Given a set S of n keys, a perfect hash function for S maps the keys in S to the first $m \geq n$ integers without collisions. It may return an arbitrary result for any key not in S and is called *minimal* if $m = n$. The most important parameters are its space consumption, construction time, and query time. Years of research now enable modern perfect hash functions to be extremely fast to query, very space-efficient, and scale to billions of keys. Different approaches give different trade-offs between these aspects. For example, the smallest constructions get within 0.1% of the space lower bound of $\log_2 e$ bits per key. Others are particularly fast to query, requiring only one memory access. Perfect hashing has many applications, for example, to avoid collision resolution in static hash tables, and is used in databases, bioinformatics, and stringology.

Since the last comprehensive survey in 1997, significant progress has been made. This survey covers the latest developments and provides a starting point for getting familiar with the topic. Additionally, our extensive experimental evaluation can serve as a guide to select a perfect hash function for use in applications.

CCS Concepts: • **Theory of computation** → **Data compression; Bloom filters and hashing**; • **Information systems** → **Point lookups**; • **General and reference** → **Surveys and overviews**;

Additional Key Words and Phrases: Compressed data structures, minimal perfect hashing, survey

This article is based on and has text overlaps with the doctoral thesis of Hans-Peter Lehmann [106]. We would like to thank the anonymous reviewers for their valuable feedback.



H.-P.L., P.S., and S.W.: This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 882500). This work was also supported by funding from the pilot program Core Informatics at KIT (KiKIT) of the Helmholtz Association (HGF). R.P.: Rasmus Pagh is part of BARC, supported by the VILLUM Foundation grant 54451. S.V.: This work is also partially supported by the project "SEcurity and Rights In the Cyberspace - SERICS" (PE00000014 - CUP H73C2200089001) under the National Recovery and Resilience Plan (NRRP) funded by the European Union - NextGenerationEU. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the Italian MUR. Neither the European Union nor the Italian MUR can be held responsible for them.

Authors' Contact Information: Hans-Peter Lehmann (corresponding author), Karlsruhe Institute of Technology, Karlsruhe, Germany; e-mail: research@hanspeter-lehmann.de; Thomas Mueller, Independent researcher, Basel, Switzerland; e-mail: thomasm@apache.org; Rasmus Pagh, BARC, University of Copenhagen, Copenhagen, Capital Region of Denmark, Denmark; e-mail: pagh@di.ku.dk; Giulio Ermanno Pibiri, Ca' Foscari University of Venice, Venice, Veneto, Italy; e-mail: giulioermanno.pibiri@unive.it; Peter Sanders, Karlsruhe Institute of Technology, Karlsruhe, Baden-Württemberg, Germany; e-mail: sanders@kit.edu; Sebastiano Vigna, Università degli Studi di Milano, Milan, Lombardy, Italy; e-mail: sebastiano.vigna@unimi.it; Stefan Walzer, Karlsruhe Institute of Technology, Karlsruhe, Baden-Württemberg, Germany; e-mail: stefan.walzer@kit.edu.



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

© 2026 Copyright held by the owner/author(s).

ACM 0360-0300/2026/03-ART251

<https://doi.org/10.1145/3797036>

ACM Reference Format:

Hans-Peter Lehmann, Thomas Mueller, Rasmus Pagh, Giulio Ermanno Pibiri, Peter Sanders, Sebastiano Vigna, and Stefan Walzer. 2026. Modern Minimal Perfect Hashing: A Survey. *ACM Comput. Surv.* 58, 10, Article 251 (March 2026), 36 pages. <https://doi.org/10.1145/3797036>

1 Introduction

Given a set $S \subseteq U$ of n keys drawn from a universe set U of size u , a perfect hash function (PHF) for S maps the keys in S to the first $m \geq n$ integers $[m] = \{0, \dots, m-1\}$ without collisions. For any key not in S , it may return an arbitrary result. Therefore, it does not need to be able to distinguish between keys in S and keys that it was not constructed with. As we will see later, this enables a constant amount of space per key regardless of the nature of keys. Most practical implementations of such functions need, in fact, less than 2 bits/key, depending on the load factor n/m . In contrast, a data structure storing a set of n keys must use at least $\log_2(eu/n) - o(1)$ bits/key, which could be much larger than 2 when $u \gg n$. A PHF is called *minimal* (MPHF) if $m = n$, in which case it represents a bijection between the input set and the first n integers. Intuitively, the smaller output range reduces the flexibility of how keys can be placed, thus reducing the number of functions that are collision-free. This increases the space consumption compared to the non-minimal case (details in a later paragraph). For an MPHF, the space lower bound is about $\log_2 e \approx 1.443$ bits per key, and practical applications can already achieve 1.444 bits per key [104]. In contrast to cryptographic hashing, properties of the inverse function are irrelevant. We follow the practice in the literature and focus on *minimal* PHFs only. Note that most of the techniques we describe have configurations to construct non-minimal PHFs as well. In fact, some techniques construct a PHF first and later convert it to an MPHF (see Section 3.1). Perfect hashing has many applications in various fields of computing, from bioinformatics to stringology. Over the past years, there has been significant progress in this research area. Many different approaches have been proposed, showing a wide variety of ideas on how to solve the perfect hashing problem.

This Survey in Context. Lewis and Cook [111] review early letter-based approaches [38, 42–44, 58, 70, 86, 96, 148, 152]. The last survey about minimal perfect hashing by Czech et al. [46] gives a comprehensive review of approaches until the year 1997 [4, 27, 28, 33–35, 37, 47, 51, 52, 65, 67, 68, 80, 87, 88, 93, 95, 111, 116, 122, 134, 144, 147, 149, 151, 154, 156, 161, 162, 164]. We do not go into detail about these here and instead focus on the *modern* approaches that have outperformed the former ones on modern hardware. In particular, our focus is on practical constructions, i.e., on those that have been implemented and shown to perform well in practice. As such, we also give an extensive evaluation that shows their performance in practice and helps to choose a fitting PHF for a given application. We nonetheless point out theoretical constructions when necessary.

Space Lower Bounds. The space lower bound for representing an MPHF is $n \log_2 e - \mathcal{O}(\log n) \approx 1.443n$ bits for large n and $u \rightarrow \infty$ [69, 116, 121]. This bound is quite simple to explain. Take a function $f : S \rightarrow [n]$, that is, minimal perfect on some set S . Because f can be evaluated with any input key from the universe and outputs only values from $[n]$, there can be additional input sets for which f is minimal perfect. More precisely, f is minimal perfect for all sets where exactly one input key is in each preimage $f^{-1}(1), f^{-1}(2), \dots, f^{-1}(n)$. The number of input sets on which f is minimal perfect is therefore $|f^{-1}(1)| \cdot |f^{-1}(2)| \cdot \dots \cdot |f^{-1}(n)|$. This expression is maximized if all preimages have the same size u/n . Therefore, a single function can be minimal perfect for at most $(u/n)^n$ different input sets. There are $\binom{u}{n}$ different possible input sets. Therefore, we need to be able to differentiate between at least $\binom{u}{n} / (u/n)^n$ different behaviors of our MPHF to be able to cover every possible input set. A minimal PHF then needs to store which of these behaviors has

to be used on the respective input set. Therefore, the number of bits needed is $\log_2 \left(\binom{u}{n} / \left(\frac{u}{n} \right)^n \right)$ which, by using Stirling's approximation of the factorial, is $\approx \log_2 \left(\left(\frac{ue}{n} \right)^n / \left(\frac{u}{n} \right)^n \right) - \mathcal{O}(\log n) = n \log_2 e - \mathcal{O}(\log n) \approx 1.443n$ for large n and $u \rightarrow \infty$. Using a similar combinatorial argument, one can show that the space lower bound to represent non-minimal PHFs with load factor $\alpha = n/m$ is $n \left(\log_2 e + \left(\frac{1}{\alpha} - 1 \right) \log_2 (1 - \alpha) \right) - \mathcal{O}(\log_2 n)$ [11, 69, 106, 122]. For example, a PHF with a load factor of $\alpha = 0.5$ already has a lower bound of $\approx 0.443n$ bits.

These lower bounds also hold for the *expected* space requirement of randomized constructions of MPHFs and even if the input set is chosen uniformly at random. For randomized MPHFs the lower bounds are tight as we will see for $n = m$ in Section 6: a randomized brute force construction [122] requires $n \log_2 e - \Omega(\log n)$ bits in expectation, even for infinite universes.

A completely independent lower bound of $\Omega(\log \log u)$ bits in total [69] holds for *deterministic* constructions of MPHFs in the *worst case* even for $n = 2$ because $\log_2 u$ functions $f : U \rightarrow [2]$ are required for each pair of elements from the universe is separated by at least one such function. Since this lower bound does not apply to randomized constructions and exceeds the more general lower bound only if $n = \mathcal{O}(\log \log u)$ we leave it aside for the rest of the paper.

Variants. There are several variants of perfect hashing. Because these use independent techniques, we only mention them briefly but do not go into detail. *k-PHF*s are a generalization that allow for up to k collisions on each output value. *Order-preserving* MPHFs retain any arbitrary order of the keys, which means that they need at least $\log n$ bits per key due to storing one of the $n!$ permutations. *Monotone* MPHFs, in contrast, use the *intrinsic* order of the keys as implied by an order of the universe. Surprisingly, MMPHFs achieve a much smaller space consumption of just $\mathcal{O}(\log \log \log u)$ bits per key [6, 9, 100]. We refer to the extended version [107] for details.

Overview. We begin our journey with an overview of different applications of perfect hashing in Section 2. In Section 3, we give an introduction to common techniques. We then categorize the modern approaches into three categories based on their main working principle in Section 4. In Sections 5 to 7 we then discuss the approaches in those categories: Perfect hashing through retrieval in Section 5, brute-force in Section 6, and fingerprinting in Section 7. In Section 8, we give theoretical details about the space consumption and construction time of different approaches. After this overview of the literature, we give a detailed practical evaluation in Section 9. There, Figure 7 can serve as a guide to select a PHF in practice, given specific space and time budgets. We conclude this article in Section 10.

2 Applications

In the following section, we outline some of the wide range of applications of perfect hashing.

Hash Tables and Retrieval. Hash tables are one of the most fundamental data structures used today. When the key set is static, a PHF can be used to directly index the hash table without the need for collision resolution. This guarantees very fast access times [70] with as few cache faults as well-engineered hash tables (like, for example, Abseil's "Swiss" hash tables [78, 79]), and fewer external memory access operations [101]. If we do not need the ability to query whether some key is stored in the table, we can store only the values and not the keys in the table cells. We thus obtain a space-efficient dictionary with a static key set and updateable values known as a *value-dynamic* retrieval data structure. As in the static version (see Section 3.2), its behavior is undefined for any key not in the original set. The best known lower and upper bounds for updateable retrieval [103] suggest that constructions based on perfect hashing are optimal. We remark that there is also work on retrieval data structures with an incremental or dynamic key set not based on perfect hashing [49, 102, 125].

AMQ Data Structures. An Approximate Membership Query (AMQ) data structure (also called *filter*) can answer membership queries, “Is $x \in S$?”, allowing false-positive answers. The space lower bound for an AMQ data structure for S with false positive rate ϵ is $n \log_2(1/\epsilon)$ bits [32]. The most widely known example is the Bloom filter [20]. With the optimal number of hash functions, it needs at least $n(\log_2(e) \cdot \log_2(1/\epsilon))$ bits to be represented [30], thus approximately 44.3% more space than the space lower bound for an AMQ data structure. Assuming a PHF f taking c bits per key, we can easily achieve an AMQ data structure with $m(c + \lceil \log_2(1/\epsilon) \rceil)$ bits by storing $\lceil \log_2(1/\epsilon) \rceil$ -bit fingerprints of the input keys in a packed array with m cells indexed by f [17, 30, 63, 82, 119]. For example, with an MPHf we can get close to $n(\log_2(e) + \log_2(1/\epsilon))$ bits, which is $\frac{\log_2(e)}{\log_2(1/\epsilon)} \cdot 100\%$ extra space over the lower bound. Since MPHfs can have very fast construction and query time, this can be practically valuable because even more space efficient AMQ constructions based on retrieval [57] have slower queries.

Databases. Perfect hashing also finds use in static databases (or *stores*) to map variable-length keys (e.g., strings) from a known universe to smaller identifiers that can be stored with less space [23, 126]. For example, this is the case for Indeed Engineering [61]. The COPR database [145] uses perfect hashing for efficient search in log files. Perfect hashing can also be used to speed up finding association rules [2] in databases [36, 91]. SNIPS [129] uses perfect hashing to synchronize the content of a distributed database without communicating the full list of files that each peer stores. Perfect hashing can also be used to enable faster joins and aggregates in OLAP databases [72]. Finally, *monotone* minimal PHFs can be useful as part of an algorithm for range queries in databases [101, 113].

Bioinformatics. The (node-centric) De Bruijn graph is a directed graph where nodes correspond to strings of length k , called k -mers, and edges model $(k-1)$ -length overlaps between the nodes [48, 77]. Note that there is a one-to-one correspondence between a De Bruijn graph and a set of k -mers. This combinatorial object is vital to many applications in bioinformatics, such as *de novo* genome assembly [112, 135], transcriptome assembly from RNA [81], and variant calling [92]. Representing such graphs in compact space is therefore very important. Perfect hashing can be used to enable fast navigation in such graphs [39] and to support membership queries for k -mers [3, 120, 136, 137, 150]. Pibiri et al. [138] also introduce the first MPHf specifically designed for k -mer sets. They exploit the fact that two consecutive k -mers overlap by $k-1$ symbols to beat the space lower bound that applies to general-purpose minimal perfect hashing, and to even achieve locality of the hash values for overlapping k -mers. Perfect hashing is also used to build a De Bruijn graph in external memory; for example, to associate some metadata to each k -mer to aid construction [98], and to efficiently implement union-find data structures for k -mers [41].

Text Indexing. Perfect hashing can be used to index children of a node of a suffix tree [160] in constant time [13]. In general, perfect hashing can be used to store trees, for example, in prefix search [10]. In a large external memory lexicon, perfect hashing can be used to ensure that the number of memory accesses is low when accessing specific words [163]. Monotone minimal perfect hashing can be used for constant-time rank queries on an input set. As a special case, one can build a monotone MPHf on the positions of 1-bits of a bit vector [128], possibly undercutting the space consumption of storing the bit vector itself. The data structure then only supports rank queries for positions that we know have a 1-bit, which can be a common use-case. Running rank queries on an alphabet is useful to decode the Burrows-Wheeler Transform (BWT) [31] using a PHF [13]. For reporting the number of occurrences of each character in a substring, Belazzougui et al. [14] store the rank of each character occurrence in an array indexed by a monotone MPHf. From a list of documents, this then enables efficiently finding the k documents with the most occurrences of a

pattern (top- k retrieval) [14, 127]. Finally, in pattern matching [12, 73, 84], monotone MPHFs are applied mostly to integer sequences representing the occurrences of certain characters in a text.

Natural Language Processing. N -gram language models are a crucial ingredient of natural language processing, machine learning, and spell checking [118]. Pibiri and Venturini [141, 142] introduce a compressed trie representation of N -gram language models and use perfect hashing to achieve faster lookups. For real-time speech recognition at Amazon, Strimel et al. [153] introduce the compressed N -gram data structure DashHashLM, which is particularly focused on fast lookups. It computes IDs of the N -gram context using a minimal PHF. In contrast to storing the IDs explicitly, this avoids the need for a more complex lookup table.

Further Applications. In flow lookup tables in routers [115], partially dynamic PHFs get rebuilt only if they change too much (see also the extended version of this article [107]). Perfect hashing helps to perform depth-first search in large implicitly defined graphs using a value-dynamic retrieval data structure that stores one bit per node while all other data structures can reside in external memory [59]. For example, this can be used for formal verification of software via model-checking. Finally, monotone MPHFs can be used for efficient queries in encrypted data [21].

3 Preliminaries

In this section, we explain general tools and techniques that we need later in the article. We also highlight basic building blocks that are used in almost every MPH construction.

3.1 Common Perfect Hashing Tools

Fully Random Hash Functions. In the course of this article, we regularly need access to a fully random hash function $h(x)$, $x \in U$. In fact, we often need a number of different hash functions, which we here model by adding a second parameter to h : we assume that $h(x, s)$ is a fully random function for all values $s \in \mathbb{N}_0 = \{0, 1, 2, \dots\}$ and for any $x \in U$. We also assume that $h(x, s)$ can be evaluated in constant time for any $x \in U$ and $s \in \mathbb{N}_0$. This is known as the “simple uniform hashing assumption” which is common in the literature about hashing [51, 101, 108, 110, 130, 131]. The assumption is an adequate model for practical hash functions and can be justified in theory [1, 18, 55, 155].

Master Hashes. Most PHFs first generate an *initial*, “*master*”, *hash* using a high quality random hash function with large output range (e.g., with range $[2^{64}]$ or $[2^{128}]$).¹ This makes the construction mostly independent of the input distribution. Another important advantage is that the construction works with fixed-size keys (i.e., the hashes themselves) instead of the original, possibly variable-length, keys. Clearly, handling fixed-length keys aids implementation and improves performance. The idea is introduced in one of the first perfect hashing papers by Sprugnoli [152] (although initially just with a modulo operation).

Partitioning. A second fundamental idea introduced by Sprugnoli [152] is to apply *partitioning* to perfect hashing. In order to determine a PHF for a larger key set, it is possible to first divide the input keys into smaller subsets of approximately the same size, for example, using a random hash function. After determining PHFs on all subsets independently, adding a prefix sum over the subset sizes to the hash values gives a PHF of the overall set. While this causes some space and time overhead due to the need to store and lookup the prefix-sum values, it can help with cache locality and therefore construction performance. Furthermore, each partition can be built independently in parallel using multiple threads.

¹Examples include MurmurHash, xxHash, CityHash, and GxHash. The SMHasher test suite [5] offers a comprehensive comparison.

Another way to partition PHFs that avoids storing prefix sums is to construct a k -PHF. PtrHash [83] takes a different approach, to avoid prefix sums, by constructing a non-minimal PHF. In particular, each partition is allocated the same output range r . If the number of partitions is P , then the function is made minimal by *remapping*, i.e., by ignoring the extra $rP - n$ output values as explained below.

Remapping. Any PHF with output range $m > n$ can be converted to an MPHf. Folklore methods include marking the actual outputs in a vector of m bits and then performing a rank query using the output of the PHF. An alternative is to store the unused output values in a predecessor data structure and then discount the output of the PHF by the number of unused output values that precede it. Sometimes, in the first case one does not even need an additional bit vector because a special encoding of the output can be used to distinguish between used and unused entries [25]. Pibiri and Trani [139] remap the values greater than n to the unused output values in $[n]$ instead. Therefore, only a small fraction of the keys need to query the remapping data structure. Constructing a non-minimal PHF first might be faster but might also increase space consumption and query time.

3.2 Retrieval Data Structures

A (*static function* [74]) *retrieval data structure* represents a function $f : S \rightarrow \{0, 1\}^r$ that maps each key in $S \subseteq U$ to an r -bit integer. Similarly to an MPHf, the behavior of a retrieval data structure is undefined for any key not in S and, because it does not need to differentiate between keys in S and keys not in S , it can represent f in close to rn bits.

Demaine et al. [49] point out that given an MPHf p for S it is trivial to construct a retrieval data structure: simply index a packed array of r -bit values using p . This results in a space consumption of at least $n \cdot (r + \log_2 e)$ bits. For small r (such as $r = 1$), the space overhead of at least $\log_2 e$ bits/key is large percentage-wise, while for large r (such as $r = 16$), it may be acceptable. However, the packed array approach issues at least two dependent memory accesses for each query whereas many kinds of retrieval data structures issue independent memory accesses [54], which are parallelized efficiently by modern processors. Classically, three accesses for a good balance between query efficiency and space effectiveness as discussed below. The converse can be efficient: in Section 5 we use retrieval data structures to build PHFs.

One of the main approaches to the construction of retrieval data structures is due to Czech et al. [45, 117], and has been the source of most development in the area. The idea is that given the function f , one fixes an integer $k > 0$ and chooses k hash functions $h : U \times [k] \rightarrow [m]$, $m \geq n$, and uses them to write a system of n equations in k variables $\{w_i\}$

$$w_{h(x,0)} + w_{h(x,1)} + \dots + w_{h(x,k-1)} = f(x) \quad \text{for all } x \in S. \quad (1)$$

Note that “+” denotes the sum operator in an Abelian group. In the original proposal, it was modular arithmetic, but is now more commonly a bitwise XOR operation between bit vectors. Clearly, storing a solution w_0, \dots, w_{m-1} for this system is sufficient to compute f on all keys in S , but elements in $U \setminus S$ will be mapped to arbitrary values. The criterion originally used to make the system solvable was to make the ratio n/m small enough that the system can be almost always put in echelon form in linear time, instead of the typical cubic time of Gaussian elimination. For example, for $k = 2$ and any $\varepsilon > 0$ one can obtain a structure with a space of $\approx (2 + \varepsilon)rn$, i.e., more than twice the optimum, whereas for $k = 3$ the overhead decreases to $\approx 23\%$. These overheads come from *thresholds* for the *peelability* problem discussed in Section 3.3. Arbitrarily lower overheads can be obtained by increasing k and solving the linear system [74]—for example, for $k = 3$ (10% overhead) or $k = 4$ (3% overhead), as the threshold for solvability is lower [143], and tends to 1 as k grows, as we discuss in the Section 3.3. Recent retrieval data structures achieve very low space consumption:

for example, Bumped Ribbon Retrieval (BuRR) [57] is also based on solving a system of linear equations where a key's equation uses a random subset of a block of w consecutive variables, e.g., with $w = 64$. BuRR can achieve a space consumption of less than $1.017n$ and supports queries in time $O(r)$. We do not go into detail about its working principle here but still want to mention it because several approaches in Section 5.2 use BuRR as a black box retrieval data structure. As noted by Dietzfelbinger and Pagh [54], retrieval data structures can also be used to implement AMQ data structures by storing a fingerprint of each key in the set (as discussed in Section 2).

3.3 Graph Properties

Several PHF constructions that we describe later are based on *peeling* and *orienting* graphs. In this section, we give an introduction to these terms.

Peelability. Peeling is the process of iteratively taking away nodes of degree 1 from a graph, together with their adjacent edge [25, 45, 97, 117, 124, 157, 158]. We call a graph *peelable* if it can be peeled to an empty graph. Clearly peelability in ordinary graphs is equivalent to acyclicity. However, once we move to k -uniform hypergraphs, $k \geq 3$, where an edge is a subset of vertices of cardinality k , the two concepts are no longer equivalent.² Peelability is important for retrieval data structures because we can associate a k -uniform hypergraph with the system in Equation (1): vertices are elements of $[m]$, and each equation is a hyperedge containing the vertices $h(x, i)$, $0 \leq i < k$. It is easy to see that peelability implies that the system is in echelon form up to row swaps and column swaps, as each peeled edge is an equation containing a variable not appearing elsewhere. Peeling can be performed in linear time. As each edge is peeled due to a degree 1 vertex, the associated equation and variable index are put on a stack. After peeling, the equations and variables are popped from the stack and processed in reverse order. Each equation is solved by setting the variable to the required value, exploiting that the variable does not appear in previously solved equations. At the end of the process, equations are popped from the stack and solved backwards.

Czech et al. [45, 117] showed that a graph with m vertices, n uniformly random edges, and $n/m \leq 1/(2 + \epsilon)$, $\epsilon > 0$ is peelable with a probability > 0 . In practice, they suggest to use $n/m = 0.467$. For k -uniform hypergraphs, $k \geq 3$, there is a sharp *peelability threshold* peel_k : a hypergraph with m vertices and n uniformly random hyperedges of size k is peelable with probability $1 - o(1)$ if $n/m < \text{peel}_k - \epsilon$, and peelable with probability $o(1)$ if $n/m > \text{peel}_k + \epsilon$. For $k = 3$ the threshold is $\text{peel}_3 \approx 0.8185$, and for $k > 3$ the threshold decreases, so $k = 3$ is the most interesting case for practical applications, providing a 23% space overhead for retrieval. Molloy [124] proved ten years later that the thresholds conjectured in [117] were quite accurate. We remark that a slightly different distribution for the hyperedges yields larger peelability thresholds [157, 158], e.g., ≈ 0.9179 for $k = 3$, equal to the orientability thresholds we discuss next, which moreover improve for larger k . This has, to our knowledge, not yet been exploited in MPH constructions.

Orientability. Orientability is a weaker concept than peelability: we call a hypergraph *orientable* if we can *orient* each edge, that is, select a vertex from each edge so that no two edges share the same selected vertex. Clearly, a peelable graph is orientable, as we can select for each edge the vertex from which it was peeled.³ Orientability is a well known property, that is, analyzed, among others, in the context of cuckoo hashing [133]. Similarly to the peelability threshold, the *orientability threshold* describes how large m has to be with respect to n to ensure that the hypergraph can be

²Note that early papers sometimes used the term “acyclic” for “peelable”. In fact, there are several definitions of acyclicity in a hypergraph [29].

³In general, a hypergraph is (k, ℓ) -orientable if it possible to select k vertices in each edge so that no vertex is selected by more than ℓ edges. While this notion might be useful for k -perfect hashing, we do not consider it further in this article.

almost always oriented as n goes to infinity. For $k \geq 3$, this threshold is (surprisingly) equivalent to the threshold for the k -XORSAT problem [50, 143]: when considering a random system on \mathbb{F}_2 with m variables, n equations, and k variables per equation, the system is almost always solvable as n goes to infinity as long as n/m is below the k -XORSAT solvability threshold, which is the same as the orientability threshold for k -uniform hypergraphs. For example, 3-uniform hypergraphs have threshold ≈ 0.9179 , 4-uniform hypergraphs have threshold ≈ 0.9768 , and when k increases further the threshold approaches 1 [143]. The threshold has also been extended to larger fields [75]. There is no threshold for 2-XORSAT, but ordinary graphs have an orientability threshold of $1/2$, and below the same ratio they are peelable with positive probability.

4 Overview

We give an overview of the origins of perfect hashing in the extended version of this article [107]. In this survey, we focus on modern PHF constructions. For this, we divide the approaches into three categories, which we explain in the following.

Perfect hashing through retrieval (Section 5) builds on retrieval data structures (see Section 3.2). Early approaches explicitly store the desired hash value for each key, leading to a space consumption of $\mathcal{O}(\log n)$ bits per key. Later approaches achieve constant space per key by storing the index of one of several candidate output values determined by different hash functions on the key.

Simple brute-force search for a PHF achieves the optimal space consumption but has exponential running time, making it impractical. In *perfect hashing through brute-force* (Section 6) we describe techniques that still use brute-force at their core, but are practical because they first reduce the size of the sets that need to be handled with brute-force.

Perfect hashing through fingerprinting (Section 7) describes another technique, that is, focused more on construction and query performance than on space consumption. There are several implementations of the same basic technique, mostly with minor algorithmic changes.

In addition to the techniques themselves, we give a comparison of the asymptotic performance in Section 8. Figure 1 gives an overview of all approaches that we consider in this survey, and shows how the ideas influence each other. We can see that the three main categories each have their own strands of research. Recent approaches start using ideas from multiple categories, as we show in the figure by placing them in between the categories. Interestingly, the number of publications in the last few years increased massively, after several years of mostly inactivity in the early 2,000s. With so much new development in the field, this survey gives structure and an overview over the state of the art.

5 Perfect Hashing through Retrieval

We recall from Section 3.2 that a retrieval data structure represents a function $f : S \rightarrow \{0, 1\}^r$ that maps each key in $S \subseteq U$ to an r -bit integer. In the following, we explain a range of PHF constructions that use retrieval data structures (see Section 3.2). In Section 5.1, we explain PHFs that store the final output hash value in the retrieval data structure. This needs a space consumption of $\mathcal{O}(\log n)$ bits per key. Then, in Section 5.2, we explain approaches where the retrieval data structures only store the index of one of multiple choices, enabling constant space per key. Some of the approaches we present here do not have an explicit name in their original paper. We therefore follow the convention of previous papers and name the approaches after the first letters of the authors' last names.

5.1 Storing Perfect Hash Values

Retrieval data structures can be used to explicitly store a desired perfect hash value for each key using about $\log_2 n$ bits per key. Several early MPHf constructions actually implement a retrieval




Figure 1. Perfect hashing approaches and how they influence each other. While we describe them in this article, we do not evaluate the performance of approaches given in gray color because they are clearly larger than current ones or because they do not have a publicly available implementation.

Table 1. Overview of Perfect Hashing Through Retrieval

	Stored	Approach	Year	k	Graph Property	$ V $	Space Usage
Output value		CHM [45]	1992	2	Acyclic	$2.09n$	$2.09 \cdot n \log_2 n$
		MWHC [117]	1996	3	Peelable	$1.23n$	$1.23 \cdot n \log_2 n$
		BMZ [22]	2004	2	$ \text{Critical} < n/2$	$1.15n$	$1.15 \cdot n \log_2 n$
Hash function index		Bloomier filters [40]	2004	3	Peelable	$1.23n$	$2.46n^*$
		BPZ [26]	2007	3	Peelable	$1.23n$	$2.62n$
		GOV [74]	2016	3	Orientable	$1.10n$	$2.24n$
		WBPM [159]	2020	$\log_2 n$	Bipartite	n	$1.83n$
		SicHash [108]	2023	Mix 2, 4, and 8	Orientable	$(1 + \epsilon)n$	$2n$
		ShockHash [110]	2023	2	Orientable	n	$1.443n$
		MorphisHash [89]	2025	2	Orientable**	n	$1.443n$

We use the hypergraph interpretation here. The number k of nodes per edge corresponds to the number of choices for each key. * The original paper [40] discusses only PHFs, not MPHFs, and with weaker bounds. We report here the actual bounds for the hypergraph technique [117]. ** The retrieval data structure used cannot represent all orientable graphs, so some have to be skipped.

data structure and store the desired hash values [45, 117], even if the authors do not explicitly explain their approach as a retrieval data structure. In this section, we nevertheless look at those to illustrate the progress on perfect hashing made over the past years. We give an overview in Table 1. Remember from Section 3.2 that retrieval data structures can be constructed from the following system of equations: $w_{h(x,0)} + w_{h(x,1)} + \dots + w_{h(x,k-1)} = f(x)$ for all $x \in S$. All the PHF constructions in this section solve a system of that form. Remember also from Section 3.3 the

concepts of *peelability* (repeatedly taking away nodes of degree 1) and *orientability* (directing edges such that each node has indegree 1) of graphs.

CHM. Given a list of keys, Czech et al. [45] use peelability on a random graph to solve the system in Equation (1), where f maps each key to its rank in the list. If the graph is not peelable, which rarely happens below the peelability threshold, CHM retries with another set of hash functions. The construction is a general retrieval data structure, but the authors only use it to construct an MPHf. The space consumption is $2.09n \log_2 n$ bits. As we noted in Section 3.3, this article introduces the idea of peelability, that is, used in many later constructions.

MWHC. The authors of CHM then joined forces with Wormald [117] to extend the idea of CHM to 3-uniform hypergraphs. Once again, the construction is a general retrieval data structure, but the authors use it to construct an MPHf, lowering the space consumption to $1.23n \log_2 n$ bits. Belazzougui et al. [8] improve the peeling step of MWHC through better cache locality.

BMZ. Botelho et al. [22] enhance the idea of CHM into another direction than MWHC. In contrast, the graph does not need to be peelable—only orientable. Like CHM, the hash value is given by the sum of the values stored at the two ends of each edge. When assigning values to nodes, BMZ first uses the normal peeling process, that is, known from the previous approaches. When there are no more nodes of degree 1, we have arrived at the 2-core. If that 2-core consists of only nodes with degree 2, the remaining edges form loops. BMZ calls them *critical* and assigns them first using breadth-first-search. This ensures that different keys get different values, but BMZ no longer supports an arbitrary value for each key. As such, in contrast to CHM and MWHC, it is no longer a general-purpose retrieval data structure. Instead, it is a step towards smaller PHFs. After the critical edges, BMZ then assigns the remaining nodes similar to the algorithms above. At all times, it keeps a list of output values that are not assigned yet to find collision-free mapping. The fact that BMZ does not require peelability has the advantage that $|V|$ can be much closer to n . This reduces the amount of space needed to only $1.15n \log_2 n$ bits.

5.2 Multiple Choice Hashing

In the following, we present several ideas that we categorize as *multiple choice hashing*. In multiple choice hashing, each input key has a set of candidate output values. Here, these candidate values are determined by evaluating different hash functions $h(\cdot, i)$ on the input key. Let $i(x)$ be a function assigning a number to each key, indicating which of its candidate values should be selected in order to give a collision-free mapping. Then, $h(x, i(x))$ is unique for each key x , and only the hash function index $i(x)$ needs to be stored in a retrieval data structure. In contrast to Section 5.1, the approaches can therefore break the space consumption of $\Omega(\log n)$ bits per key. We illustrate this idea in Figure 2. Multiple choice hashing can also be interpreted as a hypergraph, where each key corresponds to a hyperedge, connecting its candidate output values. Then an orientation of this graph gives a mapping from keys to candidate values. In this interpretation, the similarities to the older approaches in the previous section become clear. In the following, we explain approaches from the literature in more detail. Table 1 gives an overview of the differences between the approaches. We want to mention the similarity of multiple choice hashing to *cuckoo hashing* [133], which is a collision resolution strategy in hash tables. Cuckoo hashing uses the *power of multiple choices* [123] to reduce the load of hash table cells.

Bloomier filters. Chazelle et al. [40] introduce the concept of *Bloomier filters*, which are thought of as a generalization of Bloom filter. The setting is slightly different, as they assume that each key is representable with r bits, and they create a retrieval data structure that maps each key to




Figure 2. Multiple choice hashing. Each key has the choice between several locations, in this case two. A retrieval data structure stores which of the choices to take.

itself. The purpose of the mapping is to create an approximate set membership data structure, but as in the previous case the construction can be used for general retrieval. In fact, they rediscover the idea of k -uniform hypergraph peeling introduced by Majewski et al. [117], but they use a completely different formalization, thus missing the connection, and not citing the previous results. Moreover, the peelability thresholds they provide are much weaker than the ones conjectured in [117]. However, they make an important observation that provides the first example of multiple-choice hashing: since each edge (i.e., key) is associated uniquely with the vertex it was peeled from, we can associate a unique vertex with each key by arbitrarily numbering the vertices in each edge, and storing the function i mapping each edge x to the index of the vertex $i(x)$ it was peeled from. This mapping just requires $\log k$ bits and is a PHF.

BPZ. Botelho et al. [24] rediscover the PHFs described in the Bloomier filter paper [40], but they do not cite it, likely because the different formalization. However, since their goal is minimal perfect hashing, they consider the problem of how to convert the Bloomier filter [40] idea to an MPHf. As we mentioned in the introduction, a standard technique involves a ranked bit vector: in this case, the vector would track the vertices that have been used to peel an edge, and thus are output by the PHF for some input key. The authors note, however, that in the optimal case of 3-uniform hypergraphs one uses two bits to store three values. They suggest to store 0 as such for vertices that are not output by the PHF, and to store 1, 2, or 3 for those which are. Then, after invoking the PHF, one has just to count how many nonzero pairs of bits appear before the output, which can be done in constant time and $o(n)$ space adapting standard ranking techniques [94]. In this way, they obtain a minimal PHF using just $2.62n$ bits [25].⁴

GOV. Genuzio et al. [74], reduce the space overhead of retrieval data structures and MPHfs. They exploit the fact that, as we already mentioned, the following thresholds are the same [53, 75]: (1) the threshold for orientability of k -uniform hypergraphs, (2) for solvability of random systems on \mathbf{F}_2 with k variables per equations, and (3) even for the same kind of random systems on larger fields. First they orient a random 3-uniform hypergraph, thus assigning a distinct vertex to each edge. Then they solve the system of equations on \mathbf{F}_3 representing the function that assigns to each edge the vertex selected by the orientation. To tame the cubic time of Gaussian elimination, GOV introduces a range of engineering tricks: first of all, keys are hashed into small partitions (see Section 3.1). The systems associated with each partition are first peeled as much as possible to reduce the number of equations. Then operations on equations use *broadword programming* [99] by packing multiple values into a single 64-bit word and running calculations on all of them at once. Finally, the paper introduces *lazy Gaussian elimination*, which is a heuristic for fast solution of

⁴Recall that the PHF can be stored in $2.46n$ bits but additional space is required for ranking.

random linear systems. Using three hash functions, and therefore two bits per key, and exploiting the same trick of Botelho et al. [24] to perform ranking on the output, they reduce the space consumption of minimal PHFs to $2.24n$ bits.⁵

WBPM. Weaver and Heule [159] use a larger number of $\mathcal{O}(\log n)$ candidate output values. When selecting which candidate to use for each key, they prefer choices with smaller indices. This is modeled as a weighted graph connecting each key to its candidates, where the edge of hash function i has weight i . WBPM then determines a minimum weight bipartite perfect matching (WBPM), which can be shown to reach a weight of $1.83n$. The selected edge indicates which hash function to use for each key. WBPM stores the hash function indices bit-by-bit in unary coding using a 1-bit retrieval data structure. The matching weight of $1.83n$ therefore also equals the space consumption of the final data structure, except for overheads like prefix sums due to partitioning.

SicHash. Most earlier approaches couple the task of retrieval and perfect hashing, which means that the approaches have to store the same number of bits for every input key. SicHash [108], in contrast, separates the two tasks of key placement and retrieval. This makes it possible to efficiently use more than one retrieval data structure and store values of different widths for different keys. Using the same width for every key, even with optimal retrieval, we can achieve no better than about 2.2 bits per key using four choices and remapping values to be minimal perfect. SicHash's main innovation is using a hash function to determine how many choices each key has. Some of the keys have two choices, some have four choices, and some keys have eight choices. As such, it stores a 1, 2, and 3-bit retrieval data structure, each handling a portion of the input keys. Using the space-efficient BuRR [57] retrieval data structure, this makes it possible to reduce the space consumption compared to previous approaches. SicHash achieves a favorable space-performance trade-off when being allowed 2–3 bits of space per key. It also achieves a rather limited gain in space efficiency by *overloading* the corresponding graph beyond the orientability threshold. Using small partitions, SicHash exploits the variance in the number of keys that can fit. Also, it uses the fact that the load factor at which construction likely fails converges to the load threshold from *above* as n grows.

ShockHash. ShockHash [110] can be seen as an extreme version of SicHash that directly tries to orient a graph with n nodes and n edges. Each key corresponds to one edge, connecting the two candidate output values of the key. The graph can be oriented if and only if it is a *pseudoforest*—a graph where each component contains as many edges as nodes. Because ShockHash works far above the orientability threshold, it needs to retry with many different graphs before it samples one, that is, orientable. ShockHash then stores the choice between the two candidate values of each key in a 1-bit retrieval data structure taking close to n bits. Additionally, it stores the hash function seed, which can be shown to need about $0.44n$ bits. This means that the majority of the data originates from a simple linear time orientation of the graph. Only $0.44n$ bits need to be determined by exponential time brute-force. Compared to a simple brute-force technique (see Section 6), which needs e^n tries in expectation, ShockHash needs only $(e/2)^n$ tries, while still reaching the space lower bound. We refer to Section 8 for more details on asymptotic performance. A key idea for making ShockHash practical is the introduction of a simple bit-parallel filter that checks whether all output values are hit by some key. If one output is not hit, the graph cannot be orientable. Because the construction is still exponential time, ShockHash first partitions the keys to smaller subsets using RecSplit (see Section 6.2). This makes the overall construction time linear in expectation.

⁵As in the case of BPZ, the additional $0.04n$ bits are used for the ranking structure.

ShockHash has configurations that achieve 1.52 bits per key using a construction time similar to RecSplit with 1.6 bits per key.

Bipartite ShockHash. *Bipartite ShockHash* [109] enhances the ShockHash idea. The approach stores two independent hash function seeds, one for each end of the edges. During construction, it builds a pool of hash function seeds and tests all combinations of seed pairs. By hashing each end of an edge to disjoint output ranges, the hash function pool can be filtered before building the pairs, which enables additional exponential speedups. Bipartite ShockHash-RS achieves a space consumption of 1.489 bits per key. In addition, bipartite ShockHash introduces a variant called ShockHash-Flat. The approach uses a k -PHF to partition the keys to smaller subsets instead of RecSplit. This leads to higher space consumption but much faster queries due to not having to traverse the splitting tree (see Section 6.2).

MorphisHash. ShockHash orients pseudoforests, so each component of the graph is a cycle with trees branching from it. The cycle can be oriented in two ways, which leads to about 1-bit of redundancy for every component. MorphisHash [89] constructs a special retrieval data structure that avoids this redundancy. The idea is to formulate orienting the graph as a system of linear equations indicating that each node needs an indegree of 1 (mod 2). The solution to the equation system is then a retrieval data structure that can return the orientation of each edge. By reducing the number of columns of the equation system, MorphisHash gives a trade-off between space consumption and success probability. Because it stores many small retrieval data structures next to the hash function seeds, MorphisHash has better cache locality during queries than ShockHash.

6 Perfect Hashing through Brute-Force

The fact that brute-force can be useful for constructing PHFs is mentioned early after the discovery of perfect hashing [42]. Today, most of the approaches achieving the best space efficiency are based on brute-force techniques. In the following, we first describe a very simple brute-force construction and an approach that uses SAT solving for more structured brute-force search. We then illustrate the practical approaches: perfect hashing through bucket placement in Section 6.1, and recursive splitting in Section 6.2.

A Simple Brute-Force Construction. The idea of the simple brute-force construction is to try random hash functions (identified by different seeds) until one happens to be minimal perfect [122]. The data structure simply needs to store the found seed in binary coding. We can expect to try $n^n/n! \approx e^n$ different hash function seeds because there are n^n functions, from which $n!$ are minimal perfect [62, 122]. This means that it reaches the space lower bound of minimal perfect hashing, but because of its exponential running time, it is not practical for large n . Still, brute-force is used as a building block on smaller sets in a range of practical constructions, which we explain below.

SAT Solving. Weaver and Heule [159] describes a solution based on SAT solving. The data structure consists of a simple sequence of bits, and a hash function on each key determines which of these bits (possibly flipped) should be concatenated to give the final hash function value. To determine the values of the bits in the sequence, the algorithm searches for an assignment using SAT solving. The SAT encoding is based on an *all-different* constraint and is feasible for $n \leq 40$. The SAT solver can detect correlations between the bits in the data structure to prune the search space more efficiently than brute-force. Still, due to its large search space of all possible functions, we categorize it as brute-force. While the construction using SAT encoding achieves space consumption very close to the lower bound, it is much slower in practice than other approaches.




Figure 3. Illustration of perfect hashing through bucket placement, where the $n = 6$ keys are first mapped to buckets and then, for each bucket i , a seed value v_i is computed to place the keys in the bucket without collisions to the output domain $[m = 8]$. Here, h has range $[m]$.




Figure 4. Relative sizes of different buckets when using various bucket mapping functions. The legend gives the original paper even if later techniques use the same functions. Adapted from [90].

Table 2. Overview of Perfect Hashing Approaches using Bucket Placement

Approach	Year	Bucket Assignment	Seed Compression	Search Type	Partitions
FCH [66]	1992	Skewed	Fixed width	Additive	None
CHD [11]	2009	Uniform	SDC	Double Additive	None
PtrHash [139]	2021	Skewed	Multiple options	XOR	Large [140]
PHOBIC [90]	2024	Optimized	Interleaved	Additive+Retry	Small
PtrHash [83]	2025	Approximate optimized	Byte array	Cuckoo	Medium
PHast [16]	2025	Uniform	Fixed width	Windowed	None

6.1 Bucket Placement

Perfect hashing through bucket placement is achieved via a three-step algorithm: *mapping*, *ordering*, and *searching* (MOS) [11, 42, 65–68, 132, 148]. The general idea is as follows. (1) *Mapping*. Each key x is mapped to a small bucket $B_i = \{x | b(x) = i\}$, using a *bucket assignment function* $b : U \rightarrow [t]$ and t buckets. (2) *Ordering*. The buckets $\{B_i\}_{i=0}^{t-1}$ are then sorted; usually by falling size. (3) *Searching*. For each bucket B_i in the order, we determine a value v_i , using brute-force, that the final function f uses to place all the keys in the bucket to the output domain $[n]$ *without collisions* with keys of previously processed buckets and with one another. Once the value v_i has been determined, we say that the bucket B_i has been “placed”. (4) *Encoding*. The collection of values $\{v_i\}_{i=0}^{t-1}$ is then stored in an array $V[i] = v_i$ and compressed with a mechanism that retains the ability to quickly retrieve $V[i]$ for any random $i \in [t]$. The array V is what the MPH data structure actually stores.

At the beginning of the search, almost all values in the output domain $[m]$ are available, hence the first buckets are significantly easier to place. Buckets with fewer keys are easier to place as well. This is intuitively why, to accelerate the search, it is useful to sort the buckets by falling size and process them in this order. A query for a key x then only needs to retrieve $V[i]$, where $i = b(x)$, and use it to compute the resulting hash value. Figure 3 presents an illustration of bucket placement and Table 2 outlines the approaches discussed in this section.

FCH. FCH [66] is the first PHF construction that uses the bucket placement idea. It uses $t = \lceil cn / (1 + \log_2 n) \rceil$ buckets, where $c > \log_2 e$ is a parameter that trades space usage for construction speed. To further amplify the effect that the first buckets are easier to place, FCH uses a bucket assignment function that produces a *skewed* distribution of bucket sizes: it hashes 60% of the keys

to 30% of the buckets. The information v_i used for placing a bucket consists of two parts, d_i and s_i . For each bucket B_i , a displacement value $d_i \in [n]$ is determined so that all the keys x in B_i are placed and assigned hash values $(h(x, s_i) + d_i) \bmod n$. In the following (and in Table 2), we refer to this form of hash function as *additive displacement*. The seed $s_i \in \{0, 1\}$ gives more flexibility during the search for bucket B_i : if all of the displacements cause collisions for $s_i = 0$, FCH repeats the search for $s_i = 1$. If both seeds do not work, the construction fails.

To accelerate the identification of d_i , an auxiliary data structure formed by two integer arrays, taking $\Theta(n \log n)$ bits overall, is used to maintain the output values available in $[n]$ during the search. Given a random available output value $p \in [n]$, d_i is computed by “aligning” an arbitrary key of B_i , say x_i , to p : let $q = h(x_i, s_i) \bmod n$, then $d_i = (p - q) \bmod n$. This random alignment is rather critical for the search as it guarantees that all output values have the same probability to be occupied. On the other hand, this auxiliary structure has to be updated as well during the search, involving some extra cache misses.

Since each displacement d_i can be coded in $\lceil \log_2 n \rceil$ bits, all the pairs $\{(s_i, d_i)\}$ are stored in an array of $t(1 + \lceil \log_2 n \rceil)$ bits. It follows that FCH uses about c bits/key. By decreasing the number of buckets, it is possible to lower the space usage at the cost of a larger construction time. However, in practice it is unfeasible to go below 2.5 bits/key for large values of n . On the other hand, FCH shines for its very fast query time because the pair (s_i, d_i) is coded in fixed-width. That is, FCH just spends one memory access per query (plus some inexpensive bit manipulations), which is optimal.

CHD. CHD [11] maps the keys uniformly to $t = \lceil n/\lambda \rceil$ buckets, where $\lambda > 0$ is a chosen parameter representing the average bucket size. For each bucket B_i , a displacement *pair* $(d_{i,0}, d_{i,1})$ is computed to ensure that all keys x within the bucket can be placed collision-free to output values $(h(x, 0) + d_{i,0}, h(x, 1) + d_{i,1}) \bmod m$. CHD chooses for the search an output range larger than n , hence the values are in $[m]$, where $m = \lceil n/\alpha \rceil$ for a suitable load factor $0 < \alpha < 1$. The domain of the function is then restricted back to $[n]$ using a rank query as we discussed in Section 3.1. Differently from FCH, CHD maintains a bitmap of m bits to mark values in $[m]$ that are still available.

Instead of explicitly storing a displacement pair $(d_{i,0}, d_{i,1})$ for each bucket, CHD encodes the position of the pair in the predefined sequence $(0, 0), \dots, (0, m-1), \dots, (m-1, 0), \dots, (m-1, m-1)$. Very importantly, CHD tries positions in increasing order starting from the smallest position, i.e., following the order $[0, 1, 2, \dots]$, and not at random as FCH. This intuitively helps compression because smaller numbers are tried first. The sequence of pairs’ positions is then compressed within entropy bounds using the Simple Dense Coding (SDC) mechanism by Fredriksson and Nikitin [71], while permitting to retrieve an index in constant time. With $\alpha = 0.99$ and $\lambda = 6$, CHD achieves 2 bits/key and it is much faster to build compared to FCH (albeit significantly slower to query).

PTHash. Under proper configuration, PTHash [139] combines the query time of FCH and the space effectiveness of CHD, with fast construction time. It uses the same imbalance trick of FCH by hashing 60% of the keys to 30% of $t = \lceil cn/\log_2 n \rceil$ buckets. Like CHD, it stores the values $\{v_i\}$ in compressed form but, rather than coding them using SDC, PTHash generalizes the encoding step to use a wide range of mechanisms (e.g., fixed-width, dictionary-based, Elias-Fano [60, 64], and others). Using an appropriate coding scheme, only a single memory access is required to find the hash value for a key and the remaining operations use simple arithmetic.

Since the first processed buckets are much easier to place, the values $\{v_i\}$ for the first 30% of the buckets have a much smaller entropy than those for the other buckets. The *front* of the array V , of length $\approx 0.3t$, can thus be compressed better than the remaining *back* part. This allows PTHash to tune two different compressors for the front and the back parts, respectively: a light-weight compressor for the front and a more-succinct (but slower to query) compressor for the back part.

Like CHD, it also first generates a non-minimal function with range $m = \lceil n/\alpha \rceil$ for some load factor $0 < \alpha < 1$, and uses a bitmap of m bits to mark taken output values. It then scales the output range back to $[n]$ by remapping values larger than n using an explicit list of $m - n$ free output values that can be effectively compressed using Elias-Fano.

To accelerate the search for the values $\{v_i\}$, PTHash uses the so-called *XOR displacement*: the keys of a bucket are mapped to output values given by $(h(x, s) \oplus h(v_i, s)) \bmod m$, where $(A \oplus B)$ denotes the bit-wise XOR between the integers A and B , and s is a randomly chosen seed. As also noted for CHD, the values $\{v_i\}$ are tried in linear order to improve compression, while still achieving random displacement of keys via the XOR operator. Compared to FCH, the table of random displacements is avoided altogether. Compared to CHD, instead, long clusters of already occupied slots are skipped. Also note that the seed s does not change during construction, hence all the keys can be hashed only once instead of for every seed (which is particularly important to save time when hashing long, variable-length, keys). On the other hand, this XOR-based approach, in combination with mod, only works for large output domains [90].

PTHash-HEM [140] is a parallel construction of PTHash that first partitions the input and then builds each partition (roughly holding a few million keys) independently in parallel.

PHOBIC. PHOBIC [90] is based on PTHash. It introduces an optimal bucket assignment function $b(\cdot)$ that, instead of generating two expected bucket sizes as in FCH, gives a different expected size to *every* bucket. This allows each bucket to have the same probability to be placed successfully, which minimizes the construction time (in asymptotic terms) and makes the expected value of each v_i equal. The optimal bucket assignment function practically accelerates the construction compared to PTHash for large values of λ . Figure 4 plots the relative expected bucket sizes using different bucket assignment functions. As apparent from the plot, PHOBIC makes the first (and easiest to place) buckets much larger than those generated by FCH.

PHOBIC uses partitioning as a key design feature. In particular, it uses small partitions (of about 2,000–3,000 keys), differently from PTHash-HEM. Small partitions lead to better cache locality since the bitmap used by the search is small and accesses are more localized in memory. Small partitions also enable efficient GPU acceleration. Since small partitions are created, it might be impossible to place keys without collisions using the XOR-based search. PHOBIC therefore relies on an additive displacement function and computes $v_i = s_i \cdot P + d_i$ where s_i is a seed for the bucket B_i , P is the size of the partition, and $d_i \in [P]$ is the displacement value. The hash value of a key x , relative to the partition size P , is then $(h(x, s_i) + d_i) \bmod P$. This is similar to the function computed by FCH, but the seed s_i is not necessarily restricted in $\{0, 1\}$: if no displacement $d_i \in [P]$ works for the bucket B_i , the seed s_i is incremented (and all the keys of the bucket hashed again), and the search for a new d_i continues. The pair (s_i, d_i) is computed from v_i as follows: $s_i = \lfloor v_i/P \rfloor$ because $d_i < P$, and clearly $d_i = v_i \bmod P$. Another contribution of PHOBIC is *interleaved coding* which exploits that the values of the j th buckets of all partitions have approximately the same distribution. This is a natural generalization of the *front-back* compression strategy of PTHash and gives PHOBIC a slight space advantage over PTHash without significantly compromising query time.

PtrHash. PtrHash [83] builds on both PTHash and PHOBIC with the intent of simplifying the design of the data structure to accelerate queries. The two main simplifications are: (1) a value v_i is coded using exactly 1-byte, hence PtrHash stores a plain byte array at its core; (2) the allocation of the same number of output slots for each partition. The values stored in PtrHash are therefore less than 256 by design. If none of the values $\{0, \dots, 255\}$ is able to place a bucket, PtrHash bumps out already placed buckets in a similar way to cuckoo hashing [133]. Like PHOBIC, PtrHash partitions the input into small partitions and uses not only the same number of buckets per partition but also the same number of output slots. This has the net advantage that, during a query for key x , the




Figure 5. Illustration of the overall RecSplit data structure. Within each bucket, it constructs a splitting tree. Circular nodes represent splittings, squares represent bijections.

Table 3. Overview of Approaches using Recursive Splitting

Approach	Year	ℓ	Base case	Seed encoding	Fanout
RecSplit [62]	2020	8–16	Brute-force	Golomb-Rice	Depending on ℓ
SIMDRecSplit [19]	2023	8–18	Rotation fitting	Golomb-Rice	Depending on ℓ
ShockHash-RS [110]	2023	40–60	ShockHash	Golomb-Rice	Depending on ℓ
Bip. ShockHash-RS [109]	2023	90–120	Bip. ShockHash	Golomb-Rice	Depending on ℓ
CONSENSUS-RS [104]	2025	1	None, $\ell = 1$	CONSENSUS	2

The base case column gives the technique used after splitting the input set to small subsets of size ℓ . We also include ShockHash [109, 110] again, even though it is actually explained in Section 5.2, because it uses recursive splitting to partition the keys.

size of the partition to which x belongs to does not have to be retrieved from an encoded sequence. PtrHash further uses Taylor approximations of the optimal bucket assignment function introduced in PHOBIC, which are more efficient to compute. Lastly, it supports explicit prefetching of batches of keys to increase query throughput.

PHast. Perfect hashing made fast (PHast) [16] simplifies several aspects of the bucket placement idea for fast queries. It uses a uniform bucket assignment function, fixed width coding of seeds, and no partitioning. This is possible due to a new search strategy for seeds. Each bucket only has a small window of possible output values, which overlaps with the windows of other buckets. Instead of taking the first collision-free seed, PHast evaluates all seeds of a bucket and takes the one giving the smallest hash values among the collision-free seeds. The order in which to place the buckets is determined by a combination of both bucket index (lowest first) and size (largest first). Therefore, taking the seed with the smallest hash values leaves more flexibility for later buckets. The few buckets that cannot be placed with any of their available seeds are bumped to a more expensive fallback data structure. PHast⁺ designates variants that use additive placement functions. This accelerates queries and enables bit-parallel seed search, actually making PHast as fast as, or faster than, non-brute-force techniques.

6.2 Recursive Splitting

The simple brute-force construction explained at the beginning of this section needs e^n trials and therefore is not practical for large n . The recursive splitting idea introduced in RecSplit [62] enables space-efficiently reducing the input set into small sets of size ℓ . These are then small enough that the brute-force construction becomes feasible. The following paragraphs introduce RecSplit, as well as improvements in later papers. Table 3 gives an overview.

RecSplit. RecSplit [62] uses brute-force in a novel way to enable large key sets with space close to the lower bound. It first maps all keys to buckets of expected size b , where b is a tuning parameter usually in the range 100–2,000. In each bucket, it then constructs an independent *splitting tree*, as illustrated in Figure 5. The splitting tree partitions the keys into smaller and smaller sets until

we arrive at the leaves which have a small configurable size ℓ . At each inner node, RecSplit tries random hash functions to find one that distributes the keys to the child nodes according to a *splitting strategy*. The splitting strategy also decides the *fanout* (number of child nodes) of the tree and is optimized for balancing the work between splittings and leaves. If the number of keys in a node is not a multiple of the child node sizes, the last node might be smaller. RecSplit adjusts the hash function accordingly. The lowest level of the splitting tree is called *leaf level*. Each leaf, except for possibly the last, contains exactly ℓ keys. Usual values for the leaf size are about 8–16 keys. This is small enough that it is feasible to use the simple brute-force construction. The shape of the splitting tree depends only on the leaf size ℓ and the total number of keys in the bucket. Therefore, RecSplit does not need to store structural information of the tree. It is enough to store the seed for each node in DFS order, encoded using Golomb-Rice codes [76, 146]. The key theoretical observation of RecSplit is that the probability of finding an MPHf by brute force is the same as the probability of finding splittings and MPHfs for the leaves of the tree, so the space needed to represent the MPHf does not change with splitting (modulo small overheads due to approximations), but splitting makes the construction much faster. Overall, the overhead over the lower bound is constant for each node of the splitting tree.

RecSplit can be queried by traversing the splitting tree from the root to a leaf by applying the splitting hash functions. During traversal, it accumulates the number of keys stored in children to the left of the one descended into. The final hash value is then the sum of the value of leaf bijection, the number of keys to the left in the splitting tree, and the total size of previous buckets. Due to having to traverse the tree, queries are relatively slow compared to bucket placement techniques. There are configurations that need only 1.56 bits per key with reasonable construction time.

SIMDRecSplit and GPURecSplit. SIMDRecSplit [19] enhances RecSplit to use parallelism on multiple levels to improve the construction throughput. The paper proposes a new technique for searching for bijections called *rotation fitting*. Instead of just applying hash functions on the keys in a leaf directly, rotation fitting splits the keys into two sets using a 1-bit hash function. It then hashes each of the two sets individually, forming two words where the bits indicate which hash values are occupied. Then it tries to cyclically rotate the second word, such that the empty output values left by the first set are filled by the values of the second set. Each rotation essentially gives a new chance for a bijection [19], so it is a way to quickly evaluate additional hash function seeds in a bit-parallel way. In addition to bit parallelism, SIMDRecSplit uses parallelism on the level of words, through SIMD instructions. For this, it tests different hash function seeds or rotations simultaneously. Finally, it uses multi-threading to construct different buckets in parallel.

GPURecSplit [19] is an implementation of SIMDRecSplit on the GPU. An important observation here is that all buckets of the same size also lead to splitting trees of the same shape. Therefore, they can be constructed together using the same set of kernel calls. GPURecSplit then uses CUDA *streams* to construct different *tree shapes* in parallel.

CONSENSUS-RecSplit. Perfect hashing through brute-force often stores the seed for each subtask individually using a variable-length code like Golomb-Rice [76, 146]. With CONSENSUS [104], each seed has a fixed number of choices just barely larger than its expected number of choices. If a subtask runs out of choices without finding a successful seed, CONSENSUS backtracks to a previous subtask and looks for another successful seed. The key idea making this feasible is that hashing takes into consideration not just the current seed but the *concatenation* of all previous seeds. With this, backtracking gives a completely new chance for a seed to be successful. CONSENSUS needs less space than storing the seed of each subtask individually, even when assuming an entropy-optimal coding. CONSENSUS-RecSplit uses RecSplit with a leaf size $\ell = 1$ which would otherwise be inefficient. However, it stores the seeds efficiently using CONSENSUS. Both RecSplit and CONSENSUS-RecSplit




Figure 6. Illustration of PHFs using fingerprinting. Collisions are handled in the next layer.

Table 4. Overview of Perfect Hashing Through Fingerprinting

Approach	Year	New idea
Meraculous [39]	2011	Original idea
FiPHa [126]	2014	First description as data structure
BBHash [114]	2017	First public implementation
FMPH [15]	2023	Partitioning for faster construction
FMPHGO [15]	2023	Brute-Force trials for fewer levels
FiPS [106]	2024	Sorting for faster construction, interleaved rank data structure

can achieve a space consumption of $(1 + \epsilon) \log_2 e$ bits per key, arbitrarily close to optimal. However, while RecSplit has a running time of $\mathcal{O}(n \cdot e^{1/\epsilon})$, CONSENSUS-RecSplit is the first approach to achieve this in just $\mathcal{O}(n/\epsilon)$. We refer to Section 8 for details. A practical implementation achieves a space consumption of just 1.444 bits per key, very close to the lower bound of $\log_2 e \approx 1.443$ bits per key. However, the queries are about 60% slower than RecSplit.

7 Perfect Hashing through Fingerprinting

The idea of perfect hashing through fingerprinting is to assign a small ($\Theta(n)$) fingerprint to each input key using a hash function. We then resolve collisions between the fingerprints using recursion on the colliding keys. An advantage is the very simple and easily parallelizable construction. It is originally introduced by Chapman et al. [39] in the context of bioinformatics, but not described as a perfect hashing data structure of general interest. Müller et al. [126] then enhance and describe the idea from a data structure perspective. Table 4 gives an overview of the development.

FiPHa. Perfect hashing through fingerprinting [126] hashes the n keys to γn fingerprints using an ordinary hash function, where $\gamma \geq 1$ is a tuning parameter. A bit vector of length γn indicates fingerprints to which exactly one key was mapped. At query time, when a key is the only one mapping to its location, a rank operation on the bit vector gives the MPHF value. The bit vector indicates with a 0-bit that the key was not the only one mapping to that location. In this case, an additional layer of the same data structure needs to be queried. Figure 6 illustrates this idea. The most space-efficient choice $\gamma = 1$ leads to a space consumption of e bits per key [126], which is quite far from the lower bound of $\log_2 e$ bits per key. In this case, a query traverses e layers of the data structure in expectation. FiPHa was developed in cooperation with the German company SAP, and the source code is not publicly available.

BBHash. The first publicly available implementation, BBHash [114], iterates over all keys to count the collisions. Then it iterates over the keys again to write the fingerprint bit vector and to extract colliding keys. During construction, this causes up to three random memory accesses per key and level, while we will later see that FiPS [106] achieves much higher cache locality. For parallelization, BBHash uses a large number of atomic operations in the arrays. While FiPHa already introduces

the idea to scale the bit vector of fingerprints, BBHash makes this more explicit by introducing the γ parameter that we already used in the description of FiPHa.

FMPH. FMPH [15] is a fast implementation of the idea in the Rust programming language. The approach still uses the construction algorithm with two passes. For parallelization, FMPH distributes the keys to multiple threads and relies on a large number of atomic operations, just like BBHash. FMPH achieves impressive speedups compared to BBHash, even though it essentially implements the same algorithm.

FiPS. FiPS [106]—Fingerprint Perfect hashing through Sorting—is a third implementation of the approach. In contrast to the approaches above, FiPS focuses on cache locality when filling the bit vector during construction. It does so by sorting the fingerprints and then determining collisions by a simple scan. Using integer sorting, the construction takes linear time. This approach through sorting is already described in the FiPHa paper [126] but without a publicly available implementation. Using existing sorting libraries, the construction can be parallelized efficiently. FiPS interleaves the select data structures and the bit vector. More precisely, in each cache line, FiPS stores both the bit vector indicating fingerprints that did not collide, and the number of bits set in previous cache lines. Different cache line sizes provide different trade-offs between space consumption and query performance.

FMPHGO. FMPH with group optimization (FMPHGO) [15] is a new spin on the fingerprinting idea, combining it with a few brute-force tries. The idea is to hash the keys to buckets like in FCH or CHD. FMPHGO then tries a (small) number of different hash functions for each bucket and invests additional space to store which hash function should be used. It selects the hash function that causes the least collisions in the fingerprint array. Overall, this reduces the number of required layers and the required space. Compared to FMPH, FMPHGO reduces the storage space by up to 0.7 bits/key, while the query performance stays mostly the same.

8 Asymptotic Performance

In this section, we discuss the asymptotic space consumption, construction time, and query time of different families of approaches. This includes how the configuration parameters can be tuned to achieve a space consumption that asymptotically approaches the lower bound.

We say a PHF has space overhead ε if it uses $(1 + \varepsilon)n \log_2 e + o(n)$ bits of space (see Section 1). It is called *succinct* if $\varepsilon \in o(1)$. Many of the approaches we describe in this survey admit succinct configurations, such as PTHash [139, 140], PHOBIC [90], RecSplit [19, 62], and ShockHash [109, 110]. Here, we go one step further and discuss the $o(1)$ term in more detail. We are specifically interested in approaches with tuning parameters that can bring ε arbitrarily close to 0 without the construction or query time increasing too much. In contrast, earlier papers mostly model the tuning parameters as constants. In Table 5, we give an overview of the asymptotic construction time depending on the parameters. Note that all approaches can trivially get linear time construction through partitioning (see Section 3.1). However, the partitions introduce some space overhead. If we want to get arbitrarily close to the lower bound, we might have to select larger partition sizes, leading to construction and query times that depend on the parameter. We only use partitioning where used in the original paper. We start with brute-force approaches that are only suitable for small input sets of size ℓ .

Simple Brute-Force. The simple brute-force technique (see Section 6) takes time $\mathcal{O}(e^\ell)$ to construct. The technique reaches the space lower bound but encoding the result loses a constant number of bits. Therefore, in order to decrease the space overhead, we have to increase ℓ . Queries take constant time assuming that the seed of $\approx \ell \log_2 e$ bits fits into a machine word of size w .

Table 5. Asymptotic Construction and Query Time of Different Approaches

Approach	Space overhead	Construction time	Query time
Simple Brute-Force [122]	$\mathcal{O}\left(\frac{1}{\ell}\right)$	$\mathcal{O}(e^\ell)$	$\mathcal{O}(1)$
ShockHash [110]	$\mathcal{O}\left(\frac{1}{\ell}\right)$	$\mathcal{O}(1.36^\ell)$	$\mathcal{O}(1)$
Bip. ShockHash [109]	$\mathcal{O}\left(\frac{1}{\ell}\right)$	$\mathcal{O}(1.17^\ell)$	$\mathcal{O}(1)$
Multiple Choice [24, 74, 108, 159]	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(1)$
Bucket Placement [11, 66, 90, 139, 140]	$\mathcal{O}\left(\frac{1}{\lambda}\right)$	$\mathcal{O}\left(n \frac{e^\lambda}{\lambda}\right)$	$\mathcal{O}(1)$
Recursive Splitting [19, 62] + any brute-force base case	$\mathcal{O}\left(\frac{1}{\ell} + \frac{\log b}{b}\right)$	$n(e^{\alpha\ell} + \text{poly}(b))$	$\mathcal{O}\left(\log \frac{b}{\ell}\right)$
CONSENSUS-RecSplit [104]	$\mathcal{O}(\varepsilon)$	$\mathcal{O}\left(\frac{n}{\varepsilon}\right)$	$\mathcal{O}\left(\log \frac{1}{\varepsilon}\right)$
Fingerprinting [15, 39, 106, 114, 126]	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(1)$
Hagerup and Tholey [85]	$\mathcal{O}\left(\frac{(\log \log n)^2}{\log n}\right)$	$\mathcal{O}(n)$	$\mathcal{O}(1)$

The variable ε denotes the space overhead over the lower bound. While almost all approaches have some tuning parameters, not all make it possible to get arbitrarily close to the lower bound. With w , we denote the size of a machine word. Top: Brute-force approaches for small sets of size ℓ , assuming $\ell \in \mathcal{O}(w)$. Bottom: Approaches for large sets of size n , assuming $\ell, b, \lambda \in \mathcal{O}(w)$, $\varepsilon \in \omega(1/\sqrt{n})$.

ShockHash. ShockHash [110] (see Section 5.2) is a hybrid between brute-force and multiple choice hashing. Like brute-force, its construction time is exponential in ℓ , however, with a much smaller base. Bipartite ShockHash [109] reduces the base even further. Queries of ShockHash access a seed of $\approx \ell(\log_2(e) - 1)$ bits and a 1-bit retrieval data structure. Both operations take constant time, assuming $\ell \in \mathcal{O}(w)$. While ShockHash uses the BuRR [57] retrieval data structure (see Section 3.2) in practice, we use a different construction [56] here because its space overhead shrinks more quickly. In the following, we now consider techniques for larger input sets.

Multiple Choice Hashing. Multiple choice hashing (see Section 5.2) needs a constant number of bits/key due to storing the index of the choice. Using a smart choice of hash function indices we can bring this down to 1.83 bits/key [159]. However, no matter how we tune the number of choices for each key, classical multiple choice hashing does not reach the lower bound of $\log_2 e$ bits/key.

Bucket Placement. PHOBIC [90] shows that perfect hashing through bucket placement (see Section 6.1) achieves a space consumption of $\log_2(e) + \mathcal{O}(\log(\lambda)/\lambda)$ bits per key. Therefore, by selecting large average bucket sizes λ , we can get arbitrarily close to the optimal space consumption. The construction time stays linear in n , but it increases exponentially in λ . Depending on the bucket assignment function and the order in which the buckets are placed, it is possible to achieve a construction time of $\mathcal{O}(ne^\lambda/\lambda)$. The query time is constant. Note that PtrHash [83] uses fixed size seeds for improved query time in practice, which makes it not succinct. For a similar reason, PHast [16] likely is not succinct either. However, its space consumption as a function of λ is an open problem.

Recursive Splitting. Each split in RecSplit's splitting tree loses a constant number of bits [62]. The original implementation uses the simple brute-force technique as its base case in each leaf. Then, the space overhead in each leaf of size ℓ is in $\mathcal{O}(1/\ell)$. For buckets of size b , storing the prefix sum of bucket sizes causes an additional overhead of $\mathcal{O}(\log(b)/b)$ bits. Queries traverse the splitting tree, which takes at least time $\mathcal{O}(\log(b/\ell))$, and then evaluate the base case in constant time. Note that this assumes $\ell, b \in \mathcal{O}(w)$. Constructing the splittings takes polynomial time in the bucket size, in addition to the base case construction. Note that now ℓ is a tuning parameter. This means that for any constant factor α , we get a space overhead of $\mathcal{O}(\alpha/\ell) = \mathcal{O}(1/\ell)$. However, the construction time

of the base case is $\mathcal{O}(e^{\alpha\ell}) = \mathcal{O}((e^\alpha)^\ell)$. Therefore, any exponential base case (such as ShockHash) only gives a constant factor of space improvement when included in RecSplit.

CONSENSUS-RecSplit. Bucketed CONSENSUS-RecSplit [104] is currently the only known approach where the construction time increases linearly in the inverse of the space overhead. This is because its space consumption is not influenced by an exponential-time base case. Through buckets of size $\mathcal{O}(1/\epsilon)$, it achieves a query time of $\mathcal{O}(\log(1/\epsilon))$. Consider the approaches using the parameter ℓ (or λ). Assume the word size is w , e.g., $w \in \Theta(\log n)$. If we wanted to reduce the space overhead to $o(1/w)$ by using $\ell \in \omega(w)$, then not only would the construction time become a problem (after all, it is exponential in ℓ), but the seed values of $\Theta(\ell)$ bits would also no longer fit the word size. This means that the query time would start to increase linearly in ℓ . CONSENSUS-RecSplit is currently the only approach for which this is not the case.

Fingerprinting. The fingerprint based approaches (see Section 7) have a tuning parameter γ that influences the space consumption. However, when selecting the smallest possible value of $\gamma = 1$, we get a space consumption of e bits per key, quite far from the lower bound of $\log_2 e$ bits per key. FMPHGO [15] introduces the idea of group optimization to improve the space consumption of fingerprinting. In a hypothetical case where we would encode the group seeds with variable length and retry until there are no collisions, we get perfect hashing through bucket placement. FMPHGO can therefore interpolate between Fingerprinting and bucket placement.

Hagerup and Tholey. Finally, we look at a purely theoretical result, that is, not included in the main part of the survey. Hagerup and Tholey [85] first reduce the range of input universe to $[n^3]$ by repeatedly taking the keys modulo a prime number. Then they hash the input to small sets of size $\mathcal{O}(\log n / \log \log n)$. For each possible small input set, they use brute-force to determine a minimal PHF. Then they map each actual small set to one of the pre-computed MPHFs using a lookup table. Overall, Hagerup and Tholey achieve linear construction time, constant query time, and space $1 + o(1)$ times the lower bound. More precisely, the space overhead is $\mathcal{O}((\log \log n)^2 / \log n)$ and therefore only depends on n . This means that in order to reduce the overhead, we have to increase the input size. In contrast, other approaches have a tuning parameter that can control the space overhead independently of n . Note that even for $n = 2^{256}$ we have $(\log \log n)^2 / \log n > 1$, so the overhead is large even for astronomic sizes of n . The approach has never been implemented (and is, in fact, not properly defined for $n < 2^{150}$ [25]).

9 Evaluation

In this section, we compare the performance of the different state-of-the-art PHF constructions. Our detailed measurements can help to pick the most fitting PHF for any given application. For this comparison, we look at the three most important parameters for minimal PHFs: space consumption, construction time, and query time.

Experimental Setup. We run most of our experiments on an Intel i7 11700 processor with 8 cores (16 hardware threads (HT)), pinned to its base clock speed of 2.5 GHz, and supporting AVX-512. In our experiments, pinning the clock speed does not influence the relative performance of the approaches. The machine runs Rocky Linux 9.5 with Linux 5.14.0. The sizes of the L1 and L2 data caches are 48 KiB and 512 KiB per core, and the L3 cache has a size of 16 MiB. We use the GNU C++ compiler version 14.2.0 with optimization flags `-O3 -march=native` and Rust 1.85.0 with release mode and `target-cpu=native`. We also give experiments on an ARM processor in the extended version [107]. Many of the implementations we evaluate are actively maintained and get updated with performance improvements. Our measurements refer to their state from October 12, 2025. As input data, we use strings of uniform random length in the range [10..50] containing random

characters except for the zero byte. This is because essentially all approaches support string inputs. Note that, as a first step, almost all competitors generate an initial hash of each key using a high quality hash function. This makes the remaining computation largely independent of the input distribution. The code and scripts needed to reproduce our experiments are available on GitHub under the General Public License [105]. In total, we measure 9290 data points with a cumulative duration of over 32 days.

Dominance Maps. Visual plots give a good understanding of the trade-offs between different PHFs, especially if each approach has a wide number of configurations. In order to print all three main parameters—space consumption, construction performance, and query performance—on paper, we use a type of plot that we call *dominance map* [57, 106]. We start with a 2D projection and then color each point based on the best competitor along the third dimension. This therefore provides a “front view” of the Pareto space. Pareto optimal points are the points that are not dominated in all three dimensions simultaneously. This results in a cleaner plot because it does not include data points for dominated approaches. When having a specific time and/or space budget, dominance maps show the fastest approach given that restriction.

Overview. We start our evaluation with Pareto fronts and dominance maps containing a wide range of configurations for each competitor. In Sections 9.1 and 9.2 we focus on the construction time trade-off and the query time trade-off. From these configurations, we pick representative ones for each approach, details in the extended version of this article [107]. We then look at how the approaches scale in the input size in Section 9.3 and in the number of threads in Section 9.4.

9.1 Construction Trade-Off

We start our comparison by looking at the single-threaded construction time. Figure 7 gives both a Pareto front showing all approaches, and the dominance map described above. The dominance map plots the space consumption and the construction throughput on the axes, and colors each point with the approach that has the fastest queries. For example, with a space consumption of 4 bits per key, FiPS [106] has the fastest construction. However, by sacrificing just a little bit of construction performance by looking further down, we see that we can get the faster queries of PtrHash [83]. Because some approaches are focused on very small space consumption, we additionally give a plot that uses logarithmic x -axes and y -axes. In the following, we look at the different PHF constructions in more detail. We focus on the construction time but also briefly mention the query time from the dominance maps. For details on query time, we refer to Section 9.2.

Multiple Choice Hashing. SicHash [108] is up to two times faster to construct than PTHash [139] and PHOBIC [90], up to a space consumption of about 2 bits per key. However, like most other approaches, its construction throughput stays quite far below SIMDRecSplit. Even though BPZ [25] is a lot older than the other approaches, it still holds up in terms of construction time. However, the dominance map shows that newer approaches cover it almost completely with a better trade-off. Looking at even more space-efficient approaches, mainly variants of ShockHash [110] achieve below 1.55 bits per key (0.1 bits per key overhead). Especially the logarithmic plot shows how ShockHash and bipartite ShockHash improve the space consumption significantly. Bipartite ShockHash-Flat [109] trades larger space consumption and slower construction for better query speed. MorphisHash [89], as an enhancement to ShockHash, offers similar construction time but has lower space consumption. Because its queries are slightly faster due to the integrated retrieval data structure, it mostly covers ShockHash in the dominance maps.

Bucket Placement. PTHash [139] and CHD [11] have a similar trade-off between construction time and space consumption. PHOBIC [90] improves the construction time of PTHash significantly

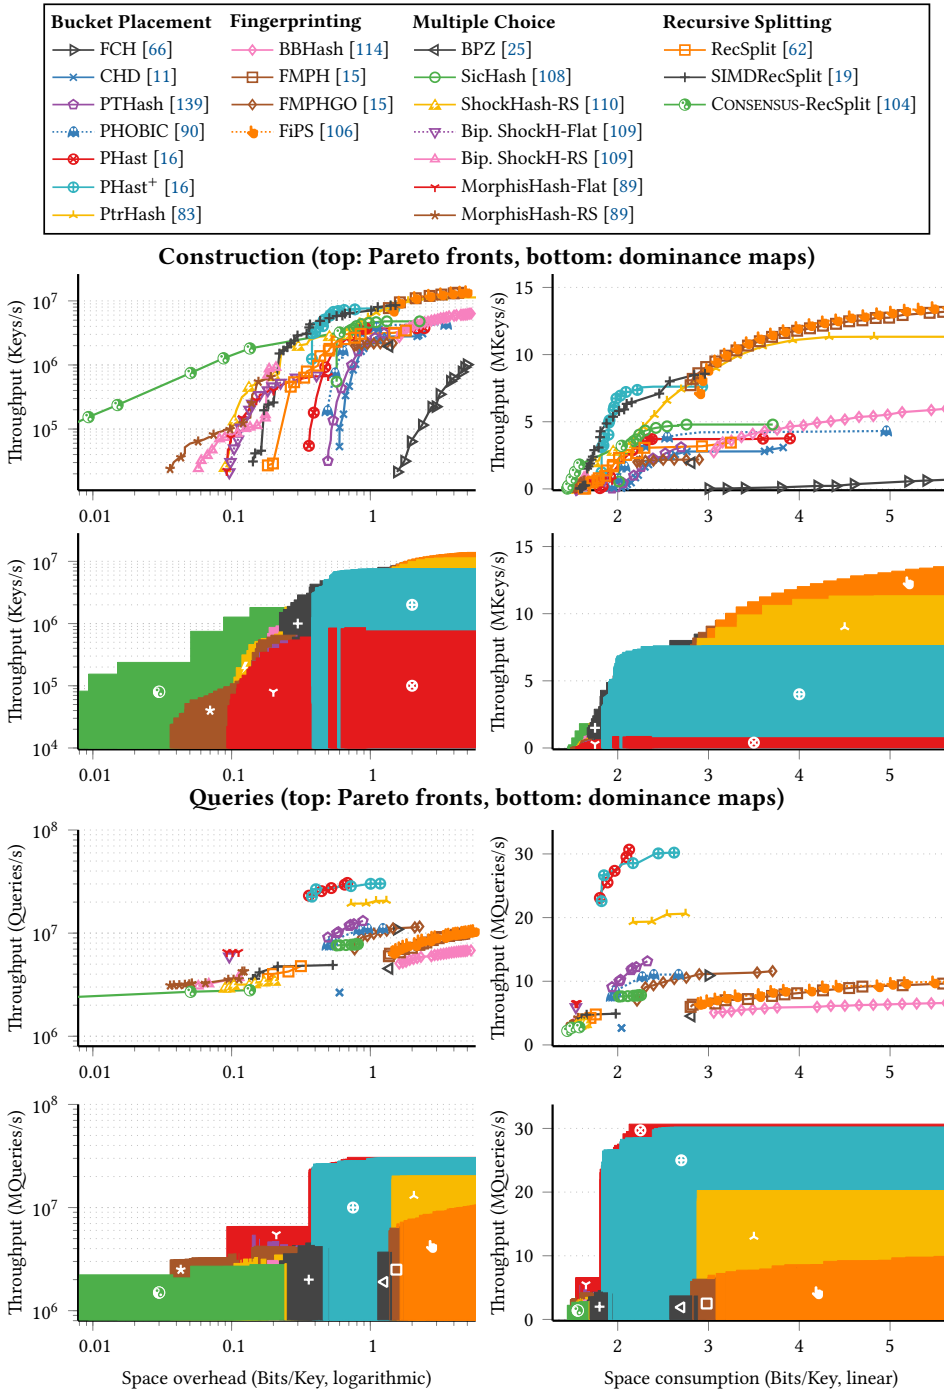


Figure 7. Trade-off construction time, query time, and space consumption. Single-threaded, $n = 100$ million string keys of length [10..50]. We do not show all markers to increase readability. The x -axis on the left is logarithmic to the space lower bound, so an overhead of 0 bits would correspond to $-\infty$.

through its optimized bucket assignment function. The construction time of FCH [66] is rather far from the other approaches. PtrHash [83], an extension of PTHash and PHOBIC, achieves significantly faster construction, although at a larger space consumption. Finally, PHast [16] is focused mainly on query speed but also achieves very good construction performance and small space consumption. Because of its fast queries that we will look at in Section 9.2, it fills most of the dominance map about space and construction time. Because PtrHash and PHast cover such a large range of configurations, they are a good initial choice in practice.

Recursive Splitting. RecSplit [62] is a construction that achieved a significant step towards the space lower bound at the time of its publication. Today, it is mostly dominated by the SIMD-parallel implementation. Just like RecSplit, SIMDRecSplit [19] is originally designed for small space consumption. We make the surprising observation that SIMDRecSplit not only wins for the most space-efficient configurations, but also has very competitive construction time for less space-efficient cases. However, being based on the splitting tree of RecSplit, it has quite slow queries, so in the dominance map, it only appears in areas that no other approaches can reach. CONSENSUS-RS [104] needs the smallest amount of space. As visible in the logarithmic versions of the plots, it significantly outperforms previous approaches in terms of space efficiency.

Fingerprinting. The goal of perfect hashing through fingerprinting [126] is to offer efficient queries and fast construction at the cost of larger space consumption. BBHash [114] is dominated by more recent implementations. FiPS [106] and recent version of FMPH [15] achieve almost twice the construction throughput and lower minimal space consumption. FMPHGO [15] reduces the space consumption of FMPH but also has a much slower construction. Its construction performance is similar to SicHash but has slower queries.

9.2 Query Trade-Off

We now discuss the approaches again, this time focusing on the query time. Figure 7 also includes a Pareto front and a dominance map for the query times. The dominance map plots the space consumption and the query throughput on the axes, and colors each point with the approach that has the fastest construction. In [107], we briefly consider the cache misses per query.

Multiple Choice Hashing. ShockHash-RS [110] and bipartite ShockHash-RS [109] use the RecSplit splitting tree but need an additional access to a retrieval data structure. However, their query performance is still very close to RecSplit. The reason is that they have fewer tree layers to traverse. It shows that the overhead of the retrieval operation is small compared to the work for traversing the heavily compressed tree. Bipartite ShockHash-Flat [109] is a variant focused on faster queries. For this, it sacrifices some of the space consumption of bipartite ShockHash-RS. For the same space consumption, it achieves 30% faster queries, which brings the query performance of very space-efficient MPHFs much closer to competitors that are not focused on space consumption. For both ShockHash variants, MorphisHash [89] uses a more cache-efficient retrieval data structure but its main contribution is smaller space consumption. SicHash [108] provides a middle ground between construction and query performance.

Bucket Placement. PHOBIC [90] improves the construction time of PTHash [139] without sacrificing too much query performance. However, PHast [16] is the clear winner in terms of query performance. It is almost three times faster than most other competitors, and only PtrHash [83] gets close. Simultaneously, PHast achieves good space consumption. CHD [11] is much slower to query, mostly due to decoding variable-length seeds. FCH [66] has fast queries, but needs more space.

Recursive Splitting. RecSplit [62] and SIMDRecSplit [19] are slow to query because they have to traverse the splitting tree, decoding variable-bitlength data in each step. However, even though the operations are much more complex, the performance is still solid and not too far away from FiPS and FMPH. For extremely small space consumption, CONSENSUS-RS is the only competitor, so it fills a large area of the dominance maps. Its query performance is similar to SIMDRecSplit.

Fingerprinting. Perfect hashing through fingerprinting [126] is originally designed to offer fast queries. Of the different implementations, FiPS [106] and FMPH [15] offer the best query throughput, being much faster than BBHash [114]. However, it is still far away from PtrHash. This holds even for a rather large space consumption of 3.5 bits per key ($\gamma = 2$), where about 73% of the keys can be handled in the first recursion layer. This indicates that the rank operation is rather costly, even when the rank data structure is interleaved with the bit vector. FMPHGO [15] achieves a smaller space consumption than the other fingerprinting approaches through some brute-force retries.

9.3 Scaling in the Input Size

We now look at the construction and query performance of different PHFs, depending on the number of input keys n . The goal is usually to have linear construction time. To achieve this, it is always possible to partition the input to smaller MPHFs. However, partitioning comes with a performance penalty, both during construction and queries. Figure 8 starts at rather small input sets of 1 million keys and goes all the way up to 200 million. In case partitions have to be used (e.g., for a parallel implementation), the plot can be used to decide on a partition size. For each approach, we use a configuration, that is, typical for it. Therefore, all competitors have a different space consumption and construction time to avoid using unfair configurations that approaches are not designed for. Due to the different space consumption, we give different plots for different space consumption. In the extended version [107] we give the exact configurations.

Construction. Figure 8 shows how the construction throughput scales in the number of input keys n . The throughput of CHD [11], BBHash [114], and FMPH [15] decreases most for larger input sizes. For the approaches based on fingerprinting, the slowdown can be explained by the random accesses to the large bit vector. FiPS [106] avoids these cache inefficient access patterns through the use of sorting, making it less sensitive to input size changes. SicHash and PHast⁺ have a large startup overhead and become relatively faster to construct with more keys. The remaining approaches scale almost linearly in n .

Queries. The query throughput of all approaches in Figure 8 drops when increasing the input size n . This is expected because the data structures get larger and do not fit in cache. In general, all approaches perform pretty well regarding their query time. For techniques below 2.5 bits/key, PHast [16] by far remains the approach with the fastest queries for the entire range of input sizes. PtrHash [83] comes close for large n . The query performance of PtrHash and PHast drops significantly when passing 10^7 keys. This can be explained by the fact that PtrHash queries are close to the RAM throughput. As the data structure becomes larger than the cache, the throughput drops. Due to PtrHash needing more space than PHast, it drops earlier. RecSplit based approaches (RecSplit [62], SIMDRecSplit [19], ShockHash-RS [109, 110], MorphisHash [89], CONSENSUS-RS [104]) have much slower queries for the entire range of input sizes but offer faster construction. Over a large range of input sizes, their query performance has a very small slope. This indicates that they are more limited by the computation than memory access.

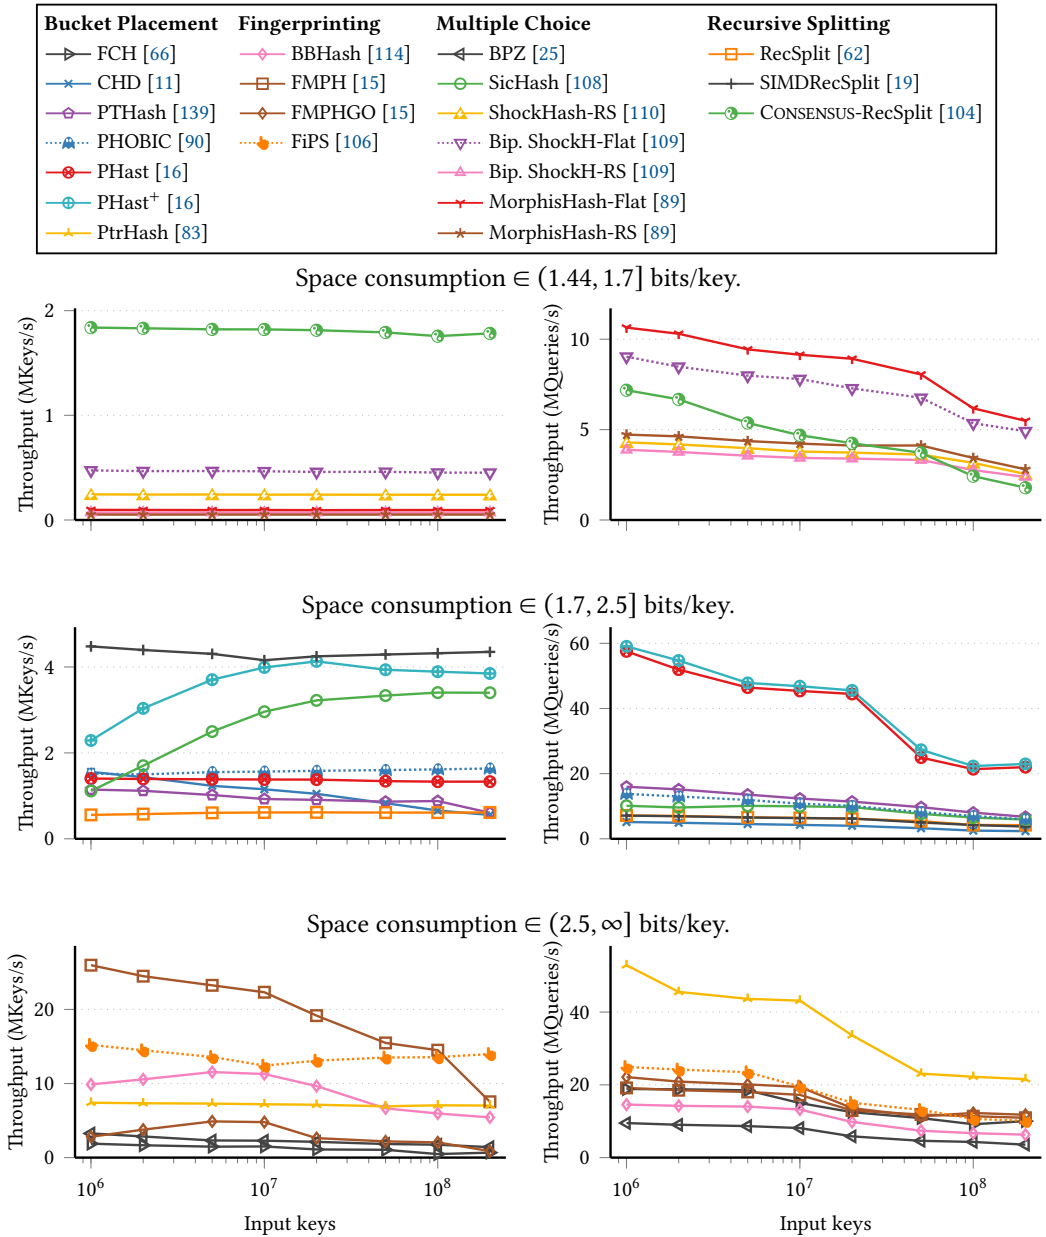


Figure 8. Comparison of construction and query performance by number of input keys n . Note that each approach has configuration parameters giving a wide trade-off between space, construction performance, and query performance. For the configurations used here, we refer to the extended version [107].

9.4 Multi-Threaded Construction

Modern processors have many cores available, and testing single-threaded code leaves a lot of processing power unused. Additionally, most data structures like PHFs are not used in isolation in actual applications. There are always other processes running on the machines. Performing

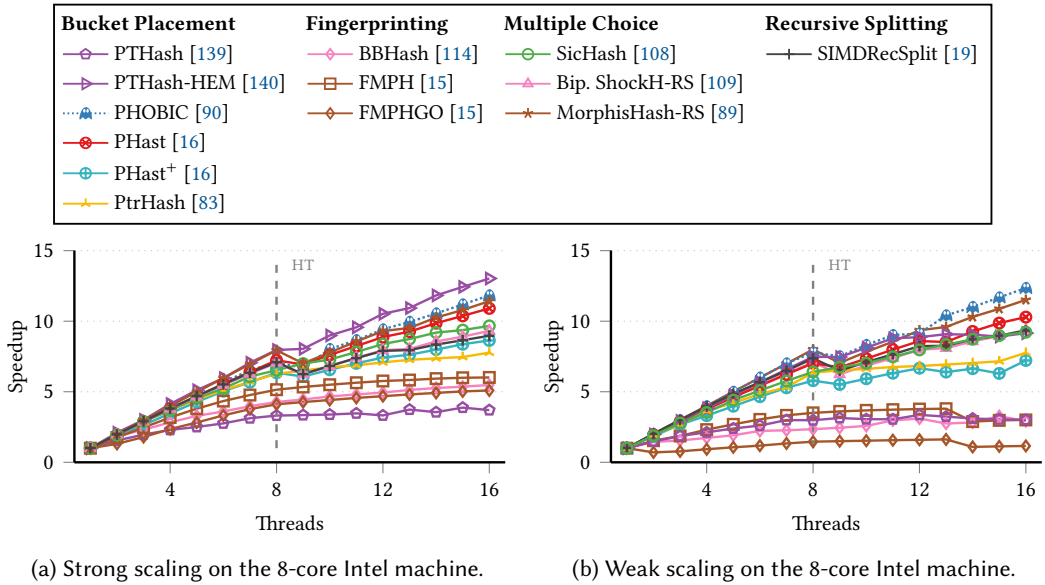


Figure 9. Multi-threaded construction by number of threads. Weak scaling with 10 million keys per thread, strong scaling with 100 million keys. We give self-speedups because each approach has a different focus.

multi-threaded measurements can, to a certain extent, account for this. Note that perfect hashing can be parallelized trivially by partitioning, which many approaches use. For these approaches, a major factor to their multi-threaded scaling is how efficiently they implement their partitioning step. This makes it less interesting algorithmically and causes a bias in the measurements. We still give multi-threaded measurements because they are very relevant in applications. Additionally, because there are fewer memory channels than threads, the cache locality of approaches is important in parallel measurements.

As a reminder, we run our experiments on an 8 core (16 HT) processor with 2 memory channels. In Figure 9, we give parallel measurements using both *weak scaling* and *strong scaling*. With weak scaling, the number of keys per thread are constant at 10 million. With strong scaling, the total number of input keys is 100 million and stays the same. For perfect hashing, strong scaling might be the more important approach, trying to construct a certain PHF as quickly as possible. However, with strong scaling there might be more threads than it is actually useful. Weak scaling ensures that each thread always has enough work to do. Like before, we use the configurations listed in the extended version [107]. As such, the absolute construction times are very different, so we only plot self-speedups. In the following, we first describe approaches with a direct parallelization before looking at the approaches using partitioning. In addition, the extended version [107] gives measurements on a 64 core machine with up to 1.28 billion keys.

Direct Parallelization. Some approaches use a parallel implementation of their internal data structures. An advantage of this technique compared to an external layer of partitioning is that it is transparent to the queries. Except for PHast [16], this generally does not seem to work well for the approaches that do it. We see that BBHash [114], PTHash [140], FMPH [15], and FMPHGO [15] only achieve a speedup of around 5 when running on 16 threads. We will later see that PTHash works better with partitioning (i.e., with the approach called PTHash-HEM). FMPHGO [15] uses a direct parallelization as well, even though it could use internal partitioning [15]. PHast, in contrast, places

all keys of a bucket in a small window of output positions. By leaving small gaps between different threads, they can work completely independently. The gaps can be filled afterwards transparently to the queries.

Internal Partitioning. Some of the approaches internally partition the input anyway, so in essence they get their parallelization for free. Note that the scaling behavior of SIMDRecSplit varies strongly depending on the configuration parameters [19]. The less space-efficient configurations that it is not actually designed for scale less well because more time is spent partitioning keys to a large number of buckets. PHOBIC [90] needs internal partitioning for its interleaved coding. It scales well with strong scaling and weak scaling, which we attribute to a well implemented partitioning step. PtrHash [83] scales well but does not profit as much from hyperthreading.

External Partitioning. Most other approaches parallelize by adding an additional layer of partitioning, which introduces small query and construction time overheads. PTHash-HEM [140] scales best among those. SicHash [108] could theoretically use internal partitioning. However, the majority of its construction time is spent constructing the BuRR retrieval data structures covering all internal partitions. While a parallel construction if BuRR is available now [7], it was not available when SicHash was presented. It therefore uses an external partitioning step.

GPU Parallelization. Even though this is not the focus of this evaluation, we mention two GPU parallel constructions. PHOBIC-GPU [90] and GPURecSplit [19] are, to our knowledge, the only PHFs with a GPU construction. Both approaches achieve a similar peak construction throughput of about 70 million keys per second on an Nvidia RTX 3090 GPU. In a reasonable construction time, PHOBIC-GPU can achieve a space consumption of about 1.7 bits per key, while GPURecSplit can achieve about 1.5 bits per key. The respective query implementations are identical to the CPU versions and can only be used on the CPU, so our measurements from Section 9.2 show that PHOBIC is much faster to query. Refer to the PHOBIC paper [90] for details.

10 Conclusion

In this article, we surveyed state-of-the-art PHF constructions. We categorized them by their working principles into approaches based on retrieval, brute-force, and fingerprinting. Historically, retrieval based methods performed the jump from logarithmic to constant number of bits per key. For this, they tuned the number of choices per key, the retrieval data structure, load factor, and search techniques. However, right now, purely retrieval-based methods cannot get below 2 bits per key. With ShockHash and its variants, the lower bound can be reached by combining it with brute-force. Several authors implement similar variants of perfect hashing through fingerprinting. For approaches based on brute-force, PTHash combines the ideas from FCH and CHD to achieve fast queries with good space consumption. PHOBIC then optimizes this approach by determining an optimal bucket assignment function. PtrHash and PHast are variants of bucket placement radically optimized for query speed, because of which they dominate a wide range of the trade-off. Finally, we looked at the more space-efficient approaches. The simple brute-force construction is made practical through the development of RecSplit, which is the first approach that goes a significant step towards the space lower bound. SIMDRecSplit, ShockHash, and MorphisHash then tune this to search in a much more structured way. They finally get beaten again by less structured brute-force again using CONSENSUS, which achieves extremely small space consumption.

Evaluation. After explaining the approaches, we performed and discussed a wide range of benchmarks. We gave plots illustrating the Pareto front of the entire trade-off between construction throughput, query throughput, and space consumption. In essence, the brute-force approaches focus on showing what is possible in terms of storage space, but sacrifice on construction and query

performance. The most space-efficient approach is CONSENSUS. Hash function construction through fingerprinting appealed through its simplicity and fast queries but is outperformed by more recent techniques now. Regarding query time, PHast offers the best performance, followed by PtrHash, both while also achieving rather small space. In general, over the past few years, perfect hashing has seen a large number of algorithmic improvements, leading to massive speedups and enabling larger and larger input sets.

Future Work. It is reasonable to assume that approaches that are more space-efficient are also slower to construct. Future work will likely give additional trade-offs and achieve even smaller space consumption. Part of this might also be enabled by additional GPU accelerated constructions. Current approaches have the tendency that the most space-efficient approaches are slower to query. There is no reason to think that this is a fundamental limitation, especially since smaller data structures should be more cache-efficient. We believe that in the future, it will be possible to construct MPHFs that are very close to the space lower bound, while still achieving fast queries. Given the large progress over the 28 years since the last survey, we are excited for the breakthroughs that the next decades might bring.

References

- [1] Anders Aamand, Jakob B. T. Knudsen, Mathias B. T. Knudsen, Peter M. R. Rasmussen, and Mikkel Thorup . 2020. Fast hashing with strong concentration bounds. In *Proceedings of the STOC*. ACM, 1265–1278. DOI : <https://doi.org/10.1145/3357713.3384259>
- [2] Rakesh Agrawal, Tomasz Imielinski, and Arun N. Swami. 1993. Mining association rules between sets of items in large databases. In *Proceedings of the SIGMOD Conference*. ACM Press, 207–216. DOI : <https://doi.org/10.1145/170035.170072>
- [3] Fatemeh Almodaresi, Hirak Sarkar, Avi Srivastava, and Rob Patro . 2018. A space and time-efficient index for the compacted colored de Bruijn graph. *Bioinform.* 34, 13 (2018), i169–i177. DOI : <https://doi.org/10.1093/BIOINFORMATICS/BTY292>
- [4] Michael R. Anderson and M. G. Anderson. 1979. Comments on Perfect Hashing Functions: A Single Probe Retrieving Method for Static Sets. *Commun. ACM* 22, 2 (1979), 104. DOI : <https://doi.org/10.1145/359060.359071>
- [5] Austin Appleby. 2010. SMHasher. Retrieved from <https://github.com/rurban/smhasher>
- [6] Sepehr Assadi, Martin Farach-Colton, and William Kuszmaul. 2023. Tight bounds for monotone minimal perfect hashing. In *Proceedings of the SODA*. SIAM, 456–476. DOI : <https://doi.org/10.1137/1.9781611977554.CH20>
- [7] Matthias Becht, Hans-Peter Lehmann, and Peter Sanders. 2024. Brief announcement: Parallel construction of bumped ribbon retrieval. DOI : <https://doi.org/10.48550/ARXIV.2411.12365>
- [8] Djamal Belazzougui, Paolo Boldi, Giuseppe Ottaviano, Rossano Venturini, and Sebastiano Vigna. 2014. Cache-oblivious peeling of random hypergraphs. In *Proceedings of the DCC*. IEEE, 352–361. DOI : <https://doi.org/10.1109/DCC.2014.48>
- [9] Djamal Belazzougui, Paolo Boldi, Rasmus Pagh, and Sebastiano Vigna. 2009. Monotone minimal perfect hashing: searching a sorted table with $O(1)$ accesses. In *SODA*. 785–794. DOI : <https://doi.org/10.1137/1.9781611973068.86>
- [10] Djamal Belazzougui, Paolo Boldi, Rasmus Pagh, and Sebastiano Vigna. 2010. Fast prefix search in little space, with applications. In *Proceedings of the ESA (1) (LNCS, Vol. 6346)*. Springer, 427–438. DOI : https://doi.org/10.1007/978-3-642-15775-2_37
- [11] Djamal Belazzougui, Fabiano C. Botelho, and Martin Dietzfelbinger. 2009. Hash, displace, and compress. In *Proceedings of the ESA (LNCS, Vol. 5757)*. Springer, 682–693. DOI : https://doi.org/10.1007/978-3-642-04128-0_61
- [12] Djamal Belazzougui, Fabio Cunial, Juha Kärkkäinen, and Veli Mäkinen. 2020. Linear-time string indexing and analysis in small space. *ACM Trans. Algorithms* 16, 2 (2020), 17:1–17:54. DOI : <https://doi.org/10.1145/3381417>
- [13] Djamal Belazzougui and Gonzalo Navarro. 2014. Alphabet-independent compressed text indexing. *ACM Trans. Algorithms* 10, 4 (2014), 23:1–23:19. DOI : <https://doi.org/10.1145/2635816>
- [14] Djamal Belazzougui, Gonzalo Navarro, and Daniel Valenzuela. 2013. Improved compressed indexes for full-text document retrieval. *J. Discrete Algorithms* 18 (2013), 3–13. DOI : <https://doi.org/10.1016/J.JDA.2012.07.005>
- [15] Piotr Beling. 2023. Fingerprinting-based Minimal Perfect Hashing Revisited. *ACM J. Exp. Algorithmics* 28 (2023), 1.4:1–1.4:16. DOI : <https://doi.org/10.1145/3596453>
- [16] Piotr Beling and Peter Sanders. 2025. PHast–Perfect hashing with fast evaluation. DOI : <https://doi.org/10.48550/arXiv.2504.17918>

- [17] Michael A. Bender, Martin Farach-Colton, Mayank Goswami, Rob Johnson, Samuel McCauley, and Shikha Singh . 2018. Bloom filters, adaptivity, and the dictionary problem. In *Proceedings of the FOCS*. IEEE, 182–193. DOI : <https://doi.org/10.1109/FOCS.2018.00026>
- [18] Ioana O. Bercea, Lorenzo Beretta, Jonas Klausen, Jakob Bæk Tejs Houen, and Mikkel Thorup. 2023. Locally uniform hashing. In *Proceedings of the FOCS*. IEEE, 1440–1470. DOI : <https://doi.org/10.1109/FOCS57990.2023.00089>
- [19] Dominik Bez, Florian Kurpicz, Hans-Peter Lehmann, and Peter Sanders . 2023. High performance construction of RecSplit based minimal perfect hash functions. In *ESA (LIPIcs, Vol. 274)*. 19:1–19:16. DOI : <https://doi.org/10.4230/LIPICS.ESA.2023.19>
- [20] Burton H. Bloom. 1970. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM* 13, 7 (1970), 422–426. DOI : <https://doi.org/10.1145/362686.362692>
- [21] Alexandra Boldyreva, Nathan Chenette, and Adam O’Neill . 2011. Order-Preserving Encryption Revisited: Improved Security Analysis and Alternative Solutions. In *Proceedings of the CRYPTO (LNCS, Vol. 6841)*. Springer, 578–595. DOI : https://doi.org/10.1007/978-3-642-22792-9_33
- [22] Fabiano C. Botelho, David M. Gomes, and Nivio Ziviani. 2004. A new algorithm for constructing minimal perfect hash functions. *Differences* 100, 2 (2004), 09.
- [23] Fabiano C. Botelho, Rasmus Pagh, and Nivio Ziviani. 2007. Perfect hashing for data management applications. DOI : <https://doi.org/10.48550/arXiv.cs/0702159>
- [24] Fabiano C. Botelho, Rasmus Pagh, and Nivio Ziviani. 2007. Simple and space-efficient minimal perfect hash functions. In *Proceedings of the WADS (LNCS, Vol. 4619)*. Springer, 139–150. DOI : https://doi.org/10.1007/978-3-540-73951-7_13
- [25] Fabiano C. Botelho, Rasmus Pagh, and Nivio Ziviani. 2013. Practical perfect hashing in nearly optimal space. *Inf. Syst.* 38, 1 (2013), 108–131. DOI : <https://doi.org/10.1016/j.is.2012.06.002>
- [26] Fabiano C. Botelho and Nivio Ziviani. 2007. External perfect hashing for very large key sets. In *Proceedings of the CIKM*. ACM, 653–662. DOI : <https://doi.org/10.1145/1321440.1321532>
- [27] Marshall D. Brain and Alan L. Tharp. 1990. Perfect hashing using sparse matrix packing. *Inf. Syst.* 15, 3 (1990), 281–290. DOI : [https://doi.org/10.1016/0306-4379\(90\)90001-6](https://doi.org/10.1016/0306-4379(90)90001-6)
- [28] Marshall D. Brain and Alan L. Tharp. 1994. Using tries to eliminate pattern collisions in perfect hashing. *IEEE Trans. Knowl. Data Eng.* 6, 2 (1994), 239–247. DOI : <https://doi.org/10.1109/69.277768>
- [29] Johann Brault-Baron. 2016. Hypergraph acyclicity revisited. *ACM Comput. Surv.* 49, 3, Article 54 (Dec. 2016), 26 pages. DOI : <https://doi.org/10.1145/2983573>
- [30] Andrei Z. Broder and Michael Mitzenmacher. 2003. Survey: Network applications of bloom filters: A survey. *Internet Math.* 1, 4 (2003), 485–509. DOI : <https://doi.org/10.1080/15427951.2004.10129096>
- [31] Michael Burrows and D. J. Wheeler. 1994. A block-sorting lossless data compression algorithm. *SRC* 124 (1994).
- [32] Larry Carter, Robert W. Floyd, John Gill, George Markowsky, and Mark N. Wegman. 1978. Exact and approximate membership testers. In *Proceedings of the STOC*. ACM, 59–65. DOI : <https://doi.org/10.1145/800133.804332>
- [33] Chin-Chen Chang. 1984. An ordered minimal perfect hashing scheme based upon Euler’s theorem. *Inf. Sci.* 32, 3 (1984), 165–172. DOI : [https://doi.org/10.1016/0020-0255\(84\)90032-X](https://doi.org/10.1016/0020-0255(84)90032-X)
- [34] Chin-Chen Chang and C. H. Chang. 1988. An ordered minimal perfect hashing scheme with single parameter. *Inf. Process. Lett.* 27, 2 (1988), 79–83. DOI : [https://doi.org/10.1016/0020-0190\(88\)90096-8](https://doi.org/10.1016/0020-0190(88)90096-8)
- [35] Chin-Chen Chang, C. Y. Chen, and Jinn-Ke Jan. 1991. On the design of a machine-independent perfect hashing scheme. *Comput. J.* 34, 5 (1991), 469–474. DOI : <https://doi.org/10.1093/COMJNL/34.5.469>
- [36] Chin-Chen Chang and Cih-Yang Lin. 2005. Perfect hashing schemes for mining association rules. *Comput. J.* 48, 2 (2005), 168–179. DOI : <https://doi.org/10.1093/COMJNL/BXH074>
- [37] Chin-Chen Chang and Tzong-Chen Wu. 1991. A letter-oriented perfect hashing scheme based upon sparse table compression. *Softw. Pract. Exp.* 21, 1 (1991), 35–49. DOI : <https://doi.org/10.1002/SPE.4380210104>
- [38] C. C. Chang. 1984. The study of an ordered minimal perfect hashing scheme. *Commun. ACM* 27, 4 (1984), 384–387. DOI : <https://doi.org/10.1145/358027.358051>
- [39] Jarrod A. Chapman, Isaac Ho, Sirisha Sunkara, Shujun Luo, Gary P Schroth, and Daniel S Rokhsar . 2011. Meraculous: De novo genome assembly with short paired-end reads. *PLOS One* 6, 8 (08 2011), 1–13. DOI : <https://doi.org/10.1371/journal.pone.0023501>
- [40] Bernard Chazelle, Joe Kilian, Ronitt Rubinfeld, and Ayellet Tal. 2004. The bloomier filter: An efficient data structure for static support lookup tables. In *Proceedings of the SODA*. SIAM, 30–39.
- [41] Rayan Chikhi, Antoine Limasset, and Paul Medvedev. 2016. Compacting de Bruijn graphs from sequencing data quickly and in low memory. *Bioinform.* 32, 12 (2016), 201–208. DOI : <https://doi.org/10.1093/BIOINFORMATICS/BTW279>
- [42] Richard J. Cichelli. 1980. Minimal perfect hash functions made simple. *Commun. ACM* 23, 1 (1980), 17–19. DOI : <https://doi.org/10.1145/358808.358813>
- [43] Curtis Cook and RR Oldehoeft. 1982. A letter oriented minimal perfect hashing function. *ACM SIGPlan* 17, 9 (1982), 18–27.

- [44] Gordon V. Cormack, R. Nigel Horspool, and Matthias Kaiserswerth. 1985. Practical perfect hashing. *Comput. J.* 28, 1 (1985), 54–58. DOI : <https://doi.org/10.1093/COMJNL/28.1.54>
- [45] Zbigniew J. Czech, George Havas, and Bohdan S. Majewski. 1992. An optimal algorithm for generating minimal perfect hash functions. *Inf. Process. Lett.* 43, 5 (1992), 257–264. DOI : [https://doi.org/10.1016/0020-0190\(92\)90220-P](https://doi.org/10.1016/0020-0190(92)90220-P)
- [46] Zbigniew J. Czech, George Havas, and Bohdan S. Majewski. 1997. Perfect hashing. *Theoretical Computer Science* 182, 1-2 (1997), 1–143. DOI : [https://doi.org/10.1016/S0304-3975\(96\)00146-6](https://doi.org/10.1016/S0304-3975(96)00146-6)
- [47] Zbigniew J. Czech and Bohdan S. Majewski. 1993. A linear time algorithm for finding minimal perfect hash functions. *Comput. J.* 36, 6 (1993), 579–587. DOI : <https://doi.org/10.1093/COMJNL/36.6.579>
- [48] Nicolaas Govert De Bruijn. 1946. A combinatorial problem. *Proceedings of the Section of Sciences of the Koninklijke Nederlandse Akademie van Wetenschappen te Amsterdam* 49, 7 (1946), 758–764.
- [49] Erik D. Demaine, Friedhelm Meyer auf der Heide, Rasmus Pagh, and Mihai Patrascu . 2006. De dictionariis dynamicis pauco spatio utentibus (lat. On dynamic dictionaries using little space). In *Proceedings of the LATIN (LNCS, Vol. 3887)*. Springer, 349–361. DOI : https://doi.org/10.1007/11682462_34
- [50] Martin Dietzfelbinger, Andreas Goerd, Michael Mitzenmacher, Andrea Montanari, Rasmus Pagh, and Michael Rink . 2010. Tight thresholds for cuckoo hashing via XORSAT. In *Proceedings of the ICALP (LNCS, Vol. 6198)*. Springer, 213–225. DOI : https://doi.org/10.1007/978-3-642-14165-2_19
- [51] Martin Dietzfelbinger and Friedhelm Meyer auf der Heide. 1990. A new universal class of hash functions and dynamic hashing in real time. In *Proceedings of the ICALP (LNCS, Vol. 443)*. Springer, 6–19. DOI : <https://doi.org/10.1007/BFB0032018>
- [52] Martin Dietzfelbinger, Joseph Gil, Yossi Matias, and Nicholas Pippenger. 1992. Polynomial hash functions are reliable (extended abstract). In *Proceedings of the ICALP (LNCS, Vol. 623)*. Springer, 235–246. DOI : https://doi.org/10.1007/3-540-55719-9_77
- [53] Martin Dietzfelbinger, Michael Mitzenmacher, and Michael Rink. 2011. Cuckoo hashing with pages. In *Proceedings of the ESA (LNCS, Vol. 6942)*. Springer, 615–627. DOI : https://doi.org/10.1007/978-3-642-23719-5_52
- [54] Martin Dietzfelbinger and Rasmus Pagh. 2008. Succinct data structures for retrieval and approximate membership. In *Proceedings of the ICALP (LNCS, Vol. 5125)*. Springer, 385–396. DOI : https://doi.org/10.1007/978-3-540-70575-8_32
- [55] Martin Dietzfelbinger and Michael Rink. 2009. Applications of a splitting trick. In *Proceedings of the ICALP (LNCS, Vol. 5555)*. Springer, 354–365. DOI : https://doi.org/10.1007/978-3-642-02927-1_30
- [56] Martin Dietzfelbinger and Stefan Walzer. 2019. Constant-time retrieval with $O(\log m)$ extra bits. In *STACS (LIPIcs, Vol. 126)*. 24:1–24:16. DOI : <https://doi.org/10.4230/LIPICS.STACS.2019.24>
- [57] Peter C. Dillinger, Lorenz Hübschle-Schneider, Peter Sanders, and Stefan Walzer. 2022. Fast succinct retrieval and approximate membership using ribbon. In *SEA (LIPIcs, Vol. 233)*. 4:1–4:20. DOI : <https://doi.org/10.4230/LIPICS.SEA.2022.4>
- [58] M. W. Du, T. M. Hsieh, K. F. Jea, and D. W. Shieh. 1983. The study of a new perfect hash scheme. *IEEE Trans. Software Eng.* 9, 3 (1983), 305–313. DOI : <https://doi.org/10.1109/TSE.1983.236866>
- [59] Stefan Edelkamp, Peter Sanders, and Pavel Simecek. 2008. Semi-external LTL model checking. In *Proceedings of the CAV (LNCS, Vol. 5123)*. Springer, 530–542. DOI : https://doi.org/10.1007/978-3-540-70545-1_50
- [60] Peter Elias. 1974. Efficient storage and retrieval by content and address of static files. *J. ACM* 21, 2 (1974), 246–260. DOI : <https://doi.org/10.1145/321812.321820>
- [61] Indeed Engineering. 2018. Indeed MPH: Fast and Compact Immutable Key-Value Stores. Retrieved from <https://engineering.indeedblog.com/blog/2018/02/indeed-mph>
- [62] Emmanuel Esposito, Thomas Mueller Graf, and Sebastiano Vigna. 2020. RecSplit: Minimal perfect hashing via recursive splitting. In *Proceedings of the ALENEX*. SIAM, 175–185. DOI : <https://doi.org/10.1137/1.9781611976007.14>
- [63] Bin Fan, David G. Andersen, Michael Kaminsky, and Michael Mitzenmacher. 2014. Cuckoo filter: Practically better than bloom. In *Proceedings of the CoNEXT*. ACM, 75–88. DOI : <https://doi.org/10.1145/2674005.2674994>
- [64] Robert Mario Fano. 1971. *On the Number of Bits Required to Implement an Associative Memory*. Technical Report. MIT, Computer Structures Group. Project MAC, Memorandum 61”.
- [65] Edward A. Fox, Qi Fan Chen, Amjad M. Daoud, and Lenwood S. Heath. 1991. Order-preserving minimal perfect hash functions and information retrieval. *ACM Trans. Inf. Syst.* 9, 3 (1991), 281–308. DOI : <https://doi.org/10.1145/125187.125200>
- [66] Edward A. Fox, Qi Fan Chen, and Lenwood S. Heath. 1992. A faster algorithm for constructing minimal perfect hash functions. In *Proceedings of the SIGIR*. ACM, 266–273. DOI : <https://doi.org/10.1145/133160.133209>
- [67] Edward A. Fox, Lenwood S Heath, and Qi-Fan Chen. 1989. An $O(n \log n)$ algorithm for finding minimal perfect hash functions. (1989).
- [68] Edward A. Fox, Lenwood S. Heath, Qi Fan Chen, and Amjad M. Daoud. 1992. Practical minimal perfect hash functions for large databases. *Commun. ACM* 35, 1 (1992), 105–121. DOI : <https://doi.org/10.1145/129617.129623>

- [69] Michael L. Fredman and János Komlós. 1984. On the size of separating systems and families of perfect hash functions. *SIAM J. Algebraic Discrete Methods* 5, 1 (1984), 61–68. DOI: <https://doi.org/10.1137/0605009>
- [70] Michael L. Fredman, János Komlós, and Endre Szemerédi. 1984. Storing a sparse table with $O(1)$ worst case access time. *J. ACM* 31, 3 (1984), 538–544. DOI: <https://doi.org/10.1145/828.1884>
- [71] Kimmo Fredriksson and Fedor Nikitin. 2007. Simple compression code supporting random access and fast string matching. In *Proceedings of the WEA (LNCS, Vol. 4525)*. Springer, 203–216. DOI: https://doi.org/10.1007/978-3-540-72845-0_16
- [72] Kevin P. Gaffney and Jignesh M. Patel. 2024. Is perfect hashing practical for OLAP systems?. In *CIDR*.
- [73] Travis Gagie, Gonzalo Navarro, and Nicola Prezza. 2020. Fully functional suffix trees and optimal text searching in BWT-runs bounded space. *J. ACM* 67, 1 (2020), 2:1–2:54. DOI: <https://doi.org/10.1145/3375890>
- [74] Marco Genuzio, Giuseppe Ottaviano, and Sebastiano Vigna. 2016. Fast scalable construction of (minimal perfect hash) functions. In *Proceedings of the SEA (LNCS, Vol. 9685)*. Springer, 339–352. DOI: https://doi.org/10.1007/978-3-319-38851-9_23
- [75] Andreas Goerdt and Lutz Falke. 2012. Satisfiability thresholds beyond k -XORSAT. In *Proceedings of the CSR (LNCS, Vol. 7353)*. Springer, 148–159. DOI: https://doi.org/10.1007/978-3-642-30642-6_15
- [76] Solomon W. Golomb. 1966. Run-length encodings. *IEEE Trans. Inf. Theory* 12, 3 (1966), 399–401. DOI: <https://doi.org/10.1109/TIT.1966.1053907>
- [77] Irving John Good. 1946. Normal recurring decimals. *Journal of the London Mathematical Society* 1, 3 (1946), 167–169.
- [78] Google. 2017. Abseil. Retrieved from <https://abseil.io>
- [79] Google. 2017. Abseil’s Swiss Table. Retrieved from <https://abseil.io/about/design/swisstables>
- [80] Marco Gori and Giovanni Soda. 1989. An algebraic approach to cichelli’s perfect hashing. *BIT* 29, 1 (1989), 2–13. DOI: <https://doi.org/10.1007/BF01932700>
- [81] Manfred G. Grabherr, Brian J Haas, Moran Yassour, Joshua Z Levin, Dawn A Thompson, Ido Amit, Xian Adiconis, Lin Fan, Raktima Raychowdhury, and Qiandong Zeng, et al. 2011. Full-length transcriptome assembly from RNA-Seq data without a reference genome. *Nature Biotechnology* 29, 7 (2011), 644–652. DOI: <https://doi.org/10.1038/nbt.1883>
- [82] Thomas Mueller Graf and Daniel Lemire. 2020. Xor filters: Faster and smaller than bloom and cuckoo filters. *J. Exp. Algorithmics* 25 (2020), 1–16. DOI: <https://doi.org/10.1145/3376122>
- [83] Ragnar Groot Koerkamp. 2025. PtrHash: Minimal perfect hashing at RAM throughput. In *SEA (LIPICs, Vol. 338)*. 21:1–21:21. DOI: <https://doi.org/10.4230/LIPICs.SEA.2025.21>
- [84] Roberto Grossi, Alessio Orlandi, and Rajeev Raman. 2010. Optimal trade-offs for succinct string indexes. In *Proceedings of the ICALP (LNCS, Vol. 6198)*. Springer, 678–689. DOI: https://doi.org/10.1007/978-3-642-14165-2_57
- [85] Torben Hagerup and Torsten Tholey. 2001. Efficient minimal perfect hashing in nearly minimal space. In *Proceedings of the STACS (LNCS, Vol. 2010)*. Springer, 317–326. DOI: https://doi.org/10.1007/3-540-44693-1_28
- [86] Gary Haggard and Kevin Karplus. 1986. Finding minimal perfect hash functions. In *Proceedings of the SIGCSE*. ACM, 191–193. DOI: <https://doi.org/10.1145/5600.5899>
- [87] George Havas and Bohdan S. Majewski. 1993. *Graph Theoretic Obstacles to Perfect Hashing*. Key Centre for Software Technology, Department of Computer Science, The University of Queensland, Australia.
- [88] George Havas, Bohdan S. Majewski, Nicholas C. Wormald, and Zbigniew J. Czech. 1993. Graphs, hypergraphs and hashing. In *Proceedings of the WG (LNCS, Vol. 790)*. Springer, 153–165. DOI: https://doi.org/10.1007/3-540-57899-4_49
- [89] Stefan Hermann. 2025. MorphisHash: Improving space efficiency of shockhash for minimal perfect hashing. In *ESA (LIPICs)*. 9:1–9:16. DOI: <https://doi.org/10.4230/LIPICs.ESA.2025.9>
- [90] Stefan Hermann, Hans-Peter Lehmann, Giulio Ermanno Pibiri, Peter Sanders, and Stefan Walzer. 2024. PHOBIC: Perfect hashing with optimized bucket sizes and interleaved coding. In *ESA (LIPICs, Vol. 308)*. 69:1–69:17. DOI: <https://doi.org/10.4230/LIPICs.ESA.2024.69>
- [91] Gwo-Jen Hwang, W. F. Tsai, and Judy C. R. Tseng. 2002. A minimal perfect hashing approach for mining association rules from very large databases. In *Proceedings of the IASTED Conference*. 80–85.
- [92] Zamin Iqbal, Mario Caccamo, Isaac Turner, Paul Flicek, and Gil McVean. 2012. De novo assembly and genotyping of variants using colored de Bruijn graphs. *Nature Genetics* 44, 2 (2012), 226–232.
- [93] C. T. M. Jacobs and Peter van Emde Boas. 1986. Two results on tables. *Inf. Process. Lett.* 22, 1 (1986), 43–48. DOI: [https://doi.org/10.1016/0020-0190\(86\)90041-4](https://doi.org/10.1016/0020-0190(86)90041-4)
- [94] Guy Jacobson. 1989. Space-efficient static trees and graphs. In *Proceedings of the FOCS*. IEEE, 549–554. DOI: <https://doi.org/10.1109/SFCS.1989.63533>
- [95] G Jaeschke. 1980. On Cichelli’s minimal perfect hash functions method. *Comm. ACM* 23, 12 (1980), 728–729.
- [96] Gerhard Jaeschke. 1981. Reciprocal hashing: A method for generating minimal perfect hashing functions. *Commun. ACM* 24, 12 (1981), 829–833. DOI: <https://doi.org/10.1145/358800.358806>
- [97] Svante Janson and Malwina J. Luczak. 2007. A simple solution to the k -core problem. *Random Struct. Algorithms* 30, 1-2 (2007), 50–62. DOI: <https://doi.org/10.1002/RSA.20147>

- [98] Jamshed Khan, Marek Kokot, Sebastian Deorowicz, and Rob Patro . 2022. Scalable, ultra-fast, and low-memory construction of compacted de Bruijn graphs with Cuttlefish 2. *Genome Biol.* 23, 1 (2022), 190. DOI : <https://doi.org/10.1186/s13059-022-02743-6>
- [99] Donald E. Knuth. 2011. *The Art of Computer Programming: Combinatorial algorithms. Part 1.* Vol. 4A. Addison-Wesley.
- [100] Dmitry Kosolobov. 2024. Simplified tight bounds for monotone minimal perfect hashing. In *CPM (LIPICs, Vol. 296)*. 19:1–19:13. DOI : <https://doi.org/10.4230/LIPICs.CPM.2024.19>
- [101] Florian Kurpicz, Hans-Peter Lehmann, and Peter Sanders. 2023. PaCHash: Packed and compressed hash tables. In *Proceedings of the ALENEX*. SIAM, 162–175. DOI : <https://doi.org/10.1137/1.9781611977561.CH14>
- [102] William Kuszmaul, Aaron Putterman, Tingqiang Xu, Hangrui Zhou, and Renfei Zhou. 2025. Tight bounds and phase transitions for incremental and dynamic retrieval. In *Proceedings of the SODA*. SIAM, 3974–3997. DOI : <https://doi.org/10.1137/1.9781611978322.135>
- [103] William Kuszmaul and Stefan Walzer. 2024. Space lower bounds for dynamic filters and value-dynamic retrieval. In *Proceedings of the STOC*. ACM, 1153–1164. DOI : <https://doi.org/10.1145/3618260.3649649>
- [104] Hans-Peter Lehmann, Peter Sanders, Stefan Walzer, and Jonatan Ziegler. 2025. Combined search and encoding for seeds, with an application to minimal perfect hashing. In *ESA (LIPICs)* 351 (2025), 109:1–109:18. DOI : <https://doi.org/10.48550/ARXIV.2502.05613>
- [105] Hans-Peter Lehmann. 2023. MPHf Experiments - GitHub. Retrieved from <https://github.com/ByteHamster/MPHF-Experiments>
- [106] Hans-Peter Lehmann. 2024. *Fast and Space-Efficient Perfect Hashing*. Ph.D. Dissertation. Karlsruher Institut für Technologie (KIT). DOI : <https://doi.org/10.5445/IR/1000176432>
- [107] Hans-Peter Lehmann, Thomas Mueller, Rasmus Pagh, Giulio Ermanno Pibiri, Peter Sanders, Sebastiano Vigna, and Stefan Walzer. 2025. Modern minimal perfect hashing: A survey. DOI : <https://doi.org/10.48550/ARXIV.2506.06536>
- [108] Hans-Peter Lehmann, Peter Sanders, and Stefan Walzer. 2023. SicHash – Small irregular cuckoo tables for perfect hashing. In *Proceedings of the ALENEX*. SIAM, 176–189. DOI : <https://doi.org/10.1137/1.9781611977561.CH15>
- [109] Hans-Peter Lehmann, Peter Sanders, and Stefan Walzer. 2023. SicHash small irregular cuckoo tables for perfect hashing. In *ALENEX*. SIAM, 176–189. DOI : <https://doi.org/10.1137/1.9781611977561.CH15>
- [110] Hans-Peter Lehmann, Peter Sanders, and Stefan Walzer. 2024. ShockHash: Towards optimal-space minimal perfect hashing beyond brute-force. In *Proceedings of the ALENEX*. SIAM, 194–206. DOI : <https://doi.org/10.1137/1.9781611977929.15>
- [111] Ted G. Lewis and Curtis R. Cook. 1988. Hashing for dynamic and static internal tables. *Computer* 21, 10 (1988), 45–56. DOI : <https://doi.org/10.1109/2.7056>
- [112] Xingyu Liao, Min Li, You Zou, Fang-Xiang Wu, Yi Pan, and Jianxin Wang. 2019. Current challenges and solutions of de novo assembly. *Quant. Biol.* 7, 2 (2019), 90–109. DOI : <https://doi.org/10.1007/S40484-019-0166-9>
- [113] Hyeontaek Lim, Bin Fan, David G. Andersen, and Michael Kaminsky. 2011. SILT: A memory-efficient, high-performance key-value store. In *Proceedings of the SOSP*. ACM, 1–13. DOI : <https://doi.org/10.1145/2043556.2043558>
- [114] Antoine Limasset, Guillaume Rizk, Rayan Chikhi, and Pierre Peterlongo. 2017. Fast and scalable minimal perfect hashing for massive key sets. In *SEA (LIPICs, Vol. 75)*. 25:1–25:16. DOI : <https://doi.org/10.4230/LIPICs.SEA.2017.25>
- [115] Yi Lu, Balaji Prabhakar, and Flavio Bonomi. 2006. Perfect hashing for network applications. In *Proceedings of the ISIT*. IEEE, 2774–2778.
- [116] Harry G. Mairson. 1983. The program complexity of searching a table. In *Proceedings of the FOCS*. IEEE, 40–47. DOI : <https://doi.org/10.1109/SFCS.1983.76>
- [117] Bohdan S. Majewski, Nicholas C. Wormald, George Havas, and Zbigniew J. Czech. 1996. A family of perfect hashing methods. *Comput. J.* 39, 6 (1996), 547–554. DOI : <https://doi.org/10.1093/COMJNL/39.6.547>
- [118] Christopher Manning and Hinrich Schütze. 1999. *Foundations of Statistical Natural Language Processing*. MIT press.
- [119] Camille Marchet, Lolita Lecompte, Antoine Limasset, Lucie Bittner, and Pierre Peterlongo . 2020. A resource-frugal probabilistic dictionary and applications in bioinformatics. *Discret. Appl. Math.* 274 (2020), 92–102. DOI : <https://doi.org/10.1016/J.DAM.2018.03.035>
- [120] Camille Marchet, Maël Kerbiriou, and Antoine Limasset. 2021. BLight: Efficient exact associative structure for k-mers. *Bioinform.* 37, 18 (2021), 2858–2865. DOI : <https://doi.org/10.1093/BIOINFORMATICS/BTAB217>
- [121] Kurt Mehlhorn. 1982. On the program size of perfect and universal hash functions. In *Proceedings of the FOCS*. IEEE, 170–175. DOI : <https://doi.org/10.1109/SFCS.1982.80>
- [122] Kurt Mehlhorn. 1984. Data structures and algorithms, Vol. 1: Sorting and searching. *EATCS Monogr. Theor. Comput. Sci.* (1984).
- [123] Michael Mitzenmacher. 2001. The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distributed Syst.* 12, 10 (2001), 1094–1104. DOI : <https://doi.org/10.1109/71.963420>
- [124] Michael Molloy. 2005. Cores in random hypergraphs and Boolean formulas. *Random Struct. Algorithms* 27, 1 (2005), 124–135. DOI : <https://doi.org/10.1002/RSA.20061>

- [125] Christian Worm Mortensen, Rasmus Pagh, and Mihai Pătraşcu. 2005. On dynamic range reporting in one dimension. In *Proceedings of the STOC*. ACM, 104–111. DOI : <https://doi.org/10.1145/1060590.1060606>
- [126] Ingo Müller, Peter Sanders, Robert Schulze, and Wei Zhou. 2014. Retrieval and perfect hashing using fingerprinting. In *Proceedings of the SEA (LNCS, Vol. 8504)*. Springer, 138–149. DOI : https://doi.org/10.1007/978-3-319-07959-2_12
- [127] Gonzalo Navarro. 2014. Spaces, trees, and colors: The algorithmic landscape of document retrieval on sequences. *ACM Comput. Surv.* 46, 4 (2014), 1–47. DOI : <https://doi.org/10.1145/2535933>
- [128] Gonzalo Navarro. 2016. *Compact Data Structures - A Practical Approach*. Cambridge University Press.
- [129] Racin Nygaard and Hein Meling. 2023. SNIPS: Succinct proof of storage for efficient data synchronization in decentralized storage systems. DOI : <https://doi.org/10.48550/ARXIV.2304.04891>
- [130] Anna Pagh and Rasmus Pagh. 2008. Uniform hashing in constant time and optimal space. *SIAM J. Comput.* 38, 1 (2008), 85–96. DOI : <https://doi.org/10.1137/060658400>
- [131] Anna Pagh, Rasmus Pagh, and Milan Ruzic. 2007. Linear probing with constant independence. In *Proceedings of the STOC*. ACM, 318–327. DOI : <https://doi.org/10.1145/1250790.1250839>
- [132] Rasmus Pagh. 1999. Hash and displace: Efficient evaluation of minimal perfect hash functions. In *Proceedings of the WADS (LNCS, Vol. 1663)*. Springer, 49–54. DOI : https://doi.org/10.1007/3-540-48447-7_5
- [133] Rasmus Pagh and Flemming Friche Rodler. 2004. Cuckoo hashing. *J. Algorithms* 51, 2 (2004), 122–144. DOI : <https://doi.org/10.1016/j.jalgor.2003.12.002>
- [134] Peter K. Pearson. 1990. Fast hashing of variable-length text strings. *Commun. ACM* 33, 6 (1990), 677–680.
- [135] Pavel A. Pevzner, Haixu Tang, and Michael S. Waterman. 2001. An eulerian path approach to DNA fragment assembly. *Proc. Natl. Acad. Sci.* 98, 17 (2001), 9748–9753. DOI : <https://doi.org/10.1073/pnas.171285098>
- [136] Giulio Ermanno Pibiri. 2022. Sparse and skew hashing of K-mers. *Bioinform.* 38, Supplement_1 (2022), i185–i194. DOI : <https://doi.org/10.1093/BIOINFORMATICS/BTAC245>
- [137] Giulio Ermanno Pibiri. 2023. On weighted k-mer dictionaries. *Algorithms Mol. Biol.* 18, 1 (2023), 3. DOI : <https://doi.org/10.1186/S13015-023-00226-2>
- [138] Giulio Ermanno Pibiri, Yoshihiro Shibuya, and Antoine Limasset. 2023. Locality-preserving minimal perfect hashing of k-mers. *Bioinform.* 39, Supplement-1 (2023), 534–543. DOI : <https://doi.org/10.1093/BIOINFORMATICS/BTAD219>
- [139] Giulio Ermanno Pibiri and Roberto Trani. 2021. PTHash: Revisiting FCH minimal perfect hashing. In *Proceedings of the SIGIR*. ACM, 1339–1348. DOI : <https://doi.org/10.1145/3404835.3462849>
- [140] Giulio Ermanno Pibiri and Roberto Trani. 2024. Parallel and external-memory construction of minimal perfect hash functions with PTHash. *IEEE Trans. Knowl. Data Eng.* 36, 3 (2024), 1249–1259. DOI : <https://doi.org/10.1109/TKDE.2023.3303341>
- [141] Giulio Ermanno Pibiri and Rossano Venturini. 2017. Efficient data structures for massive N-Gram Datasets. In *Proceedings of the SIGIR*. ACM, 615–624. DOI : <https://doi.org/10.1145/3077136.3080798>
- [142] Giulio Ermanno Pibiri and Rossano Venturini. 2019. Handling massive N-Gram datasets efficiently. *ACM Trans. Inf. Syst.* 37, 2 (2019), 25:1–25:41. DOI : <https://doi.org/10.1145/3302913>
- [143] Boris G. Pittel and Gregory B. Sorkin. 2016. The satisfiability threshold for k-XORSAT. *Comb. Probab. Comput.* 25, 2 (2016), 236–268. DOI : <https://doi.org/10.1017/S0963548315000097>
- [144] M. V. Ramakrishna and Per-Åke Larson. 1989. File organization using composite perfect hashing. *ACM Trans. Database Syst.* 14, 2 (1989), 231–263. DOI : <https://doi.org/10.1145/63500.63521>
- [145] Julian Reichinger, Thomas Krismayer, and Jan S. Rellermeier. 2024. COPR – Efficient, large-scale log storage and retrieval. DOI : <https://doi.org/10.48550/ARXIV.2402.18355>
- [146] Robert F. Rice. 1979. Some practical universal noiseless coding techniques. *JPL Publ.* (1979).
- [147] Thomas J Sager. 1984. A new method for generating minimal perfect hash functions. (1984).
- [148] Thomas J. Sager. 1985. A polynomial time generator for minimal perfect hash functions. *Commun. ACM* 28, 5 (1985), 523–532. DOI : <https://doi.org/10.1145/3532.3538>
- [149] Jeanette P. Schmidt and Alan Siegel. 1990. The spatial complexity of oblivious k-probe hash functions. *SIAM J. Comput.* 19, 5 (1990), 775–786. DOI : <https://doi.org/10.1137/0219054>
- [150] Sebastian Schmidt, Shahbaz Khan, Jarno N. Alanko, Giulio E. Pibiri, and Alexandru I. Tomescu . 2023. Matchtigs: Minimum plain text representation of k-mer sets. *Genome Biology* 24, 1 (2023), 136. DOI : <https://doi.org/10.1186/s13059-023-02968-z>
- [151] Cees F. Slot and Peter van Emde Boas. 1984. On tape versus core; An application of space efficient perfect hash functions to the invariance of space. In *Proceedings of the STOC*. ACM, 391–400. DOI : <https://doi.org/10.1145/800057.808705>
- [152] Renzo Sprugnoli. 1977. Perfect hashing functions: A single probe retrieving method for static sets. *Commun. ACM* 20, 11 (1977), 841–850. DOI : <https://doi.org/10.1145/359863.359887>
- [153] Grant P. Strimel, Ariya Rastrow, Gautam Tiwari, Adrien Piérard, and Jon Webb . 2020. Rescore in a flash: Compact, cache efficient hashing data structures for n-gram language models. In *INTERSPEECH*. DOI : <https://doi.org/10.21437/INTERSPEECH.2020-1939>

- [154] Robert Endre Tarjan and Andrew Chi-Chih Yao. 1979. Storing a sparse table. *Commun. ACM* 22, 11 (1979), 606–611.
- [155] Mikkel Thorup. 2017. Fast and powerful hashing using tabulation. *Commun. ACM* 60, 7 (2017), 94–101. DOI : <https://doi.org/10.1145/3068772>
- [156] John A. Trono. 1992. An undergraduate project to compute minimal perfect hashing functions. *ACM SIGCSE Bull.* 24, 3 (1992), 53–56. DOI : <https://doi.org/10.1145/142040.142077>
- [157] Stefan Walzer. 2021. Peeling close to the orientability threshold - spatial coupling in hashing-based data structures. In *Proceedings of the SODA*. SIAM, 2194–2211. DOI : <https://doi.org/10.1137/1.9781611976465.131>
- [158] Stefan Walzer. 2025. Peeling close to the orientability threshold – spatial coupling in hashing-based data structures. *ACM Trans. Algorithms* 21, 3 (2025), 33:1–33:23. DOI : <https://doi.org/10.1145/3711822>
- [159] Sean A. Weaver and Marijn Heule. 2020. Constructing minimal perfect hash functions using SAT technology. In *Proceedings of the AAAI*. AAAI Press, 1668–1675. DOI : <https://doi.org/10.1609/AAAI.V34I02.5529>
- [160] Peter Weiner. 1973. Linear pattern matching algorithms. In *Proceedings of the SWAT*. IEEE, 1–11. DOI : <https://doi.org/10.1109/SWAT.1973.13>
- [161] Vincent G. Winters. 1990. Minimal perfect hashing for large sets of data. In *Proceedings of the ICCI (LNCS, Vol. 468)*. Springer, 275–284. DOI : https://doi.org/10.1007/3-540-53504-7_85
- [162] Vincent G. Winters. 1990. Minimal perfect hashing in polynomial time. *BIT* 30, 2 (1990), 235–244. DOI : <https://doi.org/10.1007/BF02017345>
- [163] Ian H. Witten, Alistair Moffat, and Timothy C. Bell. 1999. *Managing Gigabytes: Compressing and Indexing Documents and Images, Second Edition*. Morgan Kaufmann.
- [164] Wei-Pang Yang and M. W. Du. 1985. A backtracking method for constructing perfect hash functions from a set of mapping functions. *BIT* 25, 1 (1985), 148–164. DOI : <https://doi.org/10.1007/BF01934995>

Received 6 June 2025; revised 18 January 2026; accepted 4 February 2026