



# ShockHash: Near Optimal-Space Minimal Perfect Hashing Beyond Brute-Force

Hans-Peter Lehmann<sup>1</sup>  · Peter Sanders<sup>1</sup>  · Stefan Walzer<sup>1</sup> 

Received: 21 May 2024 / Accepted: 24 April 2025 / Published online: 2 August 2025  
© The Author(s) 2025

## Abstract

A minimal perfect hash function (MPHF) maps a set  $S$  of  $n$  keys to the first  $n$  integers without collisions. There is a lower bound of  $n \log_2 e - \mathcal{O}(\log n) \approx 1.44n$  bits needed to represent an MPHF. This can be reached by a *brute-force* algorithm that tries  $e^n$  hash function seeds in expectation and stores the first seed that leads to an MPHF. The most space-efficient previous algorithms for constructing MPHFs all use such a brute-force approach as a basic building block. In this paper, we introduce ShockHash – Small, heavily overloaded **cuckoo hash** tables for minimal perfect hashing. ShockHash uses two hash functions  $h_0$  and  $h_1$ , hoping for the existence of a function  $f : S \rightarrow \{0, 1\}$  such that  $x \mapsto h_{f(x)}(x)$  is an MPHF on  $S$ . It then uses a 1-bit retrieval data structure to store  $f$  using  $n + o(n)$  bits. In graph terminology, ShockHash generates  $n$ -edge random graphs until stumbling on a *pseudoforest* – where each component contains as many edges as nodes. Using cuckoo hashing, ShockHash then derives an MPHF from the pseudoforest in linear time. We show that ShockHash needs to try only about  $(e/2)^n \approx 1.359^n$  seeds in expectation. This reduces the space for storing the seed by roughly  $n$  bits (maintaining the asymptotically optimal space consumption) and speeds up construction by almost a factor of  $2^n$  compared to brute-force. *Bipartite* ShockHash reduces the expected construction time again to about  $1.166^n$  by maintaining a pool of candidate hash functions and checking all possible pairs. Using ShockHash as a building block within the RecSplit framework we obtain ShockHash-RS, which can be constructed up to 3 orders of magnitude faster than competing approaches. ShockHash-RS can build an MPHF for 10 million keys with 1.489 bits per key in about half an hour. When instead using ShockHash after an efficient  $k$ -perfect hash function, it achieves space usage similar to the best competitors, while being significantly faster to construct and query.

---

✉ Hans-Peter Lehmann  
hans-peter.lehmann@kit.edu

Peter Sanders  
sanders@kit.edu

Stefan Walzer  
stefan.walzer@kit.edu

<sup>1</sup> Karlsruhe Institute of Technology, Karlsruhe, Germany

**Keywords** Compressed data structure · Perfect hashing · Bit parallelism · Vector instructions

**Mathematics Subject Classification** Theory of computation · Data compression · Information systems · Point lookups

## 1 Introduction

A perfect hash function (PHF) maps a set of  $N$  keys to the first  $M$  integers without collisions. If  $M = N$ , the hash function is called *minimal* perfect (MPHF) and is a bijection between the keys and the first  $N$  integers  $[N]$ . Minimal perfect hashing has many applications. For example, it can be used to implement static hash tables with guaranteed constant access time [28]. Storing only payload data in the hash table cells, we obtain an updatable retrieval data structure [45], and storing only fingerprints [7, 22], we obtain an approximate membership data structure. The hashes can also be used as small identifiers of the input keys [10], which are more efficient to deal with than large and complex keys. Finally, there is a range of applications in bioinformatics [1, 13–15] and text indexing [3, 5, 59].

**Related Work.** There is a lower bound of about  $N \log_2 e \approx 1.44N$  bits needed to represent an MPHF. There is also a theoretical construction matching this bound that runs in linear time and allows constant query time [32]. However, this construction does not work for realistic  $N < 2^{150}$  [11], so that (minimal) perfect hashing remains an interesting topic for algorithm engineering. A long sequence of previous work has developed a range of practical approaches with different space-time tradeoffs.

Many approaches first construct an outer hash function  $g : S \rightarrow [b]$  that partitions the input set  $S$  into small subsets  $S_1, S_2, \dots, S_b$  of sizes  $s_1 \approx s_2 \approx \dots \approx s_b$  and then construct a perfect hash function  $h_i : S_i \rightarrow \{1, \dots, s_i\}$  for  $i \in [b]$ . Given  $(s_i)_{i \in [b]}$ , or better yet the prefix sums  $p_i = s_1 + \dots + s_i$  for  $i \in [b]$ , an MPHF on  $S$  is given as  $x \mapsto p_{g(x)} + h_{g(x)}(x)$ .

On the one hand, there are holistic methods where such a partitioning step is not essential (though still possibly useful). These are (so far) all a constant factor away from the space lower bound (e.g. [6, 13, 39, 45]). One of the most space-efficient approaches among these is *SicHash* [39] that maps  $N$  keys to  $N(1 + \varepsilon)$  unique table entries using a generalization of cuckoo hashing [24, 49]. The choice of hash function for each key is then stored in a retrieval data structure and a small fallback data structure turns the constructed PHF into an MPHF. SicHash is not space-optimal partly because the cuckoo table tends to admit many valid placements of its keys, meaning a single input set is redundantly handled by many distinct states of the PHF data structure.

On the other hand, there are methods that use brute-force trial-and-error of hash functions on subsets of size  $n$  [4, 9, 21, 27, 50], which takes roughly  $e^n \approx 2.718^n$  trials. Hence, an aggressive partitioning step is required to obtain an acceptable overall running time. The previously most space efficient approach, *RecSplit* [21], is of this kind and recursively splits the input set into very small ( $n \approx 16$ ) *leaf* subsets. Surprisingly, when tuned thoroughly, this enables higher construction throughput than

the best holistic methods even when being fairly far away from the space lower bound [9].

**Contribution.** In this paper, we introduce *ShockHash* – Small, heavily overloaded cuckoo **hash** tables, which can be seen as an extreme version of *SicHash* where we use two hash functions for each key and retry construction until we can completely fill the cuckoo hash table. That way, we achieve an MPHF without an intermediate non-minimal PHF. In graph terminology, *ShockHash* repeatedly generates an  $n$ -edge random graph where each key corresponds to one edge, connecting the candidate positions of the key. The table can be filled if and only if the graph is a *pseudoforest* – a graph where no component contains more edges than nodes. While the *ShockHash* idea is straightforward in principle, we can *prove* that when using binary cuckoo hashing with two choices (and thus 1-bit retrieval) there is only an insignificant amount of redundancy. Therefore, *ShockHash* approaches the information theoretic lower bound for large  $n$  and has running time  $(e/2)^n \cdot \text{poly}(n) \approx 1.359^n$  (nearly a factor  $2^n$  faster than brute-force).

In *bipartite* *ShockHash*, further exponential improvements are possible. Instead of using a pair of fresh hash functions for each construction attempt, we build a growing pool of hash functions and consider all pairs that can be formed from this pool. Also, we let the two hash functions hash to disjoint ranges, meaning we effectively sample a bipartite graph where each edge has one endpoint in both partitions. In this bipartite setting, the hash functions of both partitions need to be *individually* surjective. We can therefore filter the set of candidate hash functions in each partition individually – before testing all combinations. This improves the construction time by an additional exponential factor, to about  $1.166^n \cdot \text{poly}(n)$ .

**Evaluation.** Still being an exponential time algorithm, we use *ShockHash* as a building block after partitioning the input. We obtain *ShockHash-RS* by using *ShockHash* instead of brute-force as a base case within the *RecSplit* framework. Though there is a small penalty in query time due to the additional access to a retrieval data structure, *ShockHash-RS* construction is about two orders of magnitude faster than tuned *RecSplit* [9] for space efficient configurations, using the same architecture. Bipartite *ShockHash-RS* improves this by a factor of 20 again. An important step in this harmonization of theory and practice is the observation that only an exponentially small fraction of the hash functions tried by *ShockHash* require the construction of a cuckoo hash table. The other cases can be covered with a simple bit-parallel filter that checks whether all entries of the cuckoo table are hit by some key. This removes much of the time overhead which made brute-force seemingly superior [9].

On large instances and for a space consumption of 1.56 bits per key, the most space efficient competitor, tuned *RecSplit* [9], requires  $137\mu\text{s}$  per key. In a similar amount of time, 1.499 bits per key can be achieved by massive parallelization of the approach using a GPU [9]. Bipartite *ShockHash-RS* now achieves a new record of just 1.489 bits per key with a similar construction time, while using only a single CPU thread (and efficient GPU parallelization seems not too difficult in future work).

We also demonstrate that *ShockHash* is useful outside the *RecSplit* framework. When using  $k$ -perfect hashing for partitioning the input, we obtain *ShockHash-Flat*, which achieves a space usage similar to the most space efficient competitors. At the same time, it is significantly faster to construct and reduces the query time by about

30%, which brings it closer to the query performance of way less space efficient approaches.

**Outline.** In Section 2, we explain preliminary ingredients needed to understand ShockHash, and in Section 3, we discuss related work. We then explain the ShockHash algorithm in Section 4 and a bipartite variant in Section 5. In Section 6, we then analyze both approaches. To make ShockHash feasible for large input sizes, we show how to use it as a building block in different partitioning schemes in Section 7. In Section 8, we give additional variants and refinements that improve the construction in practice. We conduct detailed experiments in Section 9 and conclude the paper in Section 10.

## 2 Preliminaries

In this section, we explain basic ingredients of ShockHash. This also includes the two perfect hash function constructions SicHash [39] and RecSplit [21] that ShockHash-RS is based on. Finally, we describe pairing functions that we use to encode the seeds in bipartite ShockHash.

**Cuckoo Hashing.** Cuckoo Hashing [49] is a well known approach to handle collisions in hash tables. Each object gets two candidate cells via two hash functions. A query operation looks at the two cells and searches for the object. If an insertion operation tries to insert an object into a cell that is already full, the object already stored in the cell is taken out and recursively inserted using its other candidate position. Cuckoo hashing can be extended to use more than two hash functions [24], or cells with more than one object in them [18]. In this paper, we are only interested in the basic version with two hash functions and one object per cell. The load threshold of a cuckoo hash table [25, 26, 41] is the percentage of cells that can be filled before insertion likely fails. For cuckoo hashing with two candidate cells, the load threshold is  $c = 0.5$ . Despite this, in ShockHash we consider cuckoo hash tables that are filled completely using many retries.

**Cuckoo Graph and Pseudoforests.** Cuckoo hashing can be modeled as a random graph  $G$ , where each node represents a table cell and each edge represents one object, connecting its two candidate cells. It is easy to see that a cuckoo hash table can be constructed successfully if and only if the edges of  $G$  can be directed such that the indegree of each node is  $\leq 1$ . In the following, we call this a 1-orientation. A 1-orientation exists if and only if  $G$  is a *pseudoforest*, i.e. every connected component of  $G$  is a pseudotree. A pseudotree is either a tree or a cycle with trees branching from it. A way to check whether a graph is a pseudoforest is to check whether each component contains at most as many edges as nodes.

**Retrieval Data Structures.** For a given set  $S$  of  $N$  keys and an integer  $r$ , a retrieval data structure (or *static function* data structure) stores a function  $f : S \rightarrow \{0, 1\}^r$  that maps each key to a specific  $r$ -bit value. Because it may return arbitrary values for keys not in  $S$ , it is possible to represent the function without representing  $S$  itself. Representing a retrieval data structure needs at least  $rN$  bits of space and there are practical data structures that need  $rN + o(rN)$  bits allowing linear construction time and constant query time. In particular, for  $r = 1$ , *Bumped Ribbon Retrieval (BuRR)* [19] reduces function evaluation to ANDing a hash function value with a segment of a

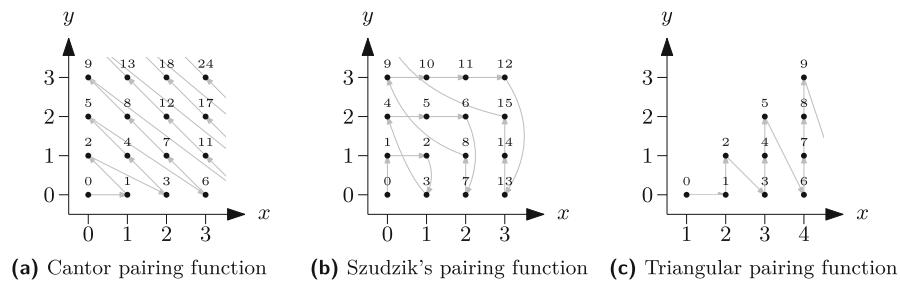
precomputed table and reporting the parity of the result. This table can be determined by solving a nearly diagonal system of linear equations (a “ribbon”). In practice, BuRR has a space overhead below 1%.

**SicHash.** Small irregular cuckoo tables for perfect Hashing [39] constructs perfect hash functions through cuckoo hashing. It constructs a cuckoo hash table and then uses a retrieval data structure to store which of the hash function choices was finally used for each key. The main innovation of SicHash is using a careful mix of 1–3 bit retrieval data structures, corresponding to 2, 4, or 8 choices for the keys. It achieves a favorable space-performance tradeoff when being allowed 2–3 bits of space per key. It cannot go below this because using only 1-bit retrieval seems to lead far from minimality while using 2 or more bits for retrieval allows redundant choices that cannot achieve space-optimality. SicHash achieves a rather limited gain in space efficiency by *overloading* the table beyond the load thresholds and trying multiple hash functions. This mainly exploits the variance in the number of keys that can fit. SicHash leaves the success probability of constructing overloaded tables as an open question. ShockHash drives the idea of overloading to its extreme and gives a formal analysis for this case.

**RecSplit.** RecSplit [21] is a minimal perfect hash function that is mainly focused on space efficiency. First, all keys are hashed to buckets of constant expected size  $b$ . A bucket’s set of keys is partitioned into different subsets recursively in a tree-like structure by *splitting* hash functions that are chosen by brute-force to ensure that all leaves except the last have size exactly  $n$  (in Ref. [21], the leaf size is called  $\ell$ ). Within the leaves, RecSplit then performs brute-force search for a minimal perfect hash function (also called *bijection*). The tree structure is based only on the size of the current bucket. This makes it possible to store only the seed values for the hash functions without storing structural information. Apart from encoding overheads for the seeds, this makes RecSplit information theoretically optimal within a bucket. The number of child nodes (*fanout*) in the two lowest levels is selected such that the amount of brute-force work is balanced between splittings and bijections.

There also is a parallel implementation using multi-threading and SIMD instructions or the GPU [9]. The paper also 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 up 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 find a way to cyclically rotate the second word, such that the empty positions left by the first set are filled by the positions of the second set. The paper shows that each rotation has a similar success probability as a new hash function seed, so it is a way to quickly evaluate additional hash function seeds.

**Pairing Functions.** A pairing function encodes two natural numbers in a single natural number. More precisely, a pairing function is a bijection between the grid  $\mathbb{N}_0^2$  and  $\mathbb{N}_0$ . We are interested in pairing functions that can be inverted efficiently. The most popular pairing function is the Cantor pairing function, which enumerates the 2D grid diagonally (see Figure 1a). It can be calculated by  $\text{pair}_c(x, y) = (x + y)(x + y + 1)/2 + y$ . Another pairing function is the one by Szudzik [56], which enumerates the 2D grid following the edges of a square (see Figure 1b). The pairing function can be calculated by  $\text{pair}_s(x, y) = y^2 + x$  if  $x \leq y$  and  $\text{pair}_s(x, y) = x^2 + x + y$  otherwise.



**Fig. 1** Illustrations of different pairing functions

In this paper, we require a function that enumerates only those coordinates of the 2D grid with  $x > y$ . We will still call it a pairing function in slight abuse of traditional terminology. Our *triangular pairing function* (see Figure 1c) can be calculated by  $\text{pair}_t(x, y) = x(x - 1)/2 + y$  with the intuition stemming from the *Little Gauss* formula. The basic idea for inverting our function  $(x', y') = \text{pair}_t^{-1}(z)$  is to set  $y = 0$  in the definition and solve for  $x$ . This gives  $x' = \lfloor 1/2 + \sqrt{1/4 + 2z} \rfloor$  and  $y' = z - \text{pair}_t(x', 0)$ . In our bipartite implementation, we use both our triangular pairing function and Szudzik's pairing function, depending on the distribution of the numbers we want to encode.

While pairing uses only integer operations, all three pairing functions rely on the square root operation and rounding for inverting. This means that inverting the functions in practice can lead to problems due to floating point inaccuracies. Whether inverting  $z$  succeeded can easily be checked by verifying that  $\text{pair}(x', y') = z$ . In our implementation, we check invertibility at construction time, so we do not get a run-time overhead during queries.

### 3 More Related Work

In addition to RecSplit and ShockHash, which we describe in the preliminaries, there is a range of other minimal perfect hash functions. We give an overview over the approaches in the following paragraphs.

**Hash-and-Displace.** Perfect hashing with Hash-and-Displace [4, 27, 50] allows fast queries and asymptotically optimal space consumption. Each key  $x$  is first hashed to a small bucket  $b(x)$  of keys. For each bucket  $b$ , an index  $i(b)$  of a hash function  $f_{i(b)}$  is stored such  $x \mapsto f_{i(b(x))}(x)$  is an injective function. For a particular bucket, this index is searched in a brute-force way. To accelerate the search, buckets are first sorted by their size. FCH [27] uses an asymmetric bucket assignment: it hashes 60% of the keys to 30% of the buckets. This is efficient because when placing the largest (and therefore hardest to place) buckets, most positions are still available. CHD [4] uses buckets of expected equal size but compresses the stored hash function seeds. PTHash [50] combines the ideas by using an asymmetric bucket assignment and storing the seeds in compressed form.

**Fingerprinting.** Perfect hashing through fingerprinting [13, 45] hashes the  $N$  keys to  $\gamma N$  positions using an ordinary hash function, where  $\gamma$  is a tuning parameter. The most space efficient choice  $\gamma = 1$  leads to a space consumption of  $e$  (not  $\log_2 e$ ) bits per key. A bit vector of length  $\gamma N$  indicates positions to which exactly one key was mapped. Keys that caused collisions are handled recursively in another layer of the same data structure. At query time, when a key is the only one mapping to its location, a rank operation on the bit vector gives the MPH value. Publicly available implementations include BBHash [42] and the significantly faster FMPH [6]. FMPHGO [6] combines the idea with a few brute-force tries to select a hash function that causes fewer collisions.

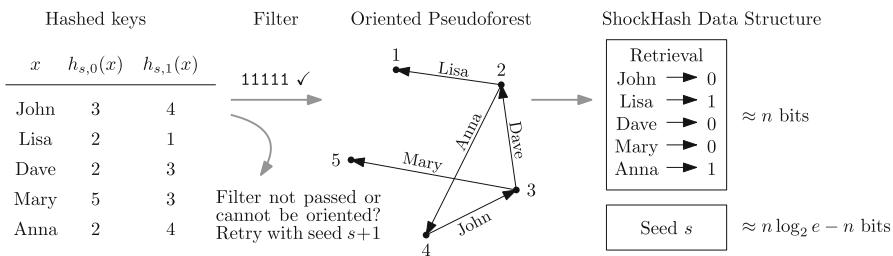
**Table Lookup.** A tempting way to replace expensive brute-force search is precomputation of solutions with subsequent table lookup – a standard technique used in many compressed data structures. For a rough idea, suppose for a subproblem with  $n$  keys, we first map them injectively to a range of size  $U' \in \Omega(n^2)$  using an intermediate hash function (less would lead to collisions – birthday paradox). Then, using a lookup table of size  $2^{U'}$ , we can find precomputed perfect hash functions in constant time. However, polynomial running time limits the subproblem size to  $n \in \mathcal{O}(\sqrt{\log N})$ , where  $N$  is the size of the overall input set. Computing concrete values for realistic values of  $N$ , one gets subproblem size much smaller than what can be easily handled even with plain RecSplit. Nevertheless, Hagerup and Tholey [32] develop a similar idea to a comprehensive theoretical solution of the perfect hashing problem yielding linear construction time, constant query time, and space  $1 + o(1)$  times the lower bound. However, this method is not even well-defined for  $N < 2^{150}$  [11]. A variant of RecSplit with rotation fitting [19] uses lookup tables of size  $2^n$  to find feasible rotations in constant time. Unfortunately, this turns out to be slower than trying all rotations directly.

## 4 ShockHash

We now introduce the main idea of this paper, ShockHash. The asymptotic load threshold of a binary cuckoo hash table is  $c = 0.5$  (see Section 2), so the success probability of constructing a table with  $n$  cells and more than  $n/2$  keys tends to zero. ShockHash overloads a cuckoo hash table far beyond its asymptotic load threshold – it inserts  $n$  keys into a binary cuckoo hash table of size  $n$ . As we will see in Theorem 8, the construction succeeds after  $(e/2)^n \text{poly}(n)$  tries in expectation. We then record the successful seed

$$s = \min\{s \in \mathbb{N} \mid \exists f \in \{0, 1\}^S : x \mapsto h_{s, f(x)}(x) \text{ is MPHF}\} \quad (h_{i,0})_{i \in \mathbb{N}}, (h_{i,1})_{i \in \mathbb{N}} : S \rightarrow [n]$$

and a successful choice  $f$  between the two candidate positions of each key. The seed needs  $n \cdot \log(e/2) + o(n) \approx 0.44n + o(n)$  bits in expectation using Golomb-Rice codes [30, 51]. The choices are stored in a 1-bit retrieval data structure, requiring  $n + o(n)$  bits (see Section 2). This means that the majority of the MPH description is not stored in the seed, like with the brute-force construction, but in the retrieval data structure. A query for key  $x$  retrieves  $f(x)$  from the retrieval data structure and returns  $h_{s, f(x)}(x)$ . Figure 2 gives an illustration of the ShockHash construction.



**Fig. 2** Illustration of the ShockHash construction. Functions  $h_0$  and  $h_1$  are randomly sampled hash functions using a seed  $s$ . Here,  $s$  is a seed value where the resulting graph is a pseudotree. During construction, many seeds need to be tried

The beauty of ShockHash is that it can check  $2^n$  different possible hash functions (determined by the  $2^n$  different functions representable by the retrieval data structure) in  $\mathcal{O}(n)$  time. This enables significantly faster construction than brute-force while still consuming the same amount of space up to lower order terms. As discussed in Section 2, a seed leads to a successful cuckoo hash table construction if and only if the corresponding random (multi)graph with edges  $\{(h_{s,0}(x), h_{s,1}(x)) \mid x \in S\}$  forms a pseudoforest. Each component of size  $c$  is a pseudotree if and only if it contains no more than  $c$  edges. This can be checked in linear time using connected components algorithms, or in close to linear time using an incremental construction of a cuckoo hash table. However, compared to the simple bit-parallel perfectness test of brute-force [21], each individual check is slower by a large constant factor. In the following paragraph, we discuss a way to address this bottleneck.

**Filter by Bit Mask.** To reduce the time spent checking if a graph is a pseudoforest, we use a filter to quickly reject most seeds. More specifically, we reject seeds for which some table cell is not a candidate position of any of the keys. If there is such a cell, we already know that cuckoo hashing cannot succeed and we can skip the full test. Otherwise, cuckoo hashing might succeed. The filter can be implemented using simple shift and comparison operations. Also, the filter can use registers, in contrast to the more complex full construction. It is one of the main ingredients making ShockHash practical and is easily proven to be very effective:

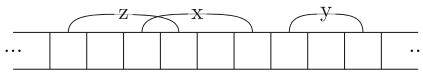
**Lemma 1** *The probability for a seed to pass the filter, i.e. for every table cell to be hit by at least one key, is at most  $(1 - e^{-2} + o(1))^n \approx 0.864^n$ .*

**Proof** Let  $X_i$  denote the number of times that cell  $i \in [n]$  is hit. Then  $(X_1, \dots, X_n)$  follows a multinomial distribution. The variables  $X_1, \dots, X_n$  are *negatively associated* in the sense introduced in [35] and satisfy

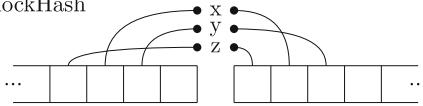
$$\Pr(\forall i \in [n] : X_i \geq 1) \leq \prod_{i=1}^n \Pr(X_i \geq 1),$$

the intuition being that since the sum  $X_1 + \dots + X_n = 2n$  is fixed, the events  $\{X_i \geq 1\}$  for  $i \in [n]$  are less likely to co-occur compared to corresponding independent events.

## ShockHash



## Bipartite ShockHash



```

repeat
  sample  $h_0, h_1 : S \rightarrow [n]$ 
  if  $\{(h_0(x), h_1(x)) \mid x \in S\}$  is pseudoforest
  return Naive
repeat without filtering
  sample  $h_0, h_1 : S \rightarrow [n/2]$ 
  if  $h_0$  surjective and  $h_1$  surjective
    if  $\{(h_0(x), h_1(x) + \frac{n}{2}) \mid x \in S\}$ 
      is pseudoforest
    return
  
```

**Fig. 3** ShockHash and bipartite ShockHash. The pseudocode illustrates the overall idea but does not lead to any performance improvements yet

Since  $X_i \sim \text{Bin}(2n, \frac{1}{n})$  for all  $i \in [n]$  we have

$$\Pr(X_i \geq 1) = 1 - (1 - \frac{1}{n})^{2n} = 1 - e^{-2} + o(1) \approx 0.864$$

and the claim follows.  $\square$

A more careful analysis [58] reveals that the probability to pass the filter is around  $b^n$  where  $b = 2e^\lambda/(\lambda e^2) \approx 0.836$  and where  $\lambda \approx 1.597$  is the solution to  $2 = \lambda/(1 - e^{-\lambda})$ .

**Enhancements.** In Section 8, we explain additional enhancements that improve the construction performance in practice. This includes applying the idea of *rotation fitting* [9] to ShockHash, as well as faster orientability checks.

## 5 Bipartite ShockHash

Bipartite ShockHash is an extension of the ShockHash idea. It enables significantly faster construction compared to plain ShockHash. In turn, this enables more aggressive parameter choices, thereby leading to improved space-efficiency. While ShockHash samples random graphs, *bipartite* ShockHash now samples *bipartite* random graphs. Figure 3 gives an illustration and very simple pseudocode. In plain ShockHash, each edge is connected to two nodes using two independent hash functions. In bipartite ShockHash, the hash functions have a range of  $[n/2]$ , but we shift the hashes of one of the hash functions by  $n/2$ , meaning each edge gets one endpoint in  $[n/2]$  and one in  $n/2 + [n/2]$ . This is similar to the original implementation of cuckoo hashing using two independent hash tables [49]. The idea might sound not very helpful at first, but opens up several ways of pruning the search space. In the following, we assume that  $n$  is an even number. We give an extension to uneven numbers in Section 8.4.

**Filtering Seed Candidates.** We show in Section 6 that testing about  $(e/2)^n \approx 1.359^n$  pairs of hash functions is sufficient for plain ShockHash. The idea of bipartite ShockHash is that it is almost as good to consider roughly  $\sqrt{(e/2)^n} = (e/2)^{n/2}$  hash functions and all pairs that can be formed from them. This already improves the practical running time because fewer hash functions need to be evaluated. However, the asymptotic construction time is not improved much because we still need to test

---

```

Function construct( $S$ )
    surjectiveCandidates  $\leftarrow \emptyset$ 
    for  $s_0 = 0$  to  $\infty$ 
        if  $h_{s_0}$  is surjective on  $S$ 
            for  $s_1 \in$  surjectiveCandidates
                if  $\exists f \in \{0, 1\}^S : x \mapsto \frac{n}{2} \cdot f(x) + h_{s_f(x)}(x)$  is a bijection
                    return  $f$  as retrieval data structure,  $s_0, s_1$ 
            surjectiveCandidates  $\leftarrow$  surjectiveCandidates  $\cup \{s_0\}$ 
    Function evaluate( $x$ )
        return  $\frac{n}{2} \cdot f(x) + h_{s_f(x)}(x)$ 

```

**Fig. 4** Pseudocode of bipartite ShockHash

all combinations. A key realization is that in the bipartite case, a pair  $(h_0, h_1)$  of hash functions can only work if *both*  $h_0$  and  $h_1$  (both with range  $[n/2]$ ) are *individually* surjective. In the non-bipartite case, in contrast, the check was that  $h_0$  and  $h_1$  (both with range  $[n]$ ) *together* hit each position in  $[n]$  at least once. This means that we can filter the list of hash functions *before* pairing them up.

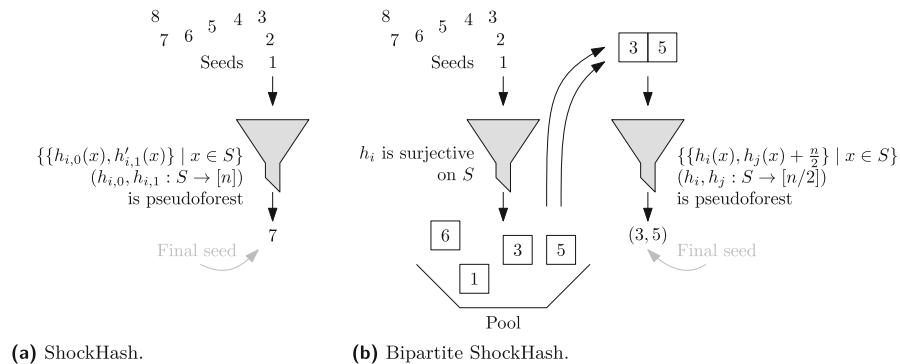
In each of the partitions, we look at  $n$  keys mapping to  $n/2$  positions. Similar to Lemma 1, the probability of passing the filter is  $0.836^{n/2}$  [58]. This suggests that if we pair up only the hash functions passing the filter then we will be considering at most  $((e/2)^{n/2} \cdot 0.836^{n/2})^2 \approx 1.136^n$  pairs. Refer to Section 6 for details.

**The Bipartite ShockHash Algorithm.** The following paragraph describes our new bipartite ShockHash algorithm. We maintain a pool of seed candidates that are surjective on  $[n/2]$ . To find a new candidate, we linearly check hash function seeds until we find a seed  $s_0$  that gives a surjective hash function. Given that new candidate, we try to combine it with all previous candidates  $s_1$  from the pool. More precisely, we check if the graph defined by the nodes  $[n]$  and the edges  $\{\{h_{s_0}(x), h_{s_1}(x) + n/2\} \mid x \in S\}$  is a pseudoforest. If it is a pseudoforest, we have found a perfect hash function. We only need to store the assignment from keys to their candidate hash function ( $h_{s_0}$  or  $h_{s_1}$ ) in a retrieval data structure, as well as the two seeds  $s_0$  and  $s_1$ . If the combination with none of the previous seed candidates leads to a pseudoforest, we add the new candidate  $s_0$  to the set of surjective candidates and search for the next one. Figure 4 gives a pseudocode for this algorithm and Figure 5 illustrates the idea of filtering hash functions before putting them in the pool.

Note that it does not matter which of the two hash functions we use for which partition of the graph. Switching the partitions just gives an isomorphic graph and does not influence orientability. We therefore always use the newly determined candidate directly and shift the old candidate by  $n/2$  to be in the second partition. Also, we neglect the possibility that a hash function combined with itself on both partitions leads to successful construction.<sup>1</sup> This allows us to store the two seeds  $s_0$  and  $s_1$ , knowing that  $s_1 < s_0$ . We do so in one integer using our triangular pairing function that we explain in Section 2. Note that the pairing function enumerates the seed pairs in exactly the same order that we test them in. Compared to storing two variable-length integers, pairing reduces constant overheads in the encoding.

---

<sup>1</sup> The function would have to map exactly two keys to each of the  $n/2$  positions, which happens with probability  $(\frac{n}{2} \cdot 2^{n-2}) \cdot (\frac{n}{2})^{n-2} = e^{-n} 2^{n/2} \text{poly}(n)$ .



**Fig. 5** Illustration of the filtering involved in ShockHash and bipartite ShockHash. The construction is complete if we find one final seed. ShockHash determines both hash functions from the same seed. Bipartite ShockHash uses independent seeds for the two hash functions and filters the seeds before combining them

**Enhancements.** In Section 8, we give additional enhancements that improve the construction performance significantly in practice. This includes other ways of coming up with a stream of hash function candidates, bit-parallel filtering, and support for uneven input sizes.

## 6 Analysis

In this section, we analyze the space usage and construction time of ShockHash. The main challenge is to lower bound the probability that a hash function seed enables successful construction of the heavily overloaded cuckoo hash table. First, in Section 6.1, we give a very simple analysis of the success probability of plain ShockHash. It is less tight than our more complex proof, but it is significantly shorter. In Section 6.2, we explain tools used in the analysis and prove small building blocks of the full proof. In Sections 6.3 and 6.4, we then analyze the success probability of plain and bipartite ShockHash, respectively. We then show in Section 6.5 that a pool containing about  $(\sqrt{e/2})^n$  hash function candidates is usually sufficient. Finally, we give the construction time and space consumption of ShockHash and bipartite ShockHash in Section 6.6. In the following we assume that a seed is given. We suppress it in notation.

It will be useful to consider the graph

$$G = ([n], \{\{h_0(x), h_1(x)\} \mid x \in S\}).$$

While similar to an Erdős-Renyi random graph,  $G$  may have self-loops<sup>2</sup> and multi-edges. Our model matches Model A in [29].

<sup>2</sup> Graphs with self-loops are easier to analyze here but we avoid them in practice for better performance.

## 6.1 A Simple Proof

First, in Theorem 2, we give a simple combinatorial argument showing that the probability for  $G$  to be a pseudotree is at least  $(e/2)^{-n} \sqrt{\pi/(2n)}$ . This lower bounds the probability of  $G$  being a pseudoforest consisting of potentially more than one tree. Section 6.3 then shows that the probability is at least  $(e/2)^{-n} \pi/e$ . Therefore, the simple argument is only a factor of  $\mathcal{O}(\sqrt{n})$  less tight than the much more complex proof.

**Theorem 2** *Let  $G$  be a multigraph with  $n$  nodes and  $n$  edges. The probability space underlying  $G$  is that of sampling  $2n$  nodes (with replacement) and creating an edge from the samples  $2i - 1$  and  $2i$  for each  $i \in [n]$ . Then the probability that  $G$  is a pseudotree is at least  $(e/2)^{-n} \sqrt{\pi/(2n)}$ .*

**Proof** For  $G$  to be a pseudotree it is sufficient (though not necessary) that the first  $n - 1$  created edges form a tree. There are  $n^{n-2}$  labeled  $n$ -node trees (Cayley's Formula [12]). Since the ordering of the edges and the order of the two samples forming an edge does not matter, each of the trees can be generated in  $2^{n-1}(n - 1)!$  ways. The last two samples can be anything, giving us  $n^2$  choices. By applying Stirling's approximation, namely

$$n! \in \left[ \left(\frac{n}{e}\right)^n \sqrt{2\pi n} \cdot e^{1/(12n+1)}, \left(\frac{n}{e}\right)^n \sqrt{2\pi n} \cdot e^{1/(12n)} \right],$$

we can show that the total probability to draw a pseudotree is at least

$$\frac{n^{n-2} 2^{n-1} (n - 1)! n^2}{n^{2n}} \geq \left(\frac{e}{2}\right)^{-n} \sqrt{\pi/(2n)}.$$

□

## 6.2 Tools

In this section, we explain tools that are needed later in the analysis, such as the configuration model and graph peeling. We start with proving two small lemmas that we use in the remaining analysis but are very generic in nature.

**Lemma 3** *Let  $X \in \mathbb{N}_0$  be a random variable. Then the probability that  $X$  is at least 1 is*

$$\Pr(X > 0) = \frac{\mathbb{E}(X)}{\mathbb{E}(X \mid X > 0)}.$$

**Proof** For any non-negative random variable  $X$  we can apply the law of total expectation to get  $\mathbb{E}(X) = \Pr(X = 0) \cdot \mathbb{E}(X \mid X = 0) + \Pr(X > 0) \cdot \mathbb{E}(X \mid X > 0) = \Pr(X > 0) \cdot \mathbb{E}(X \mid X > 0)$ . Rearranging this for  $\Pr(X > 0)$  yields the desired result. □

**Lemma 4** *For  $n, c \in \mathbb{N}$ , it holds that  $\binom{cn}{n} \leq (ec)^n$*

**Proof** We first upper bound all factors  $(cn - k)$  by  $cn$  and then apply Stirling's approximation.

$$\binom{cn}{n} = \frac{(cn)(cn-1)(cn-2)\dots(cn-n+1)}{n!} \leq \frac{(cn)^n}{n!} \leq \frac{(cn)^n}{\sqrt{2\pi n} (n/e)^n} \leq (ec)^n$$

□

**Configuration Model.** The *configuration model* [46] is a way to describe distributions of random graphs. In the model, we can fix the exact degree of every node in the graph by giving each node a number of *stubs* (*half-edges*). The graph is obtained by repeatedly sampling, uniformly at random, two unconnected stubs and connecting them. In other words, if we take one stub and look at its partner, all other stubs are equally likely.

**Graph Peeling.** In several sections of this paper, we are interested in *peeling* [33, 44, 57] graphs. For this, we iteratively take any node of degree 1 and remove it together with its corresponding edge. The process continues until all nodes have degree  $> 1$ . A graph is 1-orientable or a pseudoforest, if all nodes in the remaining graph have degree 2, meaning that the graph consists of only cycles.

**Graph Peeling in the Configuration Model.** To analyze the peeling process, it will be useful to reveal  $G$  in two steps. First the degree of each node is revealed by randomly distributing  $2n$  *stubs* among the  $n$  nodes. This yields a configuration model from which the edges are then obtained by randomly matching the stubs. The following lemma should clarify what exactly we need.

**Lemma 5** *Let  $x_1, \dots, x_{2n} \in [n]$  be independent and uniformly random. The graphs  $G_1, G_2, G_3$  defined in the following have the same distribution as  $G$ .*

1.  $G_1 = ([n], \{\{x_{2i-1}, x_{2i}\} \mid i \in [n]\})$ .
2.  $G_2 = ([n], \{\{x_i, x_j\} \mid \{i, j\} \in M\})$  where  $M$  is a uniformly random perfect matching of  $[2n]$ , i.e. a partition of  $[2n]$  into  $n$  sets of size 2.
3.  $G_3$  is defined like  $G_2$ , except that  $M$  is obtained in a sequence of  $n$  rounds. In each round an unmatched number  $i \in [2n]$  is chosen arbitrarily and matched to a distinct unmatched number  $j$ , chosen uniformly at random. The choice of  $i$  may depend on  $x_1, \dots, x_{2n}$  and on the set of numbers matched previously.

The reason for considering these alternative probability spaces for  $G$  is that they permit conditioning on partial information about  $G$  (such as its degree sequence implicit in  $x_1, \dots, x_{2n}$ ) but retaining a clean probability space for the remaining randomness.

**Proof** Compared to  $G$ , the definition of  $G_1$  simply collects the  $2n$  relevant hash values in a single list.<sup>3</sup> Concerning  $G_2$ , imagine that  $M$  is revealed first. Conditioned on  $M$ ,  $G_2$  is composed of  $n$  uniformly random edges like  $G_1$ . Concerning  $G_3$ , the key observation is that  $M$  is a uniformly random matching even if the number to be randomly matched in every round is chosen by an adversary. A formal proof could consider any arbitrary adversarial strategy and use induction. □

<sup>3</sup> Here, we assume that  $h_0$  and  $h_1$  are fully random hash functions and given for free, which is common in previous papers (Simple Uniform Hashing Assumption) [16, 17, 47, 48].

Therefore, when peeling in the configuration model, we can interleave the peeling process and the process of uncovering the sampled graph. To peel, we take a node with degree 1 and look at the other endpoint of its adjacent edge, which is uniformly distributed between all other stubs. If the node connected to it has degree 2, removing the edge gives a new degree-1 node that we can directly continue peeling. Otherwise, we have to start with a new node of degree 1 in a next iteration.

### 6.3 Success Probability in Plain ShockHash

In this section, we give the tighter analysis of the success probability of ShockHash.

Given two hash functions  $h_0, h_1 : S \rightarrow [n]$  and a function  $f : S \rightarrow \{0, 1\}$ , let  $\text{ori}(f)$  be the event that  $x \mapsto h_{f(x)}(x)$  is bijective. We are now interested in the probability that there exists such a function  $f$  that leads to a bijective function, namely  $\Pr(\exists f : \text{ori}(f))$ . There is a one-to-one correspondence between functions  $f$  with  $\text{ori}(f)$  and 1-orientations of  $G$ , i.e. ways of directing  $G$  such that each node has indegree at most 1.<sup>4</sup>

We write  $\text{PF}(G)$  for the event that  $G$  is a pseudoforest. As pointed out in Section 2:

$$\text{PF}(G) \Leftrightarrow \exists f : \text{ori}(f). \quad (1)$$

In our case with  $n$  nodes and  $n$  edges,  $\text{PF}(G)$  implies that  $G$  is a *maximal* pseudoforest, where every component is a pseudotree and not a tree. Note that a pseudotree that is not a tree admits precisely two 1-orientations because the unique cycle can be directed in two ways and all other edges must be directed away from the cycle. A useful observation is therefore

$$\text{PF}(G) \Rightarrow \#\{f : \text{ori}(f)\} = 2^{c(G)} \quad (2)$$

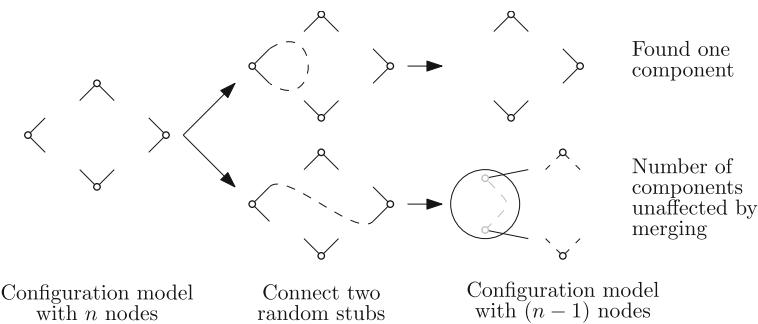
where  $c(G)$  is the number of connected components of  $G$ .

The basic idea of our proof is as follows. The probability that a random function is minimal perfect is  $e^{-n} \text{poly}(n)$  (see Lemma 7). Each of the  $2^n$  functions  $f : S \rightarrow \{0, 1\}$  has that chance of satisfying  $\text{ori}(f)$  and yielding an MPH. However, simply multiplying  $e^{-n} \text{poly}(n)$  by  $2^n$  does not necessarily yield an approximation for the probability that such an  $f$  exists. The key point here is that the  $2^n$  functions  $(x \mapsto h_{f(x)}(x))_{f \in \{0,1\}^S}$  determined by the  $2^n$  different options for  $f$  are correlated. If there are some graphs that permit many different 1-orientations, we may find many MPHs at once and the probability that there is at least one 1-orientation is reduced. This intuition is stated more formally in Lemma 3 using  $X = \#\{f : \text{ori}(f)\}$ :

$$\Pr(\exists f : \text{ori}(f)) = \frac{\mathbb{E}(\#\{f : \text{ori}(f)\})}{\mathbb{E}(\#\{f : \text{ori}(f)\} \mid \exists f : \text{ori}(f))}.$$

As discussed above, a key step in the analysis is to show that we usually find only a few MPHs at once. This amounts to analyzing the distribution of the number of

<sup>4</sup> Note that in our model, there are two ways of directing a self-loop.



**Fig. 6** Number of components in the configuration model

components in random maximal pseudoforests, which we do in Lemma 6. The main proof in Theorem 8 then formally bounds the probability that a random graph is a pseudoforest, juggling different probability spaces.

**Lemma 6** *Let  $G_n$  be the random graph sampled from the configuration model with  $n$  nodes of degree 2, i.e. the 2-regular graph obtained by randomly joining  $2n$  stubs that are evenly distributed among  $n$  nodes. Then the number  $c(G_n)$  of components of  $G$  satisfies  $\mathbb{E}(2^{c(G_n)}) \leq e \cdot \sqrt{2n}$ .*

We remark that a similar proof shows that  $\mathbb{E}(c(G_n)) \in \mathcal{O}(\log n)$ . Note also the similarity to the locker puzzle, which analyzes the length of the largest cycle in a random permutation [54].

**Proof** We will find a recurrence for  $d_n := \mathbb{E}(2^{c(G_n)})$ . Consider an arbitrary node  $v$  of  $G_n$  and one of the stubs at  $v$ . This stub forms an edge with some other stub. We have  $n-1$  other nodes, each with 2 stubs, and we have the second stub at  $v$ . Each of these  $2n-1$  stubs is matched with  $v$  with equal probability. Therefore, the probability that  $v$  has a self-loop is  $\frac{1}{2n-1}$ .

(1) Conditioned on  $v$  having a self-loop, we have found an isolated node. The distribution of the remaining graph is that of  $G_{n-1}$  and the conditional expectation of  $2^{c(G)}$  is therefore  $\mathbb{E}(2^{1+c(G_{n-1})}) = 2d_{n-1}$ .

(2) Now condition on the formed edge connecting  $v$  to  $w \neq v$ . We can now merge the nodes to a single one without affecting the number of components. The merged node inherits two unused stubs, one from  $v$  and one from  $w$ . The distribution of the remaining graph is that of  $G_{n-1}$ . Therefore, in this case, the conditional expectation of  $2^{c(G)}$  is simply  $d_{n-1}$ .

These two cases are illustrated in Figure 6 and lead us to the following recurrence:

$$d_n = \frac{1}{2n-1} 2d_{n-1} + \left(1 - \frac{1}{2n-1}\right) d_{n-1} = \left(1 + \frac{1}{2n-1}\right) d_{n-1}.$$

With the base case  $d_0 = 1$ , we can solve the recurrence and bound its value as follows, using that  $\ln(1 + x) \leq x$  for  $x \geq 0$  as well as  $H_n := \sum_{i=1}^n \frac{1}{i} \leq 1 + \ln n$ :

$$\begin{aligned} d_n &= \prod_{i=1}^n \left(1 + \frac{1}{2i-1}\right) = \exp\left(\sum_{i=1}^n \ln\left(1 + \frac{1}{2i-1}\right)\right) \leq \exp\left(\sum_{i=1}^n \frac{1}{2i-1}\right) \\ &= \exp\left(1 + \sum_{i=2}^n \frac{1}{2i-1}\right) = \exp\left(1 + \frac{1}{2} \sum_{i=2}^n \left(\frac{1}{2i-1} + \frac{1}{2i-1}\right)\right) \\ &\leq \exp\left(1 + \frac{1}{2} \sum_{i=2}^n \left(\frac{1}{2i-1} + \frac{1}{2i-2}\right)\right) = \exp\left(1 + \frac{1}{2} \sum_{i=2}^{2n-1} \frac{1}{i}\right) \\ &= \exp((1 + H_{2n-1})/2) \leq \exp(1 + \ln(2n)/2) \leq e \cdot \sqrt{2n}. \end{aligned}$$

□

Let us now re-state the known result proven in RecSplit [21], which bounds the probability that a random function is minimal perfect. We will use this later when proving Theorem 8.

**Lemma 7** [see [21]] *A random function  $h : S \rightarrow [n]$  on a set  $S$  of  $n$  keys, is minimal perfect (i.e. is a bijection) with probability  $e^{-n} \sqrt{2\pi n} \cdot (1 + o(1))$ .*

**Proof** Given  $S$ , there are  $n^n$  possible functions from  $S$  to  $[n]$  and  $n!$  of them are bijective. Therefore, the probability that a randomly selected function is minimal perfect is  $n!/n^n$ . The claim is obtained by applying Stirling's approximation. □

Finally, we can derive the success probability of the ShockHash search as follows.

**Theorem 8** *Let  $h_0, h_1 : S \rightarrow [n]$  be uniformly random functions. The probability that there exists  $f : S \rightarrow \{0, 1\}$  such that  $x \mapsto h_{f(x)}(x)$  is bijective is at least  $(e/2)^{-n} e^{-1} \sqrt{\pi}$ .*

**Proof** Recall our shorthand  $\text{ori}(f)$  for the event that  $x \mapsto h_{f(x)}(x)$  is bijective. As given in Lemma 3, we can calculate the success probability as follows.

$$\Pr(\exists f : \text{ori}(f)) = \frac{\mathbb{E}(\#\{f : \text{ori}(f)\})}{\mathbb{E}(\#\{f : \text{ori}(f)\} \mid \exists f : \text{ori}(f))}.$$

We will consider the numerator and denominator in turn.

**Numerator: Expectation.** Linearity of expectation (holding even for dependent variables) yields

$$\mathbb{E}(\#\{f : \text{ori}(f)\}) = \sum_f \Pr(\text{ori}(f)) = \sum_f \Pr(x \mapsto h_{f(x)}(x) \text{ is bijective on } S).$$

For any fixed  $f$ , the function  $x \mapsto h_{f(x)}(x)$  assigns independent random numbers to each  $x \in S$ , i.e. is a random function as considered in Lemma 7 and hence bijective

with probability  $e^{-n}\sqrt{2\pi n} \cdot (1 + o(1))$ . We therefore get

$$\mathbb{E}(\#\{f : \text{ori}(f)\}) \geq 2^n \cdot e^{-n}\sqrt{2\pi n}. \quad (3)$$

**Denominator: Conditional Expectation.** Using observations (1) and (2) we can shift our attention to the graph  $G$ :

$$\mathbb{E}(\#\{f : \text{ori}(f)\} \mid \exists f : \text{ori}(f)) = \mathbb{E}(2^{c(G)} \mid \text{PF}(G)).$$

By virtue of Lemma 5 we can moreover move to a configuration model à la  $G_3$ . We first reveal the locations  $x_1, \dots, x_{2n}$  of the  $2n$  stubs (hence the degree sequence of  $G_3$ ) and then consider the following peeling process (see Section 6.2) that reveals edges of  $G_3$  and simplifies  $G_3$  in a step-by-step fashion.

As long as there exists a node  $v$  with only one stub, firstly, match it to a random stub to form a corresponding edge  $\{v, w\}$  (consuming the two stubs) and, secondly, remove the node  $v$  and the newly formed edge  $\{v, w\}$ . These removals do not affect the number of components of the resulting graph (since  $v$  was connected to  $w$ ), nor whether the resulting graph is a pseudoforest (since the component of  $w$  lost one node and one edge).

Let  $n'$  be the number of nodes that remain after peeling and let  $G'$  be the graph obtained by matching the remaining stubs. As discussed we have  $\text{PF}(G_3) \Leftrightarrow \text{PF}(G')$  and  $c(G') = c(G_3)$ . Since the average degree of  $G_3$  is 2 and since we removed one node and one edge in every round, the average degree of  $G'$  is also 2. There are two cases.

(1) Some node of  $G'$  has degree 0. Then  $\neg\text{PF}(G')$  because some component of  $G'$  must have average degree  $> 2$ .

(2) No node of  $G'$  has degree 0. Since we ran the peeling process, there is also no node of  $G'$  with degree 1. Hence, every node of  $G'$  has degree 2. This makes  $G'$  a collection of cycles. In particular  $\text{PF}(G')$  holds. Moreover, the generation of  $G'$  is precisely the situation discussed in Lemma 6.

Because the two cases imply opposite results on  $G'$  being a pseudoforest, we know that  $\text{PF}(G')$  holds if and only if we arrive in Case 2. While we have no understanding of the distribution of  $n'$ , we can nevertheless compute:

$$\begin{aligned} \mathbb{E}(2^{c(G)} \mid \text{PF}(G)) &= \mathbb{E}(2^{c(G_3)} \mid \text{PF}(G_3)) = \mathbb{E}(2^{c(G')} \mid \text{PF}(G')) = \mathbb{E}(2^{c(G')} \mid \text{Case 2}) \\ &\leq \max_{1 \leq i \leq n} \mathbb{E}(2^{c(G')} \mid \text{Case 2 with } n' = i) \leq \max_{1 \leq i \leq n} e\sqrt{2i} = e\sqrt{2n}. \end{aligned} \quad (4)$$

**Putting the Observations Together.** Combining our bounds on the numerator (3) and the denominator (4) gives the final result

$$\Pr(\exists f : \text{ori}(f)) \geq 2^n e^{-n}\sqrt{2\pi n} / (e\sqrt{2n}) = (e/2)^{-n} e^{-1} \sqrt{\pi}.$$

□

## 6.4 Success Probability in Bipartite ShockHash

In the bipartite case, we can basically perform the same steps as in the non-bipartite version. For simplicity, we restrict ourselves to even  $n$  in the analysis. While our implementation does support uneven  $n$  (see Section 8.4), this complicates the analysis and can be largely avoided when ShockHash is integrated into a partitioning framework like RecSplit (see Section 7). In this section, we again suppress the seed of the hash functions. Let  $h_0, h_1 : S \rightarrow [n/2]$  be hash functions and  $f : S \rightarrow \{0, 1\}$  be a function that selects between the two hash functions.

We now look at the effect of testing all correlated choices of the function  $f$ . As argued in Section 6.3, we start with peeling the corresponding graph until there are no nodes with degree 1 left. Conditioned on the graph being a pseudoforest, this leaves us with a graph where each node has degree 2. The distribution of this graph is captured again by a configuration model (see Section 6.2), namely giving a random bipartite matching between the stubs. Additionally, remember that we started with a bipartite graph, so both partitions have the same size. Similar to Lemma 6, we can now show in Lemma 9 that the number of components  $c$  in the remaining graph satisfies  $\mathbb{E}(2^c) \leq e \cdot \sqrt{n}$ . Because the peeling process does not change the number of components, the same applies also to the original bipartite graph. This gives us a bound for the expected number of orientations of the graph, e.g., the number of different functions  $f$  that all make the hash function pair  $(h_0, h_1)$  minimal perfect.

**Lemma 9** *Let  $n$  be an even number, and let  $G_n$  be a random bipartite graph with  $n/2$  nodes in each partition, where all nodes have degree 2, sampled from the corresponding bipartite configuration model. Hence, the stubs from one partition are matched to the stubs of the other partition uniformly at random. Then the number  $c(G_n)$  of components of  $G_n$  satisfies  $\mathbb{E}(2^{c(G_n)}) \leq e \cdot \sqrt{n}$ .*

**Proof** We will find a recurrence for  $d_n := \mathbb{E}(2^{c(G_n)})$ . Consider an arbitrary node  $v$  from the first partition of  $G_n$  and one of the stubs at  $v$ . Because the graph is bipartite, this stub forms an edge to a node  $r$  from the other partition of the graph. The node  $r$  has a second stub that is connected back to the first partition. We now have  $n/2 - 1$  other nodes in the first partition, each with 2 stubs, and we have the second stub at  $v$ . Each of these  $n - 1$  stubs is matched with equal probability. Therefore, the probability that this edge closes a cycle is  $\frac{1}{n-1}$ .

(1) Conditioned on closing the cycle, the distribution of the remaining graph is that of  $G_{n-2}$  and the conditional expectation of  $2^{c(G_n)}$  is therefore  $\mathbb{E}(2^{1+c(G_{n-2})}) = 2d_{n-2}$ .

(2) Now condition on the edge not closing a cycle. We can now merge the three considered nodes to a single one without affecting the number of components. The merged node inherits two unused stubs, and the graph is now bipartite with  $n/2 - 1$  nodes in each partition. The distribution of the remaining graph therefore is that of  $G_{n-2}$ . Therefore, the conditional expectation of  $2^{c(G)}$  is simply  $d_{n-2}$ .

These two cases are similar to the non-bipartite case illustrated in Figure 6 and lead us to the following recurrence:

$$d_n = \frac{1}{n-1} 2d_{n-2} + \left(1 - \frac{1}{n-1}\right) d_{n-2} = \left(1 + \frac{1}{n-1}\right) d_{n-2}.$$

With the base case  $d_0 = 1$ , we can solve the recurrence and bound its value as follows, using that  $\ln(1 + x) \leq x$  for  $x \geq 0$  as well as  $H_n := \sum_{i=1}^n \frac{1}{i} \leq 1 + \ln n$ :

$$\begin{aligned}
 d_n &= \prod_{i=1}^{n/2} \left(1 + \frac{1}{2i-1}\right) = \exp \left( \sum_{i=1}^{n/2} \ln \left(1 + \frac{1}{2i-1}\right) \right) \leq \exp \left( \sum_{i=1}^{n/2} \frac{1}{2i-1} \right) \\
 &= \exp \left( 1 + \sum_{i=2}^{n/2} \frac{1}{2i-1} \right) = \exp \left( 1 + \frac{1}{2} \sum_{i=2}^{n/2} \left( \frac{1}{2i-1} + \frac{1}{2i-1} \right) \right) \\
 &\leq \exp \left( 1 + \frac{1}{2} \sum_{i=2}^{n/2} \left( \frac{1}{2i-1} + \frac{1}{2i-2} \right) \right) = \exp \left( 1 + \frac{1}{2} \sum_{i=2}^{n-1} \frac{1}{i} \right) \\
 &= \exp((1 + H_{n-1})/2) \leq \exp(1 + \ln(n)/2) \leq e \cdot \sqrt{n}.
 \end{aligned}$$

□

We can lower bound the success probability by applying Lemma 3 similarly as in Theorem 8. In contrast to Theorem 8, we now use  $\text{ori}(f)$  adapted for the bipartite case. Given two hash functions  $h_0, h_1 : S \rightarrow [n/2]$  and a function  $f : S \rightarrow \{0, 1\}$ , let  $\text{ori}(f)$  be the event that  $x \mapsto h_{f(x)}(x) + f(x) \cdot \frac{n}{2}$  is bijective. We write  $\text{PF}(h_0, h_1)$  for the event that the graph defined by the two hash functions  $h_0$  and  $h_1$  is a pseudoforest, and again get  $\text{PF}(h_0, h_1) \Leftrightarrow \exists f : \text{ori}(f)$ .

**Theorem 10** *Let  $h_0, h_1 : S \rightarrow [n/2]$  be uniformly random functions. The probability that there exists  $f : S \rightarrow \{0, 1\}$  such that  $x \mapsto h_{f(x)}(x) + f(x) \cdot \frac{n}{2}$  is bijective is at least  $(e/2)^{-n} \sqrt{n}/e$ .*

**Proof** All bipartite ShockHash functions have the form  $(x \mapsto h_{f(x)}(x) + f(x) \cdot (n/2))$ . While it is clear that the results of different  $x$  are independent, let us first justify why the function is uniform. For this, let  $c \in [n]$  be a constant,  $x$  an input value, and  $g : S \rightarrow \{0, 1\}$  a uniform random function. Then we get

$$\Pr(h_{g(x)}(x) + g(x) \cdot \frac{n}{2} = c) = \begin{cases} \Pr(h_0(x) = c \wedge g(x) = 0), & c < \frac{n}{2} \\ \Pr(h_1(x) = c - \frac{n}{2} \wedge g(x) = 1) & c \geq \frac{n}{2} \end{cases} = \frac{1}{n/2} \cdot \frac{1}{2} = \frac{1}{n}.$$

For uniform random functions  $g$  it holds that  $\mathbb{E}(\#\{f : \text{ori}(f)\}) = 2^n \cdot \Pr(\text{ori}(g))$ . Because we now also know that bipartite ShockHash with random  $g$  gives a uniform random function, we know that  $\Pr(\text{ori}(g))$  matches the probability that a random function is a bijection. Applying Lemma 7, this gives  $\mathbb{E}(\#\{f : \text{ori}(f)\}) \geq 2^n \cdot e^{-n} \sqrt{2\pi n}$ .

To determine the overall success probability, we can now continue similar to Theorem 8. Therefore, a random pair of hash functions  $h_0, h_1$  permits at least one valid placement  $f$  with at least the following probability.

$$\begin{aligned}
 \Pr(\exists f : \text{ori}(f)) &\stackrel{\text{Lem. 3}}{=} \frac{\mathbb{E}(\#\{f : \text{ori}(f)\})}{\mathbb{E}(\#\{f : \text{ori}(f)\} \mid \exists f : \text{ori}(f))} \geq \frac{2^n \cdot e^{-n} \sqrt{2\pi n}}{\mathbb{E}(\#\{f : \text{ori}(f)\} \mid \exists f : \text{ori}(f))} \\
 &\stackrel{\text{Lem. 9}}{\geq} e^{-n} \sqrt{2\pi n} \cdot 2^n / (e \cdot \sqrt{n}) = (e/2)^{-n} \sqrt{n}/e
 \end{aligned}$$

□

Our bound is a factor of  $\sqrt{2}$  better than with plain ShockHash, which reduces our bound on the expected space consumption by  $\log_2(\sqrt{2}) = 0.5$  bits. It might be possible to save a few additional bits in the retrieval data structure because  $f$  is known to map exactly half of the keys to 1, i.e., only a  $1/\Theta(\sqrt{n})$  fraction of all functions can occur as  $f$ . However, we do not consider this in more detail here.

## 6.5 Hash Function Pools

An important ingredient for making bipartite ShockHash so much more efficient than the plain version is the fact that we can combine hash functions from a pool of candidates. This makes it possible to filter the candidates *before* combining them. In the following section, we show why testing all combinations of two hash functions from a pool of candidates has a similar success probability as always sampling two fresh functions. For this analysis, let us form *two* pools of hash functions of size  $k = (e/2)^{n/2}$ . Note that this is slightly different to the construction explained before, which uses a single, growing pool. However, this does not influence the asymptotic construction time.<sup>5</sup> The first pool  $(l_i)_{i \in [k]}$  contains uniformly drawn hash functions  $l_1, \dots, l_k$  and the second pool  $(r_i)_{i \in [k]}$  contains uniformly drawn hash functions  $r_1, \dots, r_k$ . Our algorithm tests all combinations between two hash functions  $l_i$  and  $r_j$  ( $i, j \in [k]$ ) from the pools. We are therefore interested in the probability that the pools contain two hash functions that are *compatible*, meaning that their combined graph is a pseudoforest. For easier presentation, we suppress polynomial factors in this section and assume large  $n$ .

### 6.5.1 Notation and Overview

Let  $\mathcal{H} := [n/2]^S$  be the set of all (hash) functions from  $S$  to  $[n/2]$ . Sampling two hash functions  $l, r \sim \mathcal{U}(\mathcal{H})$  uniformly at random gives  $\Pr(\text{PF}(l, r)) = (e/2)^{-n}$ , as shown in Theorem 10 (ignoring polynomial factors). In the following, it will be useful to also consider  $\mathcal{H}_{\text{promisc}}$  containing *promiscuous* functions. The intuition is that a promiscuous function has a very large number of compatible partners, which is very unlikely. For reasons we will discuss later, we define  $\mathcal{H}_{\text{promisc}}$  as the set of functions that hit more than 90% of the hash values exactly two times:

$$\mathcal{H}_{\text{promisc}} := \left\{ f \in \mathcal{H} \mid \left| \left\{ i \in [n/2] \mid |f^{-1}(i)| = 2 \right\} \right| \geq 0.9n/2 \right\}.$$

For a specific function  $l$ , we define the set  $C(l)$  as all hash functions that are compatible with  $l$ . We also define  $C^*(l)$  as only those compatible functions that are not promiscuous:

$$C(l) := \{r \in \mathcal{H} \mid \text{PF}(l, r)\}, \quad C^*(l) := C(l) \setminus \mathcal{H}_{\text{promisc}}.$$

<sup>5</sup> One could view the two pools as one single pool of twice the size, where we only test a subset of the combinations. The time for testing more seed combinations would then be a constant factor larger. In practice, however, we test all combinations within the single pool, so it does not actually need to be twice as large.

For a set  $X \subseteq \mathcal{H}$  of hash functions, let  $\|X\| := |X|/|\mathcal{H}|$ . This is the probability that a randomly sampled hash function is one of the hash functions in  $X$ . Then  $\|C(l)\|$  is the probability that a randomly sampled hash function is compatible with  $l$ . If  $l$  is a random variable, this probability is a random variable as well. The expected value of this variable for uniform random  $l$  is  $\mathbb{E}_{l \sim \mathcal{U}(\mathcal{H})}(\|C(l)\|) = \Pr_{l,r \sim \mathcal{U}(\mathcal{H})}(\text{PF}(l, r)) = (e/2)^{-n}$  (ignoring polynomial factors).

The main random variable we are interested in is  $Z := \|\bigcup_{i \in [k]} C(l_i)\|$ . It depends on the hash functions that we sampled for the pool.  $Z$  is the probability that a randomly sampled hash function is compatible with at least one function from our pool  $(l_i)_{i \in [k]}$ .

If  $Z$  was small, it would mean that the hash functions in our pool need a very specific set of partners. Therefore, in this case, it would be unlikely that one of the compatible partners was drawn for the pool  $(r_i)_{i \in [k]}$ . However, we will show that  $Z$  is large enough that  $(r_i)_{i \in [k]}$  likely contains a compatible partner. More specifically, we will show that  $Z$  is closely concentrated around  $(e/2)^{-n/2}$ . Since we have  $k = (e/2)^{n/2}$  hash functions in each pool, we get a constant probability that a compatible function is in  $(r_i)_{i \in [k]}$ .

We start our proof in Section 6.5.2, showing that it is unlikely that a hash function is promiscuous. We then show in Section 6.5.3 that functions  $\notin \mathcal{H}_{\text{promisc}}$  do not have too many compatible partners. This is a key ingredient for providing concentration bounds on  $Z$  in Section 6.5.4. Finally, we combine this with the pool  $(r_i)_{i \in [k]}$  in Section 6.5.5. Figure 7 gives an overview over the proof structure.

### 6.5.2 Promiscuous Hash Functions

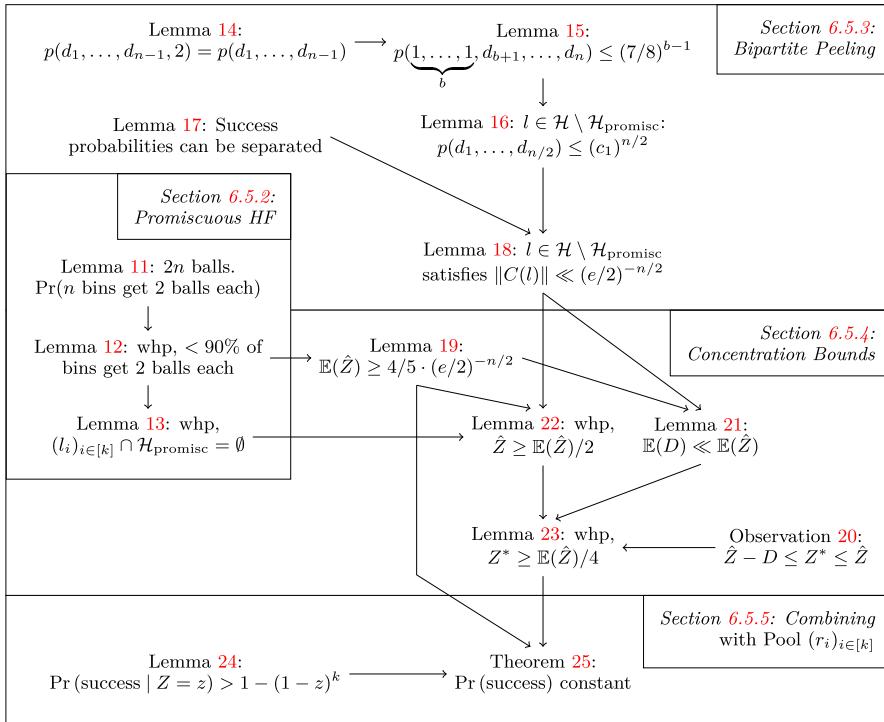
Before being able to give concentration bounds, we have to rule out a special case of hash functions with too many compatible partners. As stated before, we call these *promiscuous*. In this section, we show that these functions are very rare. To show this, we interpret the output values of our hash functions as a balls-into-bins process. In the following, we show a general property of balls-into-bins processes that we then later apply to our hash functions.

**Lemma 11** *When throwing  $2n$  balls into  $n$  bins independently and uniformly at random, the probability that each bin receives exactly two balls is  $p_n \leq 5\sqrt{n}(2e^{-2})^n$ .*

**Proof** To assign the balls to the bins such that each bin receives exactly two balls, we choose  $n$  subsets of size 2 from the set of  $2n$  balls. The number of ways we can do this is given by the multinomial coefficient  $\binom{2n}{2, 2, \dots, 2}$ . Each of these combinations has a probability of  $n^{-2n}$ . Using Stirling's approximation in the step annotated with  $*$ , this gives

$$\begin{aligned} p_n &= \binom{2n}{2, 2, \dots, 2} \cdot n^{-2n} = \underbrace{\frac{(2n)!}{2 \cdot 2 \cdot \dots \cdot 2}}_{n \text{ times}} n^{-2n} \stackrel{*}{\leq} \sqrt{2\pi 2n} \left(\frac{2n}{e}\right)^{2n} e^{\frac{1}{24n}} \cdot 2^{-n} \cdot n^{-2n} \\ &= \sqrt{4\pi n} \cdot (2e^{-2})^n \cdot e^{\frac{1}{24n}} \leq 5\sqrt{n}(2e^{-2})^n. \end{aligned}$$

□



**Fig. 7** Illustration of the proof structure showing the success probability of sampling from hash function pools instead of independent hash functions. Variables  $\hat{Z}$ ,  $D$  and  $Z^*$  are defined in Section 6.5.4. Functions  $p$  and  $q$  are defined in Section 6.5.3

**Lemma 12** *When throwing  $n$  balls into  $n/2$  bins independently and uniformly at random, the probability  $p_{90\%}$  that more than 90% of the bins receive exactly 2 balls is  $p_{90\%} \leq 0.66^n$ .*

**Proof** Let us consider a set  $A$  of  $0.9 \cdot n/2$  bins that will each receive exactly 2 balls each. We do not care which bins receive exactly 2 balls, so there are  $\binom{n/2}{0.9 \cdot n/2} = \binom{n/2}{0.1 \cdot n/2}$  ways of selecting the set  $A$ . The probability that exactly  $2|A|$  balls land in a bin in  $A$  is  $\binom{n}{0.1n} 0.1^{0.1n} 0.9^{0.9n}$ . Conditioned on this, the probability that these  $2|A|$  balls are evenly distributed among the  $|A|$  bins is  $p_{|A|} \leq 5\sqrt{0.9 \cdot n/2} (2e^{-2})^{0.9 \cdot n/2}$  by Lemma 11. Bringing this together we can bound  $p_{90\%}$  as follows, applying Lemma 4.

$$\begin{aligned}
 p_{90\%} &\leq \binom{n}{0.1n} 0.1^{0.1n} 0.9^{0.9n} \cdot \binom{n/2}{0.1 \cdot n/2} \cdot (2e^{-2})^{0.9 \cdot n/2} 5\sqrt{0.9 \cdot n/2} \\
 &\leq (10e)^{0.1n} \cdot 0.1^{0.1n} 0.9^{0.9n} \cdot (10e)^{0.1 \cdot n/2} \cdot (2e^{-2})^{0.9 \cdot n/2} 5\sqrt{0.45n} \\
 &= \left( (10e)^{0.1} \cdot 0.1^{0.1} 0.9^{0.9} \cdot (10e)^{0.05} \cdot (2e^{-2})^{0.45} \right)^n 5\sqrt{0.45n}
 \end{aligned}$$

$< 0.66^n$  for  $n$  large enough.  $\square$

**Lemma 13** *With a probability of  $1 - 0.77^n$ , none of the hash functions in the pool  $(l_i)_{i \in [k]}$  is promiscuous.*

**Proof** Each of our hash functions maps the  $n$  keys to  $n/2$  nodes in the graph. Because the hash functions are sampled at random, this can be modeled as a balls-into-bins process. We can apply Lemma 12, telling us that the probability that a function is promiscuous is  $p_{90\%} \leq 0.66^n$ . We sample a pool of  $k = (e/2)^{n/2}$  hash functions. Therefore, the probability that at least one of our hash functions is promiscuous can be upper bounded as follows using a union bound.

$$\Pr \left( \bigvee_{i \in [k]} l_i \in \mathcal{H}_{\text{promisc}} \right) \leq k \cdot p_{90\%} = ((e/2)^{1/2} \cdot 0.66)^n < 0.77^n$$

$\square$

### 6.5.3 Peeling Bipartite Graphs

For plain ShockHash, we have looked at the event  $\text{PF}(G)$  that the resulting graph  $G$  is a pseudoforest, and have given bounds for  $\Pr(\text{PF}(G))$ . In Section 6.2, we have shown that we can uncover the graph in two steps: we first reveal the degree of each node and then match these stubs at random. Staying with plain ShockHash for the moment, we now condition on the degrees of the nodes. For this, let  $(d_1, \dots, d_n)$  be the degrees of each node, and let  $p(d_1, \dots, d_n) := \Pr(\text{PF}(G) \mid \text{nodes of } G \text{ have degrees } d_1, \dots, d_n)$  be the conditioned success probability. Note that the order of the function arguments does not matter to the success probability because the stubs are matched randomly. Let  $\mathcal{B}_n$  be the distribution of balls in the balls-into-bins process with  $n$  balls and  $n/2$  bins. Before we give additional properties of  $p(d_1, \dots, d_n)$  in the following lemmas, let us look at a simple observation about the function:

$$\mathbb{E}_{(d_1, \dots, d_n) \sim \mathcal{B}_{2n}}(p(d_1, \dots, d_n)) = \Pr(\text{PF}(G)) \stackrel{\text{Thm. 8}}{=} (e/2)^{-n}. \quad (5)$$

The graph  $G$  is a pseudoforest if we can repeatedly peel away nodes of degree 1 and end up with a graph that consists of only nodes of degree 2 (forming cycles). When peeling, we repeatedly take a node of degree 1 and follow its edge to a random stub. There are now three possible outcomes.

- (1) The edge leads to a node with degree 1. Then we have found a component with more nodes than edges, meaning that the remaining graph cannot be a pseudoforest.
- (2) The edge leads to a node with degree  $\geq 3$ . Then we have peeled away a subtree of a component that is potentially still a pseudotree. We now have one less node of degree 1 in our graph. In that case, we continue the peeling process with another node of degree 1.

(3) The edge leads to a node with degree 2. Then it generates a new node with degree 1, and we can continue the peeling process with that same node.

**Lemma 14** *Nodes of degree 2 do not influence the success probability. More formally,  $p(d_1, \dots, d_i, 2) = p(d_1, \dots, d_i)$ .*

**Proof** If the peeling process arrives in case (3) where it connects to a node of degree 2, we can immediately continue peeling with that node. This always succeeds and does not cause an abort of the peeling process. Afterwards, we have one less node but did not influence the degrees of all other nodes. Therefore, the probabilities to arrive in the more interesting cases (1) and (2) stay the same, relative to each other.  $\square$

**Lemma 15** *If the graph corresponding to our hash function has  $b$  nodes of degree 1,  $p$  is exponentially small in  $b$ . More formally,  $p(\underbrace{1, \dots, 1}_b, d_{b+1}, \dots, d_n) \leq (7/8)^{b-1}$ .*

**Proof** During the peeling process, we can ignore all nodes of degree 2 because they do not influence the success probability (see Lemma 14). Let us therefore now look at the remaining nodes. If we connect to one of the nodes of degree 1, we have found a tree. This means that the construction fails because we need each component to be a pseudotree and not a tree. Because the average degree is 2, we have at most  $3b$  stubs at nodes with degree  $\geq 3$ . Therefore, the probability that we connect to a node of degree 1 (possibly indirectly through nodes of degree 2) is at most  $(b-1)/(4b)$ . We can now apply this iteratively until all nodes of degree 1 are peeled away. This gives the following probability that we never connect to a node of degree 1 (and therefore fail):

$$\begin{aligned} p(\underbrace{1, \dots, 1}_b, d_{b+1}, \dots, d_n) &\leq \prod_{j=1}^b \left(1 - \frac{j-1}{4j}\right) = \prod_{j=1}^b \left(\frac{3}{4} + \frac{1}{4j}\right) \\ &= \prod_{j=2}^b \left(\frac{3}{4} + \frac{1}{4j}\right) \leq \prod_{j=2}^b \left(\frac{7}{8}\right) = \left(\frac{7}{8}\right)^{b-1}. \end{aligned}$$

$\square$

We call a hash function promiscuous if more than 90% of the nodes in its corresponding (stub) graph have degree 2. This means that it cannot have too many nodes of degree 1. Promiscuous hash functions are rare, as we have seen in Section 6.5.2. In the following, we bound the number of compatible hash functions when our function is *not* promiscuous.

**Lemma 16** *Let  $l$  be a hash function from  $\mathcal{H} \setminus \mathcal{H}_{\text{promisc}}$  and  $(d_1, \dots, d_{n/2})$  be the corresponding degree sequence. Then  $p(d_1, \dots, d_{n/2}) \leq (c_1)^{n/2}$  where  $c_1$  is a constant  $\in (0, 1)$  and  $n$  large enough.*

**Proof** Let  $X_{\neq 2}$  be the set of nodes with degree  $\neq 2$ . Because the functions considered here are not promiscuous, we have  $|X_{\neq 2}| \geq 0.1 \cdot n/2$ . Because there are two times more stubs than nodes, the nodes in  $X_{\neq 2}$  receive 2 stubs on average. If one of the

nodes receives 0 stubs, the success probability is 0. Otherwise, at least half of the nodes in  $X_{\neq 2}$  have to receive exactly 1 stub to satisfy the average. Therefore,  $l$  has at least  $b = 0.05n/2$  nodes with degree 1. Applying Lemma 15 and selecting  $c_1 > (7/8)^{0.05} \in (0, 1)$  then concludes the proof.  $\square$

Now let us turn back to bipartite ShockHash. Here we have two random hash functions  $l$  and  $r$  that give us a bipartite graph  $G$ . Let  $(d_1, \dots, d_{n/2})$  and  $(d'_1, \dots, d'_{n/2})$  be the degrees of the nodes in the two partitions and remember that  $\mathcal{B}_n$  is the distribution of balls in the balls-into-bins process with  $n$  balls and  $n/2$  bins. Then we can define  $q$  as the following probability conditioned on the degree sequence of  $G$ :

$$\begin{aligned} q((d_1, \dots, d_{n/2}), (d'_1, \dots, d'_{n/2})) \\ := \Pr \left( \text{PF}(G) \mid G \text{ has degree sequence } (d_1, \dots, d_{n/2}), (d'_1, \dots, d'_{n/2}) \right) \\ = \Pr \left( \text{PF}(l, r) \mid l, r \text{ give degree sequence } (d_1, \dots, d_{n/2}), (d'_1, \dots, d'_{n/2}) \right) \end{aligned}$$

**Lemma 17** *For any  $d_1, \dots, d_{n/2}, d'_1, \dots, d'_{n/2}$  we have  $q((d_1, \dots, d_{n/2}), (d'_1, \dots, d'_{n/2})) \leq p(d_1, \dots, d_{n/2}) \cdot p(d'_1, \dots, d'_{n/2})$ .*

**Proof** Recall that the peeling process in the configuration model iteratively matches stubs that are the only remaining stub of their node. It can therefore be understood as the process of alternately removing a lonely stub and a random stub (and emitting a corresponding edge). Failure means that a randomly removed stub was lonely. If in the bipartite case we alternate between picking the lonely stubs on the left and on the right, then within each of the two partitions we are alternating between removing a lonely stub and a random stub. Therefore, we are basically running the peeling process within the two partitions separately. This suggests that the probability  $q((d_1, \dots, d_{n/2}), (d'_1, \dots, d'_{n/2}))$  that a bipartite graph is a pseudoforest when sampled from the configuration model with degree sequence  $((d_1, \dots, d_{n/2}), (d'_1, \dots, d'_{n/2}))$  is equal to the product of the probabilities  $p(d_1, \dots, d_{n/2})$  and  $p(d'_1, \dots, d'_{n/2})$  that two graphs are pseudoforests, namely those sampled from the configuration model with degree sequences  $(d_1, \dots, d_{n/2})$  and  $(d'_1, \dots, d'_{n/2})$ , respectively.

However, there is no strict equality due to a slight asymmetry: In the first partition we begin with removing a lonely stub while in the second partition we begin with removing a random stub. This slightly increases the failure probability for the second partition because there is always one lonely stub more when selecting a random stub than there would otherwise be. Note also that the first partition may run out of lonely stubs while the second partition has at least one lonely stub remaining. In that case we would remove a lonely stub from the second partition twice in a row, putting things back on track for the second partition.

Overall, the increase in failure probability for the second partition during (parts of) the process is accounted for by the “ $\leq$ ” in our statement.  $\square$

Both  $p(d_1, \dots, d_{n/2})$  and  $\|C(l)\|$  provide a way to rate the quality of a hash function candidate. However, there is a subtle but important difference between the two. While  $\|C(l)\|$  looks at the compatible partners in the bipartite case,  $p(d_1, \dots, d_{n/2})$  looks at

the graph when connected to itself. For example, a hash function candidate that leads to only nodes of degree 2 has  $p(2, \dots, 2) = 1$ . However, we get  $\|C(l)\| < 1$  because  $l$  is not compatible with partners that do not hit all output values.

We now formally connect the two probabilities  $p(d_1, \dots, d_{n/2})$  and  $\|C(l)\|$ . This gives a bound for the probability of drawing a compatible partner for a function in  $\mathcal{H} \setminus \mathcal{H}_{\text{promise}}$ .

**Lemma 18** *Each  $l \in \mathcal{H} \setminus \mathcal{H}_{\text{promise}}$  satisfies  $\|C(l)\| < (e/2)^{-n/2} \cdot (c_1)^{n/2}$  where  $c_1$  is a constant  $\in (0, 1)$  and  $n$  large enough.*

**Proof** Take any hash function  $l \in \mathcal{H} \setminus \mathcal{H}_{\text{promise}}$  and its corresponding degree sequence  $(d_1, \dots, d_{n/2})$ . As a reminder,  $\mathcal{B}_n$  is the distribution of balls in the balls-into-bins process with  $n$  balls and  $n/2$  bins. Then

$$\begin{aligned} \|C(l)\| &= \Pr(\text{PF}(l, r)) \\ &= \mathbb{E}_{(d'_1, \dots, d'_{n/2}) \sim \mathcal{B}_n} \left( \Pr \left( \text{PF}(l, r) \mid r \text{ has degree sequence } (d'_1, \dots, d'_{n/2}) \right) \right) \\ &= \mathbb{E}_{(d'_1, \dots, d'_{n/2}) \sim \mathcal{B}_n} \left( q((d_1, \dots, d_{n/2}), (d'_1, \dots, d'_{n/2})) \right) \\ &\stackrel{\text{Lem. 17}}{\leq} \mathbb{E}_{(d'_1, \dots, d'_{n/2}) \sim \mathcal{B}_n} \left( p(d_1, \dots, d_{n/2}) \cdot p(d'_1, \dots, d'_{n/2}) \right) \\ &= p(d_1, \dots, d_{n/2}) \cdot \mathbb{E}_{(d'_1, \dots, d'_{n/2}) \sim \mathcal{B}_n} \left( p(d'_1, \dots, d'_{n/2}) \right) \\ &\stackrel{(5)}{=} p(d_1, \dots, d_{n/2}) \cdot (e/2)^{-n/2} \stackrel{\text{Lem. 16}}{\leq} c_1^{n/2} \cdot (e/2)^{-n/2}. \end{aligned}$$

□

#### 6.5.4 Concentration Bounds

Remember variable  $Z = \|\bigcup_{i \in [k]} C(l_i)\|$ , which is the probability that a randomly selected function is compatible with a function in our pool. Also remember  $C^*(l) = C(l) \setminus \mathcal{H}_{\text{promise}}$ . Unfortunately, we cannot directly give concentration bounds on  $Z$  or even calculate  $\mathbb{E}(Z)$  because we do not know how using a pool of hash functions influences the probabilities. We therefore look at three additional random variables, defined as follows:

$$Z^* := \left\| \bigcup_{i \in [k]} C^*(l_i) \right\|, \quad \hat{Z} := \sum_{i \in [k]} \|C^*(l_i)\|, \quad D := \sum_{1 \leq i < j \leq k} \|C^*(l_i) \cap C^*(l_j)\|$$

$Z^*$  considers only partners in  $\mathcal{H} \setminus \mathcal{H}_{\text{promise}}$ .  $\hat{Z}$  is easier to calculate because it looks at each set separately. In the remainder of this section, we then determine bounds on  $Z^*$ ,  $\hat{Z}$ , and  $D$ , which we can later use to bound  $Z$ .

**Lemma 19** *Let  $Z^*$ ,  $\hat{Z}$  and  $D$  be as defined above. Then  $\mathbb{E}(\hat{Z}) > 4/5 \cdot (e/2)^{-n/2}$ .*

### Proof

$$\begin{aligned}
\mathbb{E}(\hat{Z}) &= \sum_{i \in [k]} \mathbb{E}(\|C^*(l_i)\|) \geq \sum_{i \in [k]} (\mathbb{E}(\|C(l_i)\|) - \|\mathcal{H}_{\text{promisc}}\|) \\
&= k \cdot ((e/2)^{-n} - \|\mathcal{H}_{\text{promisc}}\|) \stackrel{\text{Lem. 12}}{\geq} k \cdot \left( (e/2)^{-n} - \frac{0.66^n \cdot |\mathcal{H}|}{|\mathcal{H}|} \right) \\
&= (e/2)^{-n/2} - 0.66^n (e/2)^{n/2} = (1 - (0.66 \cdot e/2)^n) \cdot (e/2)^{-n/2} \\
&\geq (1 - 0.9^n) \cdot (e/2)^{-n/2} \geq 4/5 \cdot (e/2)^{-n/2} \text{ for } n \text{ large enough.}
\end{aligned}$$

□

**Observation 20** For  $\hat{Z}$ ,  $D$  and  $Z^*$  defined as above, it holds that  $\hat{Z} - D \leq Z^* \leq \hat{Z}$ .

**Proof** To show the bounds, we use the inclusion-exclusion principle, namely that for sets  $A$  and  $B$ ,  $|A \cup B| = |A| + |B| - |A \cap B| \leq |A| + |B|$ . We can give an upper bound for the variable  $Z^*$  by applying the inequality repeatedly using associativity of the union operation.

If we take  $\hat{Z}$  and subtract the sizes of all pairwise intersections, we get a lower bound. If all partners were compatible with at most two hash functions in our pool, subtracting the pairwise intersections would give  $Z^*$ . However, if a partner is compatible with more than two functions, this subtracts too much, therefore giving the lower bound  $\hat{Z} - D \leq Z^*$ . □

We can now show that  $\mathbb{E}(D)$  is exponentially smaller than  $\mathbb{E}(\hat{Z})$ . Intuitively, this means that  $\mathbb{E}(\hat{Z}) \approx \mathbb{E}(Z^*)$  by Observation 20. Our proof idea is to bound the intersections using the bound on  $\|C(l_i)\|$  shown in Lemma 18.

**Lemma 21** Let  $D$  and  $\hat{Z}$  be as defined above. Then  $\mathbb{E}(D) \leq (c_2)^{n/2} \cdot \mathbb{E}(\hat{Z})$ .

### Proof

$$\begin{aligned}
\mathbb{E}(D) &= \mathbb{E} \left( \sum_{1 \leq i < j \leq k} \|C^*(l_i) \cap C^*(l_j)\| \right) = \binom{k}{2} \mathbb{E}_{l_1, l_2 \sim \mathcal{U}(\mathcal{H})} (\|C^*(l_1) \cap C^*(l_2)\|) \\
&= \binom{k}{2} \frac{1}{|\mathcal{H}|} \mathbb{E}_{l_1, l_2 \sim \mathcal{U}(\mathcal{H})} (|C^*(l_1) \cap C^*(l_2)|) \\
&= \binom{k}{2} \frac{1}{|\mathcal{H}|} \mathbb{E}_{l_1 \sim \mathcal{U}(\mathcal{H})} (\mathbb{E}_{l_2 \sim \mathcal{U}(\mathcal{H})} (|C^*(l_1) \cap C^*(l_2)|)) \\
&= \binom{k}{2} \frac{1}{|\mathcal{H}|} \mathbb{E}_{l_1 \sim \mathcal{U}(\mathcal{H})} \left( \sum_{r \in C^*(l_1)} \Pr_{l_2 \sim \mathcal{U}(\mathcal{H})} (r \in C^*(l_2)) \right) \\
&\leq \binom{k}{2} \frac{1}{|\mathcal{H}|} \mathbb{E}_{l_1 \sim \mathcal{U}(\mathcal{H})} \left( \sum_{r \in C^*(l_1)} \Pr_{l_2 \sim \mathcal{U}(\mathcal{H})} (l_2 \in C(r)) \right)
\end{aligned}$$

$$\begin{aligned}
&= \binom{k}{2} \frac{1}{|\mathcal{H}|} \mathbb{E}_{l_1 \sim \mathcal{U}(\mathcal{H})} \left( \sum_{r \in C^*(l_1)} \|C(r)\| \right) \\
&\stackrel{\text{Lem. 18}}{\leq} \binom{k}{2} \frac{1}{|\mathcal{H}|} \mathbb{E}_{l_1 \sim \mathcal{U}(\mathcal{H})} \left( |C^*(l_1)| \cdot (e/2)^{-n/2} \cdot (c_1)^{n/2} \right) \\
&\leq (e/2)^n \mathbb{E}_{l_1 \sim \mathcal{U}(\mathcal{H})} (\|C(l_1)\|) \cdot (e/2)^{-n/2} \cdot (c_1)^{n/2} = (c_1)^{n/2} \cdot (e/2)^{-n/2} \\
&\stackrel{\text{Lem. 19}}{\leq} (c_2)^{n/2} \cdot \mathbb{E}(\hat{Z}) \text{ for some } c_2 \in (0, 1) \text{ and } n \text{ large enough.}
\end{aligned}$$

□

Let us now show that  $\hat{Z}$  does not get much smaller than its expected value.

**Lemma 22** *Let  $Z^*$ ,  $\hat{Z}$  and  $D$  be as defined above. Then  $\Pr(\hat{Z} \geq \mathbb{E}(\hat{Z})/2) > 1 - (c_3)^n$  where  $c_3$  is a constant  $\in (0, 1)$  and  $n$  large enough.*

**Proof** The Bernstein inequality [8] states that for independent and zero-mean random variables  $V_i$  with  $|V_i| \leq M$ , it holds that:

$$\Pr \left( \sum_{i=1}^n V_i \geq t \right) \leq \exp \left( -\frac{\frac{1}{2}t^2}{\sum_{i=1}^n \mathbb{E}(V_i^2) + \frac{1}{3}Mt} \right)$$

To apply the inequality, we now center our variables  $\|C^*(l_i)\|$  and mirror them around the value 0, giving us  $V_i = \mathbb{E}(\|C^*(l_i)\|) - \|C^*(l_i)\|$ . Centering only makes the maximum smaller. Through Lemma 13, we can assume that none of our functions  $l_i$  is promiscuous. This can be formalized by increasing  $c_3$  accordingly. Therefore, we get  $\max_{i \in [k]} (V_i) \leq \max(\|C^*(l_i)\|) \leq \max(\|C(l_i)\|) \leq (c_1)^{n/2} \cdot (e/2)^{-n/2} =: M$  through Lemma 18. Before we can get to the Bernstein inequality, we need another ingredient. Let us upper bound the value of  $\mathbb{E}((V_i)^2)$  as follows:

$$\begin{aligned}
\mathbb{E}((V_i)^2) &= \mathbb{E}((\|C^*(l_i)\| - \mathbb{E}(\|C^*(l_i)\|))^2) = \mathbb{E}(\|C^*(l_i)\|^2) - \mathbb{E}(\|C^*(l_i)\|)^2 \\
&\leq \mathbb{E}(\|C^*(l_i)\|^2) \leq \mathbb{E} \left( \max_{j \in [k]} \|C^*(l_j)\| \cdot \|C^*(l_i)\| \right) = \max_{j \in [k]} \|C^*(l_j)\| \cdot \mathbb{E}(\|C^*(l_i)\|) \\
&\stackrel{\text{Lem. 18}}{\leq} c_1^{n/2} \cdot (e/2)^{-n/2} \cdot (e/2)^{-n} = c_1^{n/2} \cdot (e/2)^{-n/2-n}
\end{aligned}$$

Setting  $t = 4/10 \cdot (e/2)^{-n/2}$  and applying the Bernstein inequality in the step annotated with  $*$ , we get:

$$\begin{aligned}
\Pr(\hat{Z} \leq \mathbb{E}(\hat{Z})/2) &= \Pr \left( \sum_{i \in [k]} \|C^*(l_i)\| \leq \mathbb{E}(\hat{Z})/2 \right) \\
&= \Pr \left( \sum_{i \in [k]} \mathbb{E}(\|C^*(l_i)\|) - \sum_{i \in [k]} V_i \leq \mathbb{E}(\hat{Z})/2 \right)
\end{aligned}$$

$$\begin{aligned}
&= \Pr \left( \mathbb{E}(\hat{Z}) - \sum_{i \in [k]} V_i \leq \mathbb{E}(\hat{Z})/2 \right) = \Pr \left( \sum_{i=1}^k V_i \geq \mathbb{E}(\hat{Z})/2 \right) \\
&\stackrel{\text{Lem. 19}}{\leq} \Pr \left( \sum_{i=1}^k V_i \geq 4/10 \cdot (e/2)^{-n/2} \right) \\
&\stackrel{*}{\leq} \exp \left( -\frac{\frac{1}{2} (4/10 \cdot (e/2)^{-n/2})^2}{k \cdot c_1^{n/2} \cdot (e/2)^{-n/2-n} + \frac{1}{3} \cdot (c_1^{n/2} \cdot (e/2)^{-n/2}) \cdot ((e/2)^{-n/2}/2)} \right) \\
&= \exp \left( -\frac{\frac{2}{25} (e/2)^{-n}}{c_1^{n/2} \cdot (e/2)^{-n} + \frac{1}{6} \cdot c_1^{n/2} \cdot (e/2)^{-n}} \right) \\
&= \exp \left( -12/175 \cdot c_1^{-n/2} \right) = \left( e^{-12/175} \right)^{\left( \left( c_1^{-1/2} \right)^n \right)} \leq (c_3)^n
\end{aligned}$$

for some  $c_3 \in (0, 1)$  and  $n$  large enough.  $\square$

Lemma 21 gives a bound on  $\mathbb{E}(D)$  and Lemma 22 gives a concentration bound on  $\hat{Z}$ . With these insights, we now give a bound on  $Z^*$  in terms of  $\mathbb{E}(\hat{Z})$ .

**Lemma 23** *Let  $Z^*$ ,  $\hat{Z}$  and  $D$  be as defined above. Then  $\Pr(Z^* < \mathbb{E}(\hat{Z})/4) \leq (c_4)^n$  where  $c_4$  is a constant  $\in (0, 1)$  and  $n$  large enough.*

**Proof** We can bound the probability that  $Z^*$  deviates too much from  $\mathbb{E}(\hat{Z})/4$  as follows. In the step annotated with  $*$ , we use the Markov inequality ( $\Pr(D \geq a) \leq \mathbb{E}(D)/a$ ).

$$\begin{aligned}
&\Pr \left( Z^* < \mathbb{E}(\hat{Z})/4 \right) \stackrel{\text{Obs. 20}}{\leq} \Pr \left( \hat{Z} - D < \mathbb{E}(\hat{Z})/4 \right) \\
&\leq \Pr \left( \hat{Z} < \mathbb{E}(\hat{Z})/2 \vee D \geq \mathbb{E}(\hat{Z})/4 \right) \\
&\leq \Pr \left( \hat{Z} < \mathbb{E}(\hat{Z})/2 \right) + \Pr \left( D \geq \mathbb{E}(\hat{Z})/4 \right) \\
&\stackrel{\text{Lem. 22}}{\leq} (c_3)^n + \Pr \left( D \geq \mathbb{E}(\hat{Z})/4 \right) \stackrel{*}{\leq} (c_3)^n + \frac{\mathbb{E}(D)}{\mathbb{E}(\hat{Z})/4} \\
&\stackrel{\text{Lem. 21}}{\leq} (c_3)^n + (c_2)^{n/2}/4 \leq (c_4)^n \text{ for some } c_4 \in (0, 1) \text{ and } n \text{ large enough.}
\end{aligned}$$

$\square$

### 6.5.5 Combining with Pool $(r_i)_{i \in [k]}$

We can now plug together the previous results, giving us the success probability of bipartite ShockHash when using a pool of size  $k$ . With  $\Pr(\text{success})$ , we denote the probability that there is a pair of compatible hash functions  $l_i$  and  $r_j$ ,  $i, j \in [n/2]$  in our pools.

**Lemma 24** Let  $Z$  be defined as above. Then  $\Pr(\text{success} \mid Z = z) = 1 - (1 - z)^k$ .

**Proof**

$$\begin{aligned}\Pr(\text{success} \mid Z = z) &= \Pr\left(\exists i \in [k] : r_i \in \bigcup_{j \in [k]} C(l_j) \mid \bigcup_{j \in [k]} C(l_j) = z\right) \\ &= \Pr_{p_1, \dots, p_k \sim \text{Ber}(z)}(\exists i \in [k] : p_i = 1) = 1 - (1 - z)^k.\end{aligned}$$

□

We already have most of the proof done. We just need to factor in the probability that the precondition of the previous lemma holds.

**Theorem 25** Let us take two pools  $(l_i)_{i \in [k]}$  and  $(r_i)_{i \in [k]}$  of size  $k = (e/2)^{n/2}$  containing randomly sampled hash functions. Then the probability that there are two hash functions  $l_i$  and  $r_j$  ( $i, j \in [k]$ ) in the pools such that  $\text{PF}(l_i, r_j)$  is  $> 0.17$  for  $n$  large enough.

**Proof**

$$\begin{aligned}\Pr(\text{success}) &\geq \Pr(\text{success} \wedge Z \geq \mathbb{E}(\hat{Z})/4) \\ &= \Pr(\text{success} \mid Z \geq \mathbb{E}(\hat{Z})/4) \cdot \Pr(Z \geq \mathbb{E}(\hat{Z})/4) \\ &\geq \Pr(\text{success} \mid Z \geq \mathbb{E}(\hat{Z})/4) \cdot \Pr(Z^* \geq \mathbb{E}(\hat{Z})/4) \\ &= \sum_{z=\mathbb{E}(\hat{Z})/4}^{\infty} \left( \Pr(\text{success} \mid Z = z) \cdot \Pr(Z = z \mid Z \geq \mathbb{E}(\hat{Z})/4) \right) \cdot \Pr(Z^* \geq \mathbb{E}(\hat{Z})/4) \\ &\stackrel{\text{Lem. 24}}{=} \sum_{z=\mathbb{E}(\hat{Z})/4}^{\infty} \left( (1 - (1 - z)^k) \cdot \Pr(Z = z \mid Z \geq \mathbb{E}(\hat{Z})/4) \right) \cdot \Pr(Z^* \geq \mathbb{E}(\hat{Z})/4) \\ &\geq \sum_{z=\mathbb{E}(\hat{Z})/4}^{\infty} \left( (1 - (1 - \mathbb{E}(\hat{Z})/4)^k) \cdot \Pr(Z = z \mid Z \geq \mathbb{E}(\hat{Z})/4) \right) \cdot \Pr(Z^* \geq \mathbb{E}(\hat{Z})/4) \\ &= \left( 1 - (1 - \mathbb{E}(\hat{Z})/4)^k \right) \cdot \underbrace{\sum_{z=\mathbb{E}(\hat{Z})/4}^{\infty} \left( \Pr(Z = z \mid Z \geq \mathbb{E}(\hat{Z})/4) \right) \cdot \Pr(Z^* \geq \mathbb{E}(\hat{Z})/4)}_{=1} \\ &\stackrel{\text{Lem. 23, 19}}{>} \left( 1 - \left( 1 - \frac{(e/2)^{-n/2}}{5} \right)^k \right) \cdot (1 - (c_4)^n) = \left( 1 - \left( 1 - \frac{1}{5k} \right)^k \right) \cdot (1 - (c_4)^n) \\ &\geq \left( 1 - e^{-1/5} \right) \cdot (1 - (c_4)^n) > 0.18 \text{ for } n \text{ large enough.}\end{aligned}$$

□

This concludes the proof of combining hash functions from a pool of candidates. We have seen that taking the pools of size  $k = (e/2)^{n/2}$  gives us a constant probability that there are two compatible functions in the pools. Note again that our actual

implementation uses one single, growing pool, not two fixed size pools. Therefore, we do not need to retry the construction if the initial pool size is not enough, but can just continue adding more functions to the pool.

## 6.6 ShockHash Construction

ShockHash tries different hash function seeds, which is equivalent to generating random graphs. Given the probability that a random graph is a pseudoforest, it is easy to determine the expected number of graphs ShockHash needs to try in order to find an MPHF. This leads directly to the space usage and construction time of ShockHash, which we state in the following Theorem.

**Theorem 26** *A ShockHash minimal perfect hash function mapping  $n$  keys to  $[n]$  needs  $\log_2(e)n + \mathcal{O}(\log n)$  bits of space in expectation and can be constructed in expected time  $\mathcal{O}((e/2)^n \cdot n)$ .*

**Proof** From Theorem 8, we know that the probability of the graph being a pseudoforest is  $\geq (e/2)^{-n} e^{-1} \sqrt{\pi}$ . We construct these graphs uniformly at random, so the expected number of seeds to try is  $\leq (e/2)^n e / \sqrt{\pi}$ . The space usage is given by the  $n + o(n)$  bits for the retrieval data structure, plus the bits to store the hash function index:

$$\mathbb{E}(\text{log}_2(\text{seed value})) \stackrel{*}{\leq} \log_2(\mathbb{E}(\text{seed value})) \leq \log_2((e/2)^n e / \sqrt{\pi}) = \log_2(e)n - n + \mathcal{O}(1).$$

In the step annotated with  $*$ , we use Jensen's inequality [34] and the fact that  $\log_2$  is concave.

For determining if at least one of the  $2^n$  functions corresponding to such a seed is valid, we can use an algorithm for finding connected components, as described in Section 4. This takes linear time for each of the seeds, resulting in an overall construction time of  $\mathcal{O}((e/2)^n \cdot n)$ . Constructing the retrieval data structure is then possible in linear time [19] and happens only once, so it is irrelevant for the asymptotic time here.  $\square$

Looking back at the simple brute-force approach, each of its  $e^n / \sqrt{2\pi n}$  expected trials needs  $n$  hash function evaluations, leading to a construction time of  $\mathcal{O}(e^n \sqrt{n})$ . Now, as shown in Theorem 26, ShockHash needs time  $\mathcal{O}((e/2)^n \cdot n)$ . This makes ShockHash almost  $2^n$  times faster than the previous state of the art. Given the observations in Ref. [9], we conjecture that ShockHash with rotation fitting reduces the number of hash function evaluations by an additional factor of  $n$ , while the space overhead tends to zero. In the following Theorem, we now give the resulting construction time and space consumption of the bipartite version.

**Theorem 27** *A bipartite ShockHash minimal perfect hash function needs  $\log_2(e)n + \mathcal{O}(\log n)$  bits of space in expectation and can be constructed in expected time  $\mathcal{O}(1.166^n)$ .*

**Proof** We know that we need to test  $(e/2)^n e / \sqrt{n}$  candidate pairs  $(h_0, h_1)$  in expectation before we find a perfect hash function. As described in Section 5, instead of

sampling two hash functions independently, we use a pool of hash functions and test all combinations of them. Theorem 25 shows that if we use a pool size of  $k = (e/2)^{n/2} \approx 1.166^n$ , we get a constant success probability. Until here, this does not improve the construction time asymptotically because all  $k^2$  combinations need to be tested. However, instead of combining all of the candidates, we can filter them directly while building the candidate pool. The filter, as with plain ShockHash, is very effective: The probability that the  $n$  hash values in a partition of size  $n/2$  hit all output positions is  $\Theta(0.836^{n/2})$  [58]. This means that we are only considering about  $((e/2)^{n/2} \cdot 0.836^{n/2})^2 = (e/2 \cdot 0.836)^n \approx 1.136^n$  pairs of hash functions in expectation. The construction time is therefore bounded by  $\max\{1.166^n, 1.136^n\}$ . Note that both of these values were rounded up anyway, so polynomial factors are dominated.

Looking at the space consumption of bipartite ShockHash, we need to encode two seeds of expected value  $\leq k$  each. Using Jensen's inequality and the retrieval data structure just like in Theorem 26, we get the resulting space usage. The fact that we suppressed polynomial factors in the analysis disappears in the  $\mathcal{O}(\log n)$  term.  $\square$

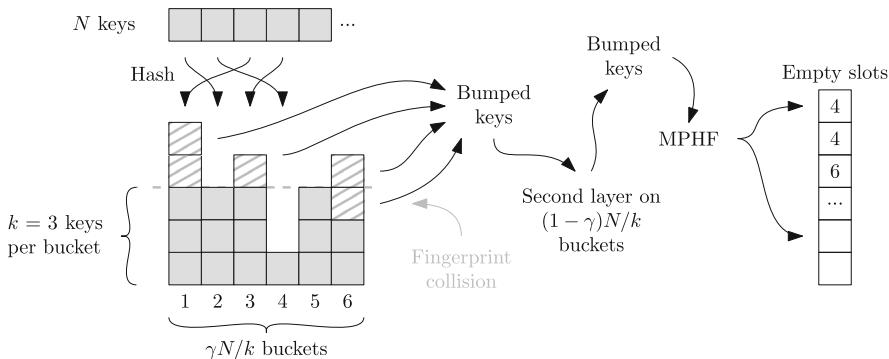
## 7 Partitioning

Even though ShockHash demonstrates significant speedups, by itself, it still needs exponential running time. As mentioned in the introduction, real world MPHF constructions usually do not search for a function for the entire input set directly. Instead, they partition the input of size  $N$  and then search on smaller subproblems of size  $n$ . In this section, we now give details on how to partition the input set efficiently before then using ShockHash as a building block.

### 7.1 ShockHash-RS = ShockHash + RecSplit

A first option is to integrate ShockHash as a base case into the highly space efficient RecSplit framework (see Section 2) and obtain ShockHash-RS. We keep the general structure of RecSplit intact. Only in the leaves, we use ShockHash instead of brute-force. We store the mapping from its keys to their hash function indices in one large retrieval data structure. Finally, after all leaves are processed, we construct the 1-bit retrieval data structure with all the  $N$  entries together.

**Fanouts.** RecSplit tries to balance the difficulty between the splittings and the bijections. ShockHash improves the performance of the bijections significantly but does not modify the way that the splittings are calculated. In this paper, we focus only on the bijections. To balance the amount of work done between splittings and bijections, we need to adapt the splitting parameters using the same techniques as the RecSplit paper. The RecSplit paper proves and uses optimal fanouts  $\lceil 0.35n + 0.5 \rceil$  and  $\lceil 0.21n + 0.9 \rceil$  for the two last splitting levels (see [21, Section 5.4]). For ShockHash-RS, we can adapt their formulas accordingly and get fanouts of  $\lfloor 0.10n + 0.5 \rfloor$  and  $\lfloor 0.073n + 0.9 \rfloor$ . However, preliminary experiments show that this is not optimal in practice. ShockHash is so much faster that the additional time invested into the splittings does not pay off. We find experimentally that setting the lowest splitting level to 4 and the second lowest to



**Fig. 8** Illustration for our bumped  $k$ -perfect hash function

3 achieves much better results in practice. To also provide faster and space-inefficient configurations, we set all fanouts to 2 when selecting leaf size  $n \leq 24$ .

## 7.2 ShockHash-Flat = ShockHash + $k$ -Perfect Hashing

An additional way to partition the input keys is to use  $k$ -perfect hashing. A minimal  $k$ -perfect hash function maps  $N$  keys to  $N/k$  output values, where each output value is hit exactly  $k$  times (assuming  $N$  divides  $k$ ). This has applications in external memory data structures [36, 37] and there are existing constructions [4]. The idea how to integrate ShockHash with  $k$ -perfect hashing is straightforward and similar to what we do in ShockHash-RS (see Section 7.1). We simply run a two-step process of first determining a  $k$ -perfect hash function, and then we construct small ShockHash data structures for the  $k$  keys hitting each output value. In contrast to ShockHash-RS (see Section 7.1), where some base cases could be smaller than  $n$ , here all of them have the same size.

**Bumped  $k$ -Perfect Hashing.** We now briefly describe a new  $k$ -perfect hash function. This function is focused on fast queries while still having rather small space consumption. Let us take  $N$  keys and hash them uniformly at random to a set of  $\gamma N/k$  buckets,  $\gamma \in (0, 1]$ . By choosing  $\gamma < 1$ , we can overload the buckets to ensure that most buckets receive at least  $k$  keys. In the experiments, we use  $\gamma = 0.9$ . We handle overflowing buckets by determining a fingerprint of each key. Each bucket then stores a threshold value using  $\log_2(k)$  bits that indicates which of the keys to bump from the bucket. This idea of bumping keys based on a threshold is inspired by bumped ribbon retrieval (BuRR) [19]. Separator hashing [31] uses a similar idea, however, without bumping keys completely and only supporting non-minimal perfect hash functions. We use a second level of the same data structure for the bumped keys, mapping them to the remaining  $(1 - \gamma)N/k$  buckets. Finally, we have a small number of keys that still get bumped in the second level. We first enumerate them by constructing a minimal perfect hash function. In our implementation, we use ShockHash-RS. With this, we then index an Elias-Fano coded sequence [20, 23] storing all empty slots in the output range. Figure 8 gives an illustration of the idea [55].

The advantage of this technique is that the majority of queries need a single access to the array of thresholds and a comparison with the key’s threshold. Few need two accesses to evaluate the second level, and only a tiny fraction of the queries needs to evaluate the explicit re-mapping.

**ShockHash-Flat.** From the bumped  $k$ -perfect hash function, we derive an MPHF that has a significantly more flat structure than ShockHash-RS. Instead of traversing a tree structure, it can perform a simple comparison with the threshold value for a majority of the input keys. Because we need to access both the threshold and then (usually) the seed of that same bucket, ShockHash-Flat stores thresholds and ShockHash seeds in an interleaved way.

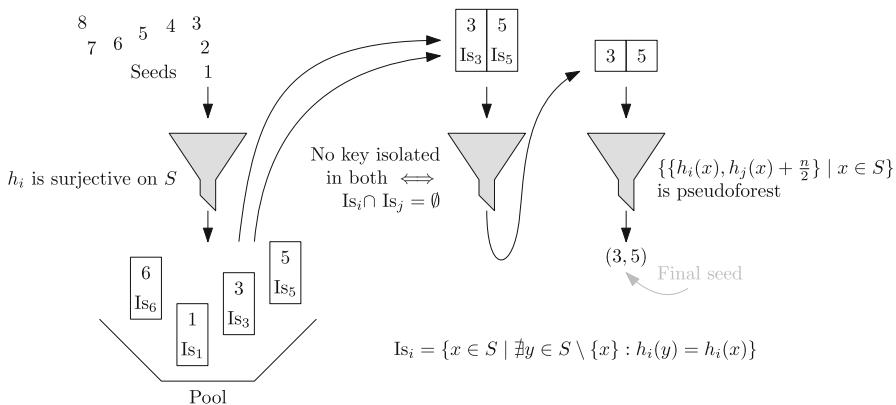
## 8 Variants and Refinements

In the following section, we describe variants and implementation details of ShockHash. In Section 8.1, we first explain how to achieve significant improvements in practice by using a bit-parallel filter. We then describe two techniques to come up with hash function candidates more efficiently, rotation fitting [9] and quad split, in Sections 8.3 and 8.2. To use bipartite ShockHash with uneven input sizes  $n$ , only small tweaks are necessary, which we describe in Section 8.4. We then continue with practical implementation tricks in Section 8.5. Finally, we describe ideas for parallelization in Section 8.6.

### 8.1 Isolated Keys Filter

In Section 5, we have already described how making the graph bipartite enables efficient filtering of hash function candidates. In the following section, we describe an additional way of filtering seeds. Bipartite ShockHash generates a set of surjective hash function candidates and then tests all combinations for orientability. By using an additional filter, we can speed up this test for orientability. The idea is to look at the case that a key is the only one mapping to a position. We refer to this key as *isolated* for that hash function seed. More formally, a key  $x$  is isolated using a hash function candidate  $h$ , if  $\{y \in S \mid h(x) = h(y)\} = \{x\}$ . If a key is isolated in *both* of the candidate hash functions, then in graph terminology this corresponds to a connected component with two nodes and one edge. Since each connected component of the final graph must have the same number of nodes and edges, there is then no need to perform the full test for orientability. We can determine bit patterns for each seed, indicating which of the keys are isolated. Then seed combinations can be ruled out using simple bit-parallel operations checking if the bit patterns are orthogonal. Note that the bit patterns used here refer not to the output positions but to the input keys (and therefore have size  $n$ ). Figure 9 illustrates the process.

A key is isolated in a partition if none of the other keys hash to its position, which happens with probability  $(1 - 1/(n/2))^{n-1} \rightarrow e^{-2}$ . A seed combination passes the filter if it has no key that is isolated in both partitions. This is approximately  $(1 - e^{-4})^n \approx 0.98^n$ , so the filter makes it possible to avoid the full check for a vast



**Fig. 9** Filtering for isolated keys. Each seed in the pool stores a bit vector of isolated keys, here annotated with Is<sub>i</sub>. A seed combination can only work if no key is isolated in both partitions, which can be efficiently checked using bit operations. Only if the filter is passed, we need to do the full orientability check

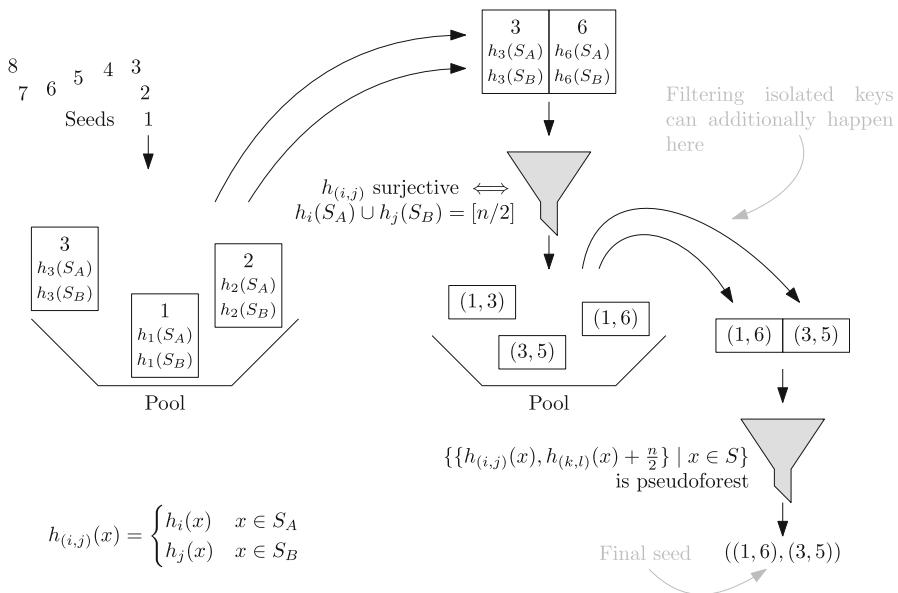
majority of seed combinations. Note that we apply the filter conditioned on the case that both functions are surjective, which should only make the filter more effective.

What makes this method interesting from a theoretical point of view is that we can be even smarter about filtering here. As stated before, if one of the hash functions has an isolated key at a position, we can skip testing it with all other hash functions that have an isolated key at the same position. We can organize all candidate hash functions in a binary trie data structure based on the isolated keys. Testing a new candidate hash function now boils down to traversing the trie. In theory, this gives additional exponential improvements in the construction time. However, preliminary experiments show that it is not helpful for the values of  $n$  we use in practice.

## 8.2 Rotation Fitting

A technique to speed up brute-force search for perfect hash functions is *rotation fitting* [9] (see Section 2). The same idea can be used in ShockHash to accelerate the search. We distribute the keys to two sets using an unseeded 1-bit hash function. We then determine the bit mask of output values that are hit in both of the sets. Like in the bit mask filter, which we use before checking for orientability (see Section 4), only if the logical OR of both masks has all bits set, it is worth testing the seed more closely. If we now cyclically rotate one of the bit masks and try again, we basically get a new chance of all output values being hit, without having to hash each key again. We then consider the distance to rotate the keys as part of the hash function seed. This corresponds to an addition modulo  $n$  to all hash values of the second set. We conjecture that – as in Ref. [9] – this reduces the number of hash function evaluations by a factor of  $n$ , while the space overhead tends to zero.

For bipartite ShockHash, rotation fitting can be applied as well, though in a slightly different way. Rotating one of the partitions of the bipartite graph within itself is not useful because it generates isomorphic graphs. Rotating the two partitions into each

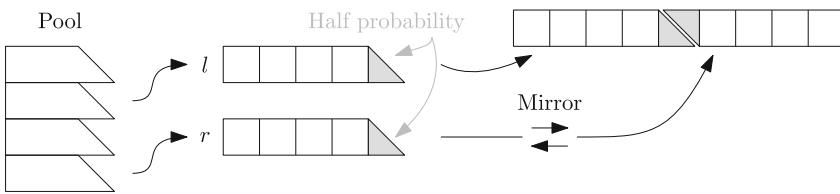


**Fig. 10** Additional filtering opportunities in quad split. We create a pool of all seeds without filtering but annotate each seed a bit vector containing hash function output values. Then we combine two seeds to form a seed for one partition. Surjectivity can be checked efficiently using bit operations. Therefore, the resulting seed is a tuple of four seeds – one for each half of the keys and each partition. In our implementation, we also add the filter for isolated keys on top (see Section 8.1 and Figure 9)

other would violate the bipartite condition, thus preventing to use the hash function pools. Instead, we can use rotation fitting to find seed candidates within each partition. More specifically, when looking for seed candidates for one partition, we distribute the keys to two subsets using an unseeded 1-bit hash function. We can then rotate one of the sets (modulo  $n/2$ ) to get additional hash function candidates. Each rotation can be tested for surjectivity using simple bit shifts that can happen in registers. In practice, this significantly improves the construction time because fewer hash functions need to be evaluated. However, the quad split technique described in the following section even enables exponential speedups.

### 8.3 Quad Split

The construction is dominated by the time spent evaluating hash function candidates (see Theorem 27), so it is natural to look at this step for improvements. For bipartite ShockHash, the quad split technique reduces the amount of time spent on finding surjective seed candidates. It basically applies the idea of bipartite ShockHash on another level of the same data structure. Like in rotation fitting, we split the input set into two sets  $S_A$  and  $S_B$  using a constant 1-bit hash function. Now we can hash each of the two sets using *independent* hash functions. In particular, we can test all combinations of assigning some hash function to each of the two sets. This reduces the number of hash function evaluations significantly.



**Fig. 11** Supporting uneven  $n$

For a seed  $i$ , let  $h_i(S_A)$  and  $h_i(S_B)$  indicate the sets of hash function output values of the two subsets  $S_A$  and  $S_B$ . A seed  $i$  for  $S_A$  can be used together with a seed  $j$  for  $S_B$  if  $h_i(S_A) \cup h_i(S_B) = [n/2]$ . By storing  $h_i(S_A)$  and  $h_i(S_B)$  as bit vectors indicating the output values, this compatibility check is a simple and efficient OR operation. In the quad split technique, we therefore annotate each hash function seed with this bit vector to enable fast search for a possible combination of seeds that is surjective. This process is illustrated in Figure 10.

Like with the isolated keys filter described in Section 8.1, we can again use a trie structure to avoid testing all combinations. We believe that this enables exponential improvements in the construction time, which could be implemented in future work. Quad split is orthogonal to the isolated keys filter, so we can actually combine both optimizations.

To encode the combined seed, we use a pairing function again. In contrast to the bipartite tries in ShockHash, the hash functions cannot be exchanged, so we cannot assume that one seed is larger than the other. We therefore need a more general pairing function. The most fitting pairing function here is Szudzik's pairing function (see Section 2), which enumerates, for all  $k \in \mathbb{N}$ , all pairs in  $[k] \times [k]$  before moving on to pairs involving numbers bigger than  $k$ . This means that we can test all combinations of previous hash functions before having to evaluate the next one. In our implementation, we make sure to try hash function combinations in linear order in the value of the pairing function.

#### 8.4 Supporting Uneven $n$

To support uneven numbers  $n$  of input keys, we can relax the bipartite property. The idea is that the output value  $\lceil n/2 \rceil$  can be hit by both hash functions, but each with half the probability. When combining the two halves, the value then gets the same probability as all other output values. For filtering candidate hash functions for surjectivity, the corresponding bit needs to be ignored – a seed candidate can be valid both if the bit is set or not set. Now, in order to use hash functions from a single pool for both the left and the right part, we have to mirror the functions used for the right part.<sup>6</sup> Refer for Figure 11 for an illustration.

<sup>6</sup> For simplicity, our current implementation uses plain shifting, accepting a doubled probability for the middle bit. Supporting uneven  $n$  then just boils down to rounding compile time constants to the right direction. When integrated into RecSplit, we expect uneven  $n$  to happen only once in every two buckets, so it has a negligible overhead of about  $1/2b$  bits per key.

## 8.5 Engineering

In the following section, we explain implementation tricks that we add to make our construction even faster in practice.

**Orientability Check.** Determining whether a given graph is a pseudoforest can be achieved in linear time using a connected components algorithm. However, this incurs significant overheads for building a graph data structure. To avoid this, our implementation therefore uses incremental cuckoo hash table construction with near linear time. For this, we keep an array representing the graph nodes. In order to avoid re-calculating a hash every time we evict a key, and to avoid case distinctions between first and second candidate cell, we use the known XOR trick. We annotate each key with the XOR of both its candidate cells. When evicting a key, we can calculate its other candidate cell efficiently by XORing the current cell with the stored value. This uses that XOR is commutative. We abort the insertion as soon as we detect a cycle.

**Partial Hash Calculation.** As described in Section 8.2, we can use rotation fitting in plain ShockHash. There we make the observation that hashing the first set of keys almost always yields a graph that, by itself, is a pseudoforest. This is not surprising because the load factor is usually close to the load threshold  $c = 0.5$  and  $n$  is small (which enables higher load [39]). We make use of this fact and reduce the number of hash function evaluations by keeping the hashes for the first set the same and just retrying hash functions for the second set. More precisely, if  $x$  is the hash function seed, we hash each key in the first set with seed  $x - (x \bmod k)$ , where  $k$  is a tuning parameter, and the keys of the second set with seed  $x$ . Therefore, the hash values of the first set can be cached over multiple iterations. In preliminary experiments, we find a value of  $k = 8$  to be a good fit – values much larger than that have diminishing returns in performance improvement and start to influence the space consumption. At  $k = 8$ , however, the influence on the space consumption is negligible when  $n$  is large. Given that hashing the keys is a bottleneck during construction, this reduces the number of keys that need to be hashed by a factor of close to 2. We only apply this optimization for large  $n > 32$ .

**Hash Cache.** In bipartite ShockHash, we regularly combine two hash function candidates from our pool to see if the resulting graph is a pseudoforest. While we skip that test for many of the candidates using the simple bit parallel filter described in Section 8.1, there is still a large number of candidates to compare. Re-evaluating the hash functions for these candidates can be a bottleneck depending on the input size  $n$ . An obvious idea is to cache the hash function output values of the seed candidates. Because the input sets and therefore the hash values are very small, we can store each hash value in a single byte. This makes the amount of space needed for each seed candidate relatively small.

**Sentinels.** For large  $n$ , the quad split technique spends most of its construction time calculating the logical OR of bit patterns looking for a result that has all bits set. This inner loop consists of only a very small number of assembly instructions. We can achieve considerable speedups here by adding a sentinel element to the end of the array that already has all bits set. Then we no longer need the repeated bounds check

for the array. When using SIMD parallelization (see Section 8.6), we use multiple sentinels depending on the number of SIMD lanes.

## 8.6 Parallelization

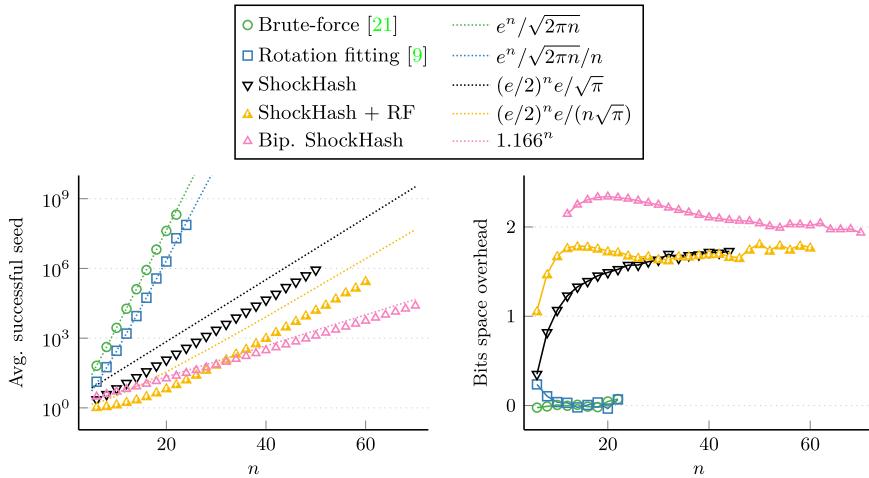
The main computational load behind ShockHash looks for seeds yielding pseudo-forests and can be parallelized on multiple levels: Over buckets when using RecSplit for partitioning, ShockHash building blocks, seeds, hash-function evaluations, and bit-parallel filters. The remaining operations are also well parallelizable: Hashing of keys to buckets can be split between processors. Parallel construction of the retrieval data structure can be done in a similar way [19]. In the following we explain one possible parallelization with respect to SIMD instructions and multi-threading. We also outline a hybrid CPU/GPU implementation.

**SIMD.** In plain ShockHash, we use SIMD parallelism in two locations. First, we use SIMD to determine the two candidate positions of all keys and to determine the bit mask for filtering. A key point here is to collect the bitwise OR of individual lanes and to only add the lanes together after all keys are done. Second, we use SIMD to evaluate the bit mask filter (see Section 4) with different rotations in parallel. Our implementation uses AVX-512 (8 64-bit values) if available and AVX2 (4 64-bit values) otherwise.

Bipartite ShockHash can also be parallelized using SIMD instructions. When using the quad split technique, we parallelize the test for candidate functions, which involves iterating over long lists of bit patterns, calculating the logical OR with each, and looking for a result that has all bits set. However, the other more involved data structures are harder to parallelize using SIMD because of more complex control flows. Therefore, we use SIMD only for checking lists of bit patterns, which is the main bottleneck for large  $n$ .

**Multi-Threading.** Because ShockHash is intended to be integrated into a partitioning framework, we can naively parallelize over the different ShockHash base cases. A simple coarse-grained source of parallelism are the RecSplit buckets. We can use any kind of load balancing to split them between threads. Even static load balancing may work because variances in construction time will average out. However, some kind of dynamic load balancing is likely to be more efficient and also works with cores of different speed that are now becoming standard in many multi-core processors. The retrieval data structure we use, BuRR [19], can be parallelized as well.

**GPUs.** A full GPU parallelization might be difficult and inefficient for cuckoo hashing as it has irregular control flow and memory access. Since filtering asymptotically dominates the computations for highly space-efficient variants, one might look at a hybrid implementation where a GPU produces a stream of seeds defining random graphs that cover all nodes and where a multicore CPU performs further stages of computation. For bipartite ShockHash with quad split, for example, the majority of the construction time is spent on comparing bit patterns to check if two hash function candidates are compatible. Here we can use the massive parallelism of GPUs to compare many patterns in parallel. Then the CPU can perform the less frequent and more complex checks for orientability.



**(a)** Average successful seed. For bipartite ShockHash, we use the implementation without quad split and plot the average of the largest seed that was evaluated before pairing up seeds. The dotted functions are upper bounds. Note that for bipartite ShockHash, we only know an upper bound of  $\mathcal{O}(1.166^n)$  and not  $1.166^n$ .

**(b)** Idealized space overhead over the lower bound  $\log_2(n^n/n!)$  in bits. If the average seed is  $s$  we charge  $\log_2(s)$  bits, plus  $n$  bits for retrieval (if applicable).

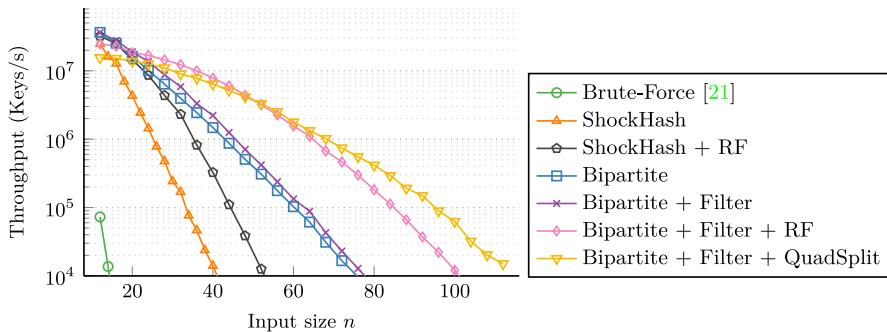
**Fig. 12** Hash function evaluations and space overhead of ShockHash compared with more simple brute-force techniques

## 9 Experiments

We run our experiments on an Intel i7 11700 processor with 8 cores and a base clock speed of 2.5 GHz. The machine runs Ubuntu 22.04 with Linux 5.15.0 and supports AVX-512 instructions. We use the GNU C++ compiler version 11.2.0 with optimization flags `-O3 -march=native`. For the competitors written in Rust, we compile in release mode with `target-cpu=native`.

Our implementation of ShockHash uses the BuRR retrieval data structure [19] with 128-bit ribbon width and 2-bit bumping information. For ShockHash-RS, we use partitioning based on RecSplit [21], in particular the SIMD-parallel implementation [9]. For partitioning keys in ShockHash-Flat, we sort them using IPS<sup>2</sup>Ra [2].

As input data, we use short strings of uniform random length  $\in [10, 50]$  containing random characters except for the zero byte. The reason for this is that all competitors natively support strings, while some only support integer keys when using their internal structs, which could give them an unfair advantage. Note that, as a first step, almost all compared codes generate a *master hash code* of each key using a high quality hash function. Any possible additional hash function can then be evaluated on the master hash code in constant time, independent of the input distribution. All experiments use a single thread. While almost all compared codes have a multi-threaded implementation, and perfect hashing can be parallelized trivially by partitioning, this is not the focus here. The code and scripts needed to reproduce our experiments are available on GitHub under the GNU General Public License [43, 53].



**Fig. 13** Construction throughput using different methods to come up with seed candidates. For comparison, the plot also includes Brute-Force search. Variants annotated with *RF* use rotation fitting. The filtered variants of bipartite ShockHash use the filter for isolated keys

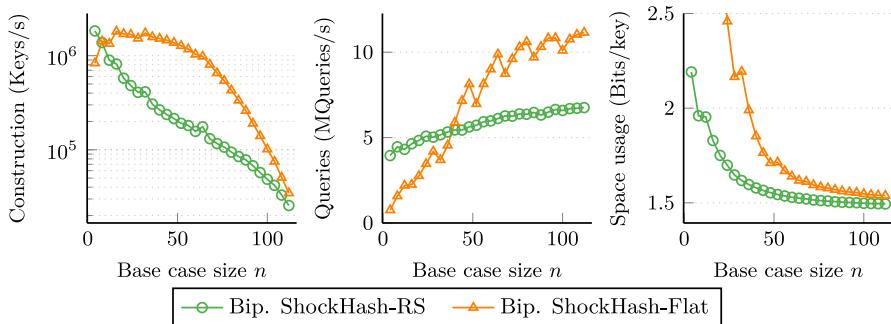
## 9.1 Number of Trials in Theory and Practice

In Figure 12, we compare the average number of hash function trials for each bijection search technique. From the different slopes of the curves, it is clearly visible that rotation fitting [9] saves a polynomial factor compared to plain brute-force, while ShockHash saves an exponential factor. Additionally, we plot the shown upper bounds for the number of trials of brute-force and ShockHash. For the rotation fitting variants, we plot the base variants divided by  $n$ , which is not formally shown to be a theoretical bound, but is an obvious conjecture. The plot shows that brute-force and rotation fitting are close to the given functions. For plain ShockHash, the measurements are even better than the theory, which suggests that our proof in Theorem 8 is not tight. Surprisingly, ShockHash seems to match the function we get when dividing our analysis by  $\sqrt{n}$ . We conjecture that the expected number of 1-orientations of a random pseudoforest might actually not be  $e \cdot \sqrt{2n}$ , but close to constant. This makes ShockHash an even better replacement for the brute-force technique. Bipartite ShockHash matches the slope of our analysis as well, but note that our analysis only shows an upper bound of  $\mathcal{O}(1.166^n)$ , not an exact value.

Figure 12 also gives the difference between the idealized space consumption and the space lower bound  $\log_2(n^n/n!)$ . It indicates that ShockHash loses space close to constant, which becomes negligible for larger  $n$ . This explains why we need to select larger  $n$  in ShockHash-RS compared to RecSplit with brute-force to achieve the same space consumption per key. Even with these larger  $n$ , ShockHash construction is significantly faster than brute-force. Bipartite ShockHash appears to have a small constant space overhead over plain ShockHash.

## 9.2 Seed Candidate Generation

Figure 13 shows different methods to generate seed candidates. For comparison, the plot also includes brute-force search. It is clearly visible that ShockHash is significantly faster than brute-force. Also, the bipartite version shows clear speedups compared



**Fig. 14** Space usage, construction performance and query performance for different partitioning schemes. The total number of input keys is  $N = 100$  million, which are then partitioned to small base cases of size  $n$ . Within the base cases we use bipartite ShockHash with quad split

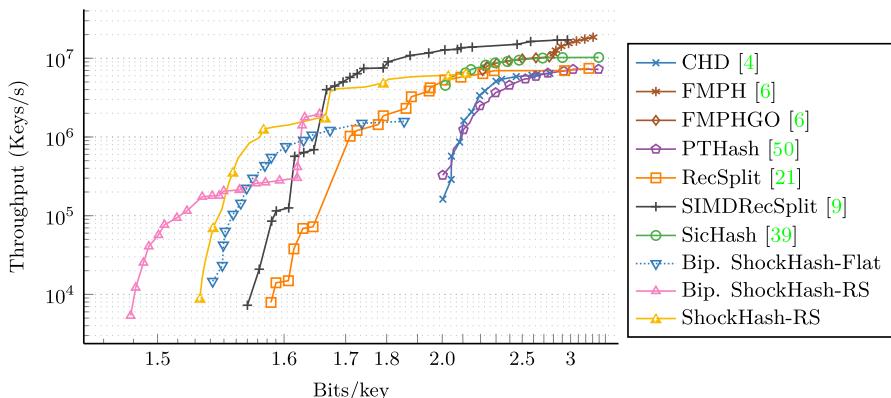
to the plain version. Filtering based on isolated keys (see Section 8.1) makes the construction about two times faster. While rotation fitting already gives impressive speedups, our quad split technique (see Section 8.3) is even faster for large  $n$ . Starting with about  $n = 60$ , the quad split technique is up to one order of magnitude faster than the basic bipartite ShockHash implementation. Note that the comparison here needs to be taken with a grain of salt because different methods have different space overheads based on  $n$ . However, for larger  $n$ , the methods have almost the same space usage. For a plot that takes space consumption into account, refer to Figure 15.

### 9.3 Partitioning

In Figure 14, we give the construction time, query time, and space consumption for the two different partitioning schemes, ShockHash-RS and ShockHash-Flat. For small  $n$ , the ShockHash-RS construction time is dominated by the splittings, which can be seen by the gap between ShockHash-RS and ShockHash-Flat. For large  $n$ , the base case dominates and both techniques have a similar throughput. Looking at the query performance, ShockHash-Flat is faster for  $n > 48$ , which is the most interesting range for good space consumption. The query throughput increases for larger base cases because we need to spend less time in the partitioning step. The jumps in the query throughput in ShockHash-Flat are caused by the fixed-width coding of ShockHash seeds, which needs a fallback data structure when a seed does not fit. For the same base case size, the space consumption of the flat partitioning scheme is higher than with RecSplit. Compared to ShockHash-RS, ShockHash-Flat trades space consumption for faster queries.

### 9.4 Comparison with Competitors

We now compare ShockHash-RS and ShockHash-Flat with competitors from the literature. Competitors include CHD [4], SicHash [39], PTHash [50], FMPHGO [6], RecSplit [21], and SIMDRecSplit [9]. We do not plot BBHash [42] because it is sig-



**Fig. 15** Pareto fronts of the space usage of different competitors.  $N = 10$  million keys. Note that both axes are logarithmic. For ShockHash-RS, we use ShockHash with SIMD and rotation fitting inside RecSplit. For ShockHash-RS, SicHash and PTHash, we plot all Pareto optimal data points but only show markers for every fourth point to increase readability. Therefore, the lines might bend on positions without markers

nificantly outperformed by FMPH [6], another implementation of the same technique. While SIMDRecSplit also includes a fast GPU implementation, we leave it out from most plots, as it would be unfair because of the different hardware architecture. While we measure ShockHash itself with  $N = 100$  million keys (see Figure 14), we perform the comparison with competitors with only  $N = 10$  million keys. This is because competitor configurations that achieve space usage close to ShockHash would otherwise need unreasonably long to compute.

**Construction.** Figure 15 gives Pareto fronts comparing the construction throughput of different competitors. Each Pareto front only gives data points that are not dominated by another configuration of the same method regarding both space and construction time. Note that we are mainly interested in configurations close to the space lower bound  $\log_2(e)$ . Only the RecSplit based competitors and ShockHash-Flat achieve space usage below 1.9 bits per key. At configurations with less than 1.65 bits per key, ShockHash-RS can significantly outperform SIMDRecSplit. ShockHash-RS is focused even more on the space efficient configurations than SIMDRecSplit. For the less space efficient configurations, it does not achieve the same throughput as SIMDRecSplit. This is not surprising because with these configurations, searching for bijections is fast, so constructing the retrieval data structure has a significant performance penalty.

Table 1 gives a selection of typical configurations. For SicHash [39], PTHash [50] and FMPHGO [6], we use the configurations given in the original papers, where FMPHGO is configured to use the hash cache. For RecSplit-based techniques, we mainly use space-efficient configurations with  $b = 2000$  and select the leaf size  $n$  such that a similar space consumption is achieved.

Comparing the configurations with a space consumption of 1.56 bit per key, ShockHash-RS is about 76 times faster than SIMDRecSplit, which is the next competitor not based on ShockHash. For space consumptions below 1.53 bits per key, bipartite ShockHash-RS is more space efficient and simultaneously three orders of magnitude

**Table 1** Query and construction performance of typical configurations of ShockHash and competitors. The table also shows GPURecSplit on an Nvidia RTX 3090

Method	Space bits/key	Construction ns/key	Query ns/query
FMPH, $\gamma=2.15$	3.529	51	54
FMPH, $\gamma=1.0$	2.804	89	69
FMPHGO, $\gamma=2.65, s=4, b=16$	3.547	86	51
FMPHGO, $\gamma=1.0, s=4, b=16$	2.212	135	69
PTHASH, $c=7.0, \alpha=0.99$ , C-C	3.524	198	20
PTHASH, $c=6.0, \alpha=0.99$ , EF	2.345	247	34
SicHash, $\alpha=0.9, p_1=21, p_2=78$	2.412	116	40
SicHash, $\alpha=0.97, p_1=45, p_2=31$	2.082	169	41
RecSplit, $n=8, b=100$	1.792	713	74
RecSplit, $n=14, b=2000$	1.585	125 521	97
SIMDRecSplit, $n=8, b=100$	1.808	117	80
SIMDRecSplit, $n=14, b=2000$	1.585	11 749	108
SIMDRecSplit, $n=16, b=2000$	1.560	137 902	100
SIMDRecSplit, $n=18, b=2000$	1.547	271 524	99
SIMDRecSplit, $n=20, b=2000$	1.535	5 569 394	97
(GPURecSplit, $n=20, b=2000$ )	1.536	55 988	100
(GPURecSplit, $n=24, b=2000$ )	1.496	508 768	94
<b>Bipartite ShockHash-RS</b> , $n = 64, b = 2000$	1.524	5 724	131
<b>Bipartite ShockHash-RS</b> , $n = 104, b = 2000$	1.496	24 406	121
<b>Bipartite ShockHash-RS</b> , $n = 128, b = 2000$	1.489	188 041	113
<b>Bipartite ShockHash-Flat</b> , $n = 100$	1.547	9 620	62
<b>Bipartite ShockHash-Flat</b> , $n = 128$	1.537	174 366	58
<b>ShockHash-RS</b> , $n = 30, b = 2000$	1.582	797	121
<b>ShockHash-RS</b> , $n = 39, b = 2000$	1.556	1 813	123
<b>ShockHash-RS</b> , $n = 58, b = 2000$	1.523	112 072	121

$N = 10$  million keys. We use SIMD instructions where available. For the bipartite versions, we use the quad split technique.

faster than SIMDRecSplit. Comparing different methods where each is given about half an hour of construction time, RecSplit is able to produce a perfect hash function with 1.58 bits per key. During the course of three years, SIMDRecSplit started a chain of work, first improving the space consumption to 1.56 bits per key. ShockHash-RS is then able to achieve 1.52 bits per key, reducing the gap to the lower space bound of  $\approx 1.442$  bits per key by about 30%. Finally, bipartite ShockHash-RS then reduces the space usage to just 1.489 bits per key, which is within 3.3% of the lower bound with practically feasible construction time. Looking at the absolute space consumption instead of the distance to the lower bound, bipartite ShockHash-RS achieves about

5% smaller perfect hash functions in the same construction time [9]. To construct an MPHF with 1.53 bits per key, SIMDRecSplit takes more than 15 hours, while bipartite ShockHash-RS achieves a better space consumption in just 57 seconds. Bipartite ShockHash-Flat trades space usage for better query speed. However, at a configuration achieving the same space usage as SIMDRecSplit, it is still about 32 times faster to construct.

**Queries.** Table 1 also shows the query throughput of typical configurations. RecSplit-based techniques have slower queries than the other techniques as they have to traverse several levels of a tree, decoding variable-bitlength data in each step. ShockHash-RS additionally needs to access a retrieval data structure. However, when comparing configurations that achieve a similar space efficiency, the query performance of ShockHash-RS is similar to competitors. This shows that the overhead of the retrieval operation is small compared to the work for traversing the heavily compressed tree.

Compared to plain ShockHash-RS, our bipartite implementation loses about 10% of the query performance at the same space usage. This is due to un-pairing seeds and having to do case distinctions based on the partitions and on the subset during quad split. However, the larger leaves in bipartite ShockHash-RS can reduce the time spent on traversing RecSplit’s splitting tree data structure, which is a major bottleneck for the queries. Bipartite ShockHash-Flat can be constructed about 32 times faster than the previously most space-efficient competitor SIMDRecSplit for the same space requirement. Simultaneously, 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.

When query performance is the main concern, PTHash [50] trades much higher space usage for much faster queries. SicHash takes the middle-ground between PTHash and bipartite ShockHash-Flat both with respect to query time and space consumption while allowing faster construction than both approaches.

## 10 Conclusion and Future Work

ShockHash is a new way to compute minimal perfect hash functions on small sets. By combining trial-and-error search with cuckoo hashing and retrieval data structures, ShockHash achieves an exponential speedup over plain brute-force (almost a factor  $2^n$ ). While the plain brute-force technique samples functions and hopes for one to be an MPHF, ShockHash samples graphs and hopes for one to be a pseudoforest. With bipartite ShockHash, we present an extension that samples bipartite graphs and performs aggressive filtering to achieve additional exponential speedups. These improvements enable the currently fastest way to achieve near space-optimal minimal perfect hash functions and breaks the dominance of the previous best methods that relied on pure brute-force for their base-case subproblems.

We integrate ShockHash as a base case into different partitioning frameworks. When using ShockHash inside RecSplit, we get ShockHash-RS, which can be constructed up to three orders of magnitude faster than the previous state of the art when comparing sequential codes. Constructing with a single thread, ShockHash-RS is even faster

than a tuned GPU implementation of the brute-force technique. Using ShockHash inside a newly developed  $k$ -perfect hash function, we get ShockHash-Flat, which can be constructed about 32 times faster than the previous state of the art, while simultaneously having 30% faster queries. This starts to close the gap to constructions that are way less space efficient and brings space efficient perfect hash functions closer to practical applications.

**Future Work.** In the future, additional filtering techniques would be interesting that could give additional speedups. We already describe one such filter, using a trie to skip incompatible hash function seeds before trying to combine them. To support even larger building blocks, a hybrid parallelization of the technique using a GPU would be an interesting future direction. Since the submission of this paper, MorphisHash was able to reduce the space overhead by introducing a retrieval data structure specifically optimized for ShockHash. The new Consensus-RecSplit technique achieves even smaller space consumption, although with slower queries.

**Author contributions** H.L. is the main author of the paper. He implemented the algorithm, performed the experiments and wrote most of the manuscript. All authors made significant contributions, to algorithm design, analysis, design &; interpretation of the experiments, and the writeup, with P.S. more on the side of design and experiments and S.W. more on the side of the analysis.

**Funding** Open Access funding enabled and organized by Projekt DEAL. 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).



**Related Version.** This journal version is based on a conference paper [40] published at ALENEX 2024. Bipartite ShockHash was previously described in a technical report [38] on arXiv without an analysis. This journal version adds an analysis of bipartite ShockHash and introduces ShockHash-Flat, a variant with faster queries.

## Declarations

**Conflicts of Interest** We have no conflicts of interest to disclose.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Almodaresi, Fatemeh, Sarkar, Hirak, Srivastava, Avi, Patro, Rob: A space and time-efficient index for the compacted colored de bruijn graph. *Bioinform.* **34**(13), i169–i177 (2018). <https://doi.org/10.1093/BIOINFORMATICS/BTY292>

2. Axtmann Michael, Witt Sascha, Ferizovic Daniel, Sanders Peter.: Engineering in-place (shared-memory) sorting algorithms. *ACM Trans. Parallel Comput.*, 9(1):2:1–2:62, 2022. <https://doi.org/10.1145/3505286>
3. Belazzougui Djamel, Boldi Paolo, Pagh Rasmus, Vigna Sebastiano.: Fast prefix search in little space, with applications. In *ESA (1)*, volume 6346 of *Lecture Notes in Computer Science*, pages 427–438. Springer, 2010. [https://doi.org/10.1007/978-3-642-15775-2\\_37](https://doi.org/10.1007/978-3-642-15775-2_37)
4. Belazzougui Djamel, Botelho Fabiano C., Dietzfelbinger Martin.: Hash, displace, and compress. In *ESA*, volume 5757 of *Lecture Notes in Computer Science*, pages 682–693. Springer, 2009. [https://doi.org/10.1007/978-3-642-04128-0\\_61](https://doi.org/10.1007/978-3-642-04128-0_61)
5. Belazzougui Djamel, Navarro Gonzalo.: Alphabet-independent compressed text indexing. *ACM Trans. Algorithms*, 10(4):23:1–23:19, 2014. <https://doi.org/10.1145/2635816>
6. Beling, Piotr: Fingerprinting-based minimal perfect hashing revisited. *ACM Journal of Experimental Algorithmics* (2023). <https://doi.org/10.1145/3596453>
7. Bender Michael A., Farach-Colton Martin, Goswami Mayank, Johnson Rob, McCauley Samuel, Singh Shikha.: Bloom filters, adaptivity, and the dictionary problem. In *FOCS*, pages 182–193. IEEE Computer Society, (2018). <https://doi.org/10.1109/FOCS.2018.00026>
8. Bernstein Sergei.: On a modification of chebyshev's inequality and of the error formula of laplace. *Ann. Sci. Inst. Sav. Ukraine, Sect. Math.*, 1(4):38–49, (1924)
9. Bez Dominik, Kurpicz Florian, Lehmann Hans-Peter, Sanders Peter.: High performance construction of RecSplit based minimal perfect hash functions. In *ESA*, volume 274 of *LIPICS*, pages 19:1–19:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. <https://doi.org/10.4230/LIPICS.ESA.2023.19>
10. Botelho Fabiano C., Pagh Rasmus, Ziviani Nivio.: Perfect hashing for data management applications. *CoRR*, abs/cs/0702159, 2007. <https://doi.org/10.48550/ARXIV.CS/0702159>
11. Botelho, Fabiano C., Pagh, Rasmus, Ziviani, Nivio: Practical perfect hashing in nearly optimal space. *Inf. Syst.* **38**(1), 108–131 (2013). <https://doi.org/10.1016/J.IS.2012.06.002>
12. Cayley, Arthur: A theorem on trees. *Quart. J. Math.* **23**, 376–378 (1878)
13. Chapman Jarrod A., Ho Isaac, Sunkara Sirisha, Luo Shujun, Schroth Gary P., Rokhsar Daniel S.: Meraculous: De novo genome assembly with short paired-end reads. *PLOS ONE*, 6(8):1–13, 08 (2011). <https://doi.org/10.1371/journal.pone.0023501>
14. Chen, Yupeng, Schmidt, Bertil, Maskell, Douglas L.: A hybrid short read mapping accelerator. *BMC Bioinform.* **14**, 67 (2013). <https://doi.org/10.1186/1471-2105-14-67>
15. Chikhi, Rayan, Lemasset, Antoine, Medvedev, Paul: Compacting de bruijn graphs from sequencing data quickly and in low memory. *Bioinform.* **32**(12), 201–208 (2016). <https://doi.org/10.1093/BIOINFORMATICS/BTW279>
16. Dietzfelbinger Martin, Heide Friedhelm Meyer auf der.: A new universal class of hash functions and dynamic hashing in real time. In *ICALP*, volume 443 of *Lecture Notes in Computer Science*, pages 6–19. Springer, (1990). <https://doi.org/10.1007/BFb0032018>
17. Dietzfelbinger Martin, Rink Michael. Applications of a splitting trick. In *Automata, Languages and Programming: 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5–12, 2009, Proceedings, Part I* 36, pages 354–365. Springer, (2009). [https://doi.org/10.1007/978-3-642-02927-1\\_30](https://doi.org/10.1007/978-3-642-02927-1_30)
18. Dietzfelbinger Martin, Weidling Christoph.: Balanced allocation and dictionaries with tightly packed constant size bins. In *ICALP*, volume 3580 of *Lecture Notes in Computer Science*, pages 166–178. Springer, (2005). [https://doi.org/10.1007/11523468\\_14](https://doi.org/10.1007/11523468_14)
19. Dillinger Peter C., Hübschle-Schneider Lorenz, Sanders Peter, Walzer Stefan. Fast succinct retrieval and approximate membership using ribbon. In *SEA*, volume 233 of *LIPICS*, pages 4:1–4:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. <https://doi.org/10.4230/LIPICS.SEA.2022.4>
20. Elias, Peter: Efficient storage and retrieval by content and address of static files. *J. ACM* **21**(2), 246–260 (1974). <https://doi.org/10.1145/321812.321820>
21. Esposito Emmanuel, Graf Thomas Mueller, Vigna Sebastiano. RecSplit: Minimal perfect hashing via recursive splitting. In *ALENEX*, pages 175–185. SIAM, (2020). <https://doi.org/10.1137/1.9781611976007.14>
22. Fan Bin, Andersen David G., Kaminsky Michael, Mitzenmacher Michael.: Cuckoo filter: Practically better than bloom. In *CoNEXT*, pages 75–88. ACM, (2014). <https://doi.org/10.1145/2674005.2674994>
23. Fano Robert Mario.: On the number of bits required to implement an associative memory. Technical report, MIT, Computer Structures Group, (1971). Project MAC, Memorandum 61"

24. Fotakis, Dimitris, Pagh, Rasmus, Sanders, Peter, Spirakis, Paul G.: Space efficient hash tables with worst case constant access time. *Theory Comput. Syst.* **38**(2), 229–248 (2005). <https://doi.org/10.1007/S00224-004-1195-X>
25. Fountoulakis, Nikolaos, Khosla, Megha, Panagiotou, Konstantinos: The multiple-orientability thresholds for random hypergraphs. *Combinatorics, Probability & Computing* **25**(6), 870–908 (2016). <https://doi.org/10.1017/S0963548315000334>
26. Fountoulakis, Nikolaos, Panagiotou, Konstantinos: Sharp load thresholds for cuckoo hashing. *Random Struct. Algorithms* **41**(3), 306–333 (2012). <https://doi.org/10.1002/RSA.20426>
27. Fox Edward A., Chen Qi Fan, Heath Lenwood S.: A faster algorithm for constructing minimal perfect hash functions. In *SIGIR*, pages 266–273. ACM, (1992). <https://doi.org/10.1145/133160.133209>
28. Fredman, Michael L., Komlós, János., Szemerédi, Endre: Storing a sparse table with 0(1) worst case access time. *J. ACM* **31**(3), 538–544 (1984). <https://doi.org/10.1145/828.1884>
29. Frieze Alan, Karoński Michał.: *Introduction to random graphs*. Cambridge University Press, (2016)
30. Golomb Solomon W.: Run-length encodings (corresp.). *IEEE Trans. Inf. Theory*, **12**(3):399–401, (1966). <https://doi.org/10.1109/TIT.1966.1053907>
31. Gonnet, Gaston H., Larson, Per-Åke: External hashing with limited internal storage. *J. ACM* **35**(1), 161–184 (1988). <https://doi.org/10.1145/42267.42274>
32. Hagerup Torben, Tholey Torsten.: Efficient minimal perfect hashing in nearly minimal space. In *STACS*, volume 2010 of *Lecture Notes in Computer Science*, pages 317–326. Springer, (2001). [https://doi.org/10.1007/3-540-44693-1\\_28](https://doi.org/10.1007/3-540-44693-1_28)
33. Janson, Svante, Luczak, Malwina J.: A simple solution to the  $k$ -core problem. *Random Struct. Algorithms* **30**(1–2), 50–62 (2007). <https://doi.org/10.1002/rsa.20147>
34. Johan Ludwig William Valdemar Jensen: Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta mathematica* **30**(1), 175–193 (1906)
35. Joag-Dev, Kumar, Proschan, Frank: Negative association of random variables with applications. *The Annals of Statistics* **11**(1), 286–295 (1983). <https://doi.org/10.1214/aos/1176346079>
36. Kurpicz Florian, Lehmann Hans-Peter, Sanders Peter.: PaCHash: Packed and compressed hash tables. In *2023 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 162–175. SIAM, (2023). <https://doi.org/10.1137/1.9781611977561.ch14>
37. Larson Per-Åke, Ramakrishna M. V.: External perfect hashing. In *SIGMOD Conference*, pages 190–200. ACM Press, (1985). <https://doi.org/10.1145/318898.318916>
38. Lehmann Hans-Peter, Sanders Peter, Walzer Stefan.: Bipartite ShockHash: Pruning ShockHash search for efficient perfect hashing. *CoRR*, abs/2310.14959v1, (2023). <https://doi.org/10.48550/ARXIV.2310.14959v1>
39. Lehmann Hans-Peter, Sanders Peter, Walzer Stefan.: SicHash – small irregular cuckoo tables for perfect hashing. In *ALENEX*, pages 176–189. SIAM, (2023). <https://doi.org/10.1137/1.9781611977561.CH15>
40. Lehmann, Hans-Peter., Sanders, Peter, Walzer, Stefan: Shockhash: Towards optimal-space minimal perfect hashing beyond brute-force. In *ALENEX*. SIAM (2024). <https://doi.org/10.1137/1.9781611977929.15>
41. Lelarge Marc.: A new approach to the orientation of random hypergraphs. In *Proc. 23rd SODA*, pages 251–264, 2012. <https://doi.org/10.1137/1.9781611973099.23>
42. Antoine Limasset, Guillaume Rizk, Rayan Chikhi, and Pierre Peterlongo.: Fast and scalable minimal perfect hashing for massive key sets. In *SEA*, volume 75 of *LIPICS*, pages 25:1–25:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. <https://doi.org/10.4230/LIPICS.SEA.2017.25>
43. MPHFs Experiments – GitHub. <https://github.com/ByteHamster/MPHF-Experiments>, 2023
44. Molloy, Michael: Cores in random hypergraphs and boolean formulas. *Random Struct. Algorithms* **27**(1), 124–135 (2005). <https://doi.org/10.1002/rsa.20061>
45. Müller Ingo, Sanders Peter, Schulze Robert, Zhou Wei.: Retrieval and perfect hashing using fingerprinting. In *SEA*, volume 8504 of *Lecture Notes in Computer Science*, pages 138–149. Springer, (2014). [https://doi.org/10.1007/978-3-319-07959-2\\_12](https://doi.org/10.1007/978-3-319-07959-2_12)
46. Newman Mark E. J.: *Networks: An Introduction*. Oxford University Press, (2010)
47. Pagh, Anna, Pagh, Rasmus: Uniform hashing in constant time and optimal space. *SIAM J. Comput.* **38**(1), 85–96 (2008). <https://doi.org/10.1137/060658400>
48. Pagh Anna, Pagh Rasmus, Ruzic Milan.: Linear probing with constant independence. In *STOC*, pages 318–327. ACM, (2007). <https://doi.org/10.1145/1250790.1250839>
49. Rasmus Pagh and Flemming Friche Rodler: Cuckoo hashing. *J. Algorithms* **51**(2), 122–144 (2004). <https://doi.org/10.1016/j.jalgor.2003.12.002>

50. Pibiri Giulio Ermanno, Trani Roberto.: PTHash: Revisiting FCH minimal perfect hashing. In *SIGIR*, pages 1339–1348. ACM, (2021). <https://doi.org/10.1145/3404835.3462849>
51. Rice, Robert F.: Some practical universal noiseless coding techniques. JPL Publication, Jet Propulsion Laboratory (1979)
52. Stéphane Boucheron Gábor LugosiPascal Massart.:Concentration Inequalities: A Nonasymptotic Theory of Independence. Oxford University Press (2013). <https://doi.org/10.1093/ACPROF:OSO/9780199535255.001.0001>
53. ShockHash – GitHub. <https://github.com/ByteHamster/ShockHash>, 2023
54. Stanley Richard P.: Enumerative combinatorics volume 1 second edition. *Cambridge studies in advanced mathematics*, (2011)
55. Stefan Hermann and Sebastian Kirmayer and Hans-Peter Lehmann and Peter Sanders and Stefan Walzer.: Engineering Minimal k-Perfect Hash Functions, CoRR. abs/2504.20001. (2025). <https://doi.org/10.48550/ARXIV.2504.20001>
56. Szudzik Matthew.: An elegant pairing function. In *Wolfram Research (ed.) Special NKS 2006 Wolfram Science Conference*, pages 1–12, (2006)
57. Walzer Stefan.: Peeling close to the orientability threshold – spatial coupling in hashing-based data structures. In *SODA*, pages 2194–2211. SIAM, (2021). <https://doi.org/10.1137/1.9781611976465.131>
58. Walzer Stefan.: The probability to hit every bin with a linear number of balls, (2024). <https://doi.org/10.48550/ARXIV.2403.00736>
59. Witten Ian H., Moffat Alistair, Bell Timothy C.: *Managing Gigabytes: Compressing and Indexing Documents and Images, Second Edition*. Morgan Kaufmann, (1999)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.