Ribbon filter: practically smaller than Bloom and Xor[1]

Peter C. Dillinger
Facebook, Inc.
Seattle, Washington, USA
peterd@fb.com

Stefan Walzer
University of Cologne
Cologne, Germany
walzer@cs.uni-koeln.de

ABSTRACT

Filter data structures over-approximate a set of hashable keys, i.e. set membership queries may incorrectly come out positive. A filter with *false positive rate* $f \in (0, 1]$ is known to require $\geq \log_2(1/f)$ bits per key. At least for larger $f \geq 2^{-4}$, existing practical filters require a space overhead of at least 20% with respect to this information-theoretic bound.

We introduce the Ribbon filter: a new filter for static sets with a broad range of configurable space overheads and false positive rates with competitive speed over that range, especially for larger $f \geq 2^{-7}$. In many cases, Ribbon is faster than existing filters for the same space overhead, or can achieve space overhead below 10% with some additional CPU time. An experimental Ribbon design with load balancing can even achieve space overheads below 1%.

A Ribbon filter resembles an Xor filter modified to maximize locality and is constructed by solving a band-like linear system over Boolean variables. In previous work, Dietzfelbinger and Walzer describe this linear system and an efficient Gaussian solver. We present and analyze a faster, more adaptable solving process we call "**R**apid **I**ncremental **B**oolean **B**anding **ON** the fly," which resembles hash table construction. We also present and analyze an attractive Ribbon variant based on making the linear system homogeneous, and describe several more practical enhancements.

1 INTRODUCTION

Background and motivation. The primary motivation for this work is optimizing data retrieval, especially in systems aggregating immutable data resources. In the example of LSM-tree storage [50], persisted key-value data is primarily split among immutable data files. A crucial strategy for reducing I/O in key look-ups is filtering accesses using an in-memory data structure. In a common configuration, each data file has an associated Bloom filter [9] representing the set of keys with associated data in that file. The Bloom filter has some false positive (FP) rate, which is the probability that querying a key not added returns TRUE (positive). For example, configuring the Bloom filter to use 10 bits of space per added key yields an FP rate of just under 1%, regardless of the size or structure of the keys

themselves[2]. Thus, the Bloom filter *filters out* almost all specific key queries[3] to data files that would find no relevant data. False negative (FN) queries would be incorrect for this application and must never occur.

Blocked Bloom filters [44, 54] are a popular Bloom variant because they are extremely fast. We do not expect to improve upon this solution[4] for short-lived applications such as database joins or the smallest levels of an LSM-tree. However, Bloom filters use at least 44% more space ("space overhead") than the information-theoretic lower bound of $\lambda = \log_2(1/f)$ bits per key for a hashed filter with FP rate f [13, Section 2.2]. Blocked Bloom can exceed 50% space overhead for small f.

In this work we focus on saving space in static filters, and optimizing the CPU time required for saving space. Our presentation and validation are kept general, but an intended application is optimizing accesses to the largest levels of an LSM-tree, where it should be worth CPU time to save space in long-lived memory-resident structures[5]. In [18] it is shown that a relatively high FP rate for these levels is best for overall efficiency. However, the space savings offered by existing practical Bloom filter alternatives is limited (dotted line in Figure 1 **(a)**), especially for higher FP rates, $\lambda \leq 5$.

Bloom filter and alternatives. We can categorize hashed filters by the logical structure of a query:

> **OR probing:** Cuckoo filters [11, 32, 34, 35], Quotient filters [5, 15, 30, 51], and most others return TRUE for a query if any one of several probed locations is a hashed data match, much like hash table look-ups. This design is great for supporting dynamic add and delete, but all known instances use $(1 + \varepsilon)\lambda + \mu$ bits per key where $\mu > 1.44$, or often $\mu = 3$ for speed[6]. Even with $\varepsilon \approx 0.05$, the space overhead is large for small λ.
>
> **AND probing:** A Bloom filter query returns TRUE iff all probed locations (bits) are a match (set to 1).

---

[1]Title inspired by [34, 35, 39].

[2]Learned filters [55] or tries [59] can take advantage of regularities in the key set. A space-efficient hash table [15] can take advantage of a densely covered key space. We focus on the general case.
[3]Static filters can also support range queries, either through *prefix Bloom* [48] or more sophisticated schemes [47].
[4]In Section 7 we mention and use another blocked Bloom implementation with new trade-offs [24].
[5]In some large-scale applications using RocksDB [31] for LSM-tree storage, we observe roughly 10% of memory and roughly 1% of CPU used in blocked Bloom filters. The size-weighted average age of a live filter is about three days. An experimental Ribbon filter option was added to RocksDB in version 6.15.0 (November 2020).
[6]$\mu > 1.44$ can be explained by these structures approximating a minimal perfect hash [36], a near-strict ordering of keys based on the location of their matching entry in the structure, on top of λ payload bits per key. This is an observation about existing structures, not necessarily a fundamental limitation.

XOR probing: Xor filters [10, 14, 20, 37, 39] return TRUE for a query iff the bitwise exclusive-or (XOR) of all probed locations is a hashed data match. XOR probing is only known to work with static filters.

Structures using XOR probing are the most promising for space efficient static filters. They are constructed by solving a linear system with one constraint per key ensuring that querying the key returns TRUE. Standard Xor filters use a fast solving process called *peeling* that limits their space efficiency to $\geq 1.22\lambda$ bits per key[7], though a variant with fast compression, the Xor+ filter [39], uses roughly $1.08\lambda + 0.5$ bits per key, which is an improvement for larger λ. Using structured Gaussian elimination instead of peeling [21, 37] offers better space efficiency, but construction times are considered impractical for many applications.

Core contribution. We introduce a faster, simplified, and more adaptable Gaussian elimination algorithm (Ribbon) for the static function data structure from [22]. Based on Ribbon, we develop a family of practical and highly space-efficient XOR-probed filters.

Results and comparison. Figure 1 **(a)** summarizes extensive benchmarking data by indicating which structure is fastest for satisfying various space and FP rate requirements for a static filter. For "fastest" we consider the sum of the construction time per key and three query times (measured for $x \in S$, $x \notin S$ and a mixed data set).[8]

Although we compare Ribbon with many approaches implemented in the fastfilter benchmark library [38], only variants of Bloom, Cuckoo, Xor, and Ribbon emerge as winners. Specifically, the color at point (x, y) indicates the fastest filter with space overhead at most x and FP rate $f \in [y/2, y \cdot 2]$ for $n = 10^7$ keys. Diagonal shading indicates different winners for $n = 10^6$ and $n = 10^8$. The timings for the winning approach are also shown in Figure 1 **(b)**. We observe the following.

- Ribbon wins to the right of the dotted line because none of the competing approaches achieve space overhead this low.
- Ribbon wins in some territory previously occupied by Xor and Xor+ filters, mostly for $f > 2^{-8}$ ($\lambda < 8$) from Xor and $f > 2^{-12}$ ($\lambda < 12$) for Xor+. In these cases, Ribbon-style Gaussian elimination is faster than peeling.
- Blocked Bloom filters are still the fastest whenever applicable, though Cuckoo and Xor take some of the remaining >44% territory (and nearby) for small FP rates.

Outline. The paper is structured as follows.

Section 2. We briefly review data structures for **static functions** and how they give rise to filters.

Section 3. We describe and analyze the new Ribbon construction algorithm, which preserves asymptotic guarantees from [22]. We also show how to improve the space efficiency of

Figure 1: (a) Fastest filter with the given combination of space overhead and false positive rate, considering a mix of construction and query times. (b) Construction and query times for fastest approaches in (a).

small Ribbon structures ("smash"). These features go into the Standard Ribbon filter, which in practice has increasing space overhead or running time as the number of keys increases.[9]

Section 4. We present the Homogeneous Ribbon filter, which shares many desirable properties with blocked Bloom filters: construction success is guaranteed, and scaling to any number of keys is efficient. Homogeneous Ribbon does not build on static functions in the standard way, which simplifies implementation but complicates analysis.

Section 5. We describe some practical enhancements and issues for Ribbon filters, including (1) efficiently utilizing any amount of memory for any number of keys, (2) laying out data for efficient queries, (3) efficiently satisfying hashing requirements, and (4) scaling Standard Ribbon with data sharding.

Section 6. We present Balanced Ribbon, an experimental extension of Standard Ribbon that uses a greedy load balancing scheme within a contiguous ribbon. This scales the extreme space efficiency of small Standard Ribbon filters to very large n, such as only $1.005\lambda + 0.008$ bits per key with practical construction and query times.

Section 7. We present more experimental validation.

---

[7]A new spatially-coupled construction for Xor filters [57] promises lower space overheads and in some cases slightly improved construction time with peeling. Simulations indicate a number of keys on the order of 10^6 or more is needed for most of the benefit, limiting the known generality of the approach. Configuration in practice is not well understood.

[8]The "right" weighing of construction and query time clearly depends on the use case. Because LSM-trees are especially useful for write-heavy workloads requiring good read latency, we find this a reasonable ratio for at least that use case. If a 4KB filter memory page has a lifetime as long as two weeks and at least one negative (useful) query per added key (n roughly 2^{12}) is seen, that satisfies the current rule of five minutes for caching SSD storage in RAM [2].

[9]If the processor word size is assumed to be $\Omega(\log n)$, query time and space overhead can be kept constant. We make no such assumption here.

2 FROM STATIC FUNCTIONS TO FILTERS

Approximate membership queries and filters. An approximate membership query filter – *filter* for short – represents a set $S \subseteq \mathcal{U}$ from some universe $\mathcal{U}$. A membership query with $x \in S$ must return TRUE, while a query with $x \in \mathcal{U} \setminus S$ may return TRUE with probability at most f where $f > 0$ is the *false positive (FP) rate*.

Static functions. A *static function* is a data structure (sometimes called "retrieval data structure") representing a function $b : S \to \{0, 1\}^r$ for some set $S \subseteq \mathcal{U}$ of keys. A query for $x \in S$ must return $b(x)$ but a query for $x \in \mathcal{U} \setminus S$ may return any value from $\{0, 1\}^r$. Membership queries ("is $x \in S$?") are not supported.

Static Functions from Linear Systems. A well-known way for constructing static functions [3, 10, 14, 20, 37, 53] uses a hash function to associate each key $x \in \mathcal{U}$ with a set $h(x) \subseteq [m]$ for some $m \geq n = |S|$. By $\vec{h}(x) \in \{0, 1\}^m$ we denote the characteristic (row) vector of $h(x)$. If $(\vec{h}(x))_{x \in S}$ are linearly independent in the vector space $\{0, 1\}^m$ over the two-element field then the system $(\vec{h}(x) \cdot Z = b(x))_{x \in S}$ of linear equations has a solution $Z \in \{0, 1\}^{m \times r}$. The static function is then given by h and Z. Most memory is used for the mr bits of Z, which takes $\frac{m}{n}r$ bits per key. A query for $x \in \mathcal{U}$ returns

$$\text{query}(x) := \vec{h}(x) \cdot Z = \bigoplus_{i \in h(x)} Z_i. \tag{1}$$

where Z_i denotes the i-th row of Z. Since queries involve $|h(x)| \cdot r$ bits from Z fast query times require sparse $\vec{h}(x)$.

Several constructions choose h such that $\vec{h}(x)$ contains exactly three 1-bits in random positions [10, 37]. In this case $\frac{n}{m}$ must not exceed the corresponding XORSAT threshold [19, 52] $c_3^* \approx 0.92$. If a greedy algorithm is used for solving the linear system, then $\frac{n}{m}$ must not exceed the *peeling threshold* $c_3^\Delta \approx 0.82$.

The space usage is roughly $(1+\varepsilon)r$ bits per key when $m = (1+\varepsilon)n$. The first paper to achieve $\varepsilon = o(1)$ is [53]. Even $\varepsilon = O(\log n/n)$ is possible, albeit with mediocre construction time [21]. A recent more practical contribution that (more humbly) aims for small *constant* $\varepsilon > 0$ [22] will be the starting point of our own construction.

Xor filters. There is a straightforward way to obtain a filter with FP rate 2^{-r} from an r-bit static function as pointed out in [20, Observation 1]. Simply pick a random fingerprint function $b : \mathcal{U} \to \{0, 1\}^r$ (a hash function) and store its restriction $b_S : S \to \{0, 1\}^r$ as a static function. Then for any $x \in S$ the static function reproduces $b_S(x) = b(x)$ while for $x \in \mathcal{U} \setminus S$ the returned value will match $b(x)$ only with probability 2^{-r} (because $b(x)$ is random and plays no role in the construction of the static function).

Such filters inherit the performance of the underlying static function, giving them the potential to be "Faster and Smaller Than Bloom and Cuckoo Filters" as claimed in [39], when dynamic insertions and deletions are not required. A standard construction with $h(x)$ being a fully random set of size 3 is appropriately named *Xor filter* [39].

3 RIBBON RETRIEVAL AND RIBBON FILTERS

By enriching the SGAUSS static function[10] from [22], we obtain the *Ribbon static function* which can be used as a *Ribbon filter*. Since

---

[10]Strictly speaking, SGAUSS is the name of the construction algorithm of the otherwise unnamed data structure.
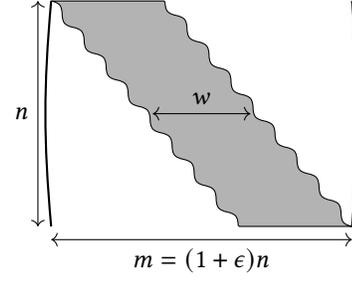
Figure 2: Typical shape of the random matrix with rows $(\vec{h}(x))_{x \in S}$ sorted by starting positions. The shaded "ribbon" region contains random bits. Gaussian elimination never causes any fill-in outside of the ribbon.

Ribbon filters also retrieve fingerprints using Equation (1)—just with a different choice of h—they can be seen as (non-standard) Xor filters.

The SGAUSS construction. For a parameter $w \in \mathbb{N}$ that we call the *ribbon width*, the vector $\vec{h}(x) \in \{0, 1\}^m$ is given by a random *starting position* $s(x) \in [m - w - 1]$ and a random *coefficient vector* $c(x) \in \{0, 1\}^w$ as $\vec{h}(x) = 0^{s-1}c(x)0^{m-s-w+1}$. Note that even though m-bit vectors like $\vec{h}(x)$ are used to simplify mathematical discussion, such vectors can be represented using $\log(m) + w$ bits.

The matrix with rows $(\vec{h}(x))_{x \in S}$ sorted by $s(x)$ has all of its 1-entries in a "ribbon" of width w that randomly passes through the matrix from the top left to the bottom right, as in Figure 2. The authors of [22] showed that a solution $Z \in \{0, 1\}^{m \times r}$ to $(\vec{h}(x) \cdot Z = b(x))_{x \in S}$ can be computed quickly:

THEOREM 3.1 ([22, THM 2]). *For any constant* $0 < \varepsilon < \frac{1}{2}$, $\vec{w} = \frac{\log n}{\varepsilon}$ *and* $\frac{n}{m} = 1 - \varepsilon$, *with high probability the linear system* $(\vec{h}(x) \cdot Z = b(x))_{x \in S}$ *is solvable for any* $r \in \mathbb{N}$ *and any* $b : S \to \{0, 1\}^r$. *Moreover, after sorting* $(\vec{h}(x))_{x \in S}$ *by* $s(x)$, *Gaussian elimination can compute a solution* Z *in expected time* $O(n/\varepsilon^2)$.

Boolean banding on the fly. For Ribbon we start with the same hash function $\vec{h}$ as in SGAUSS. For slightly improved presentation, execution speed, and chance of construction success, we force coefficient vectors $c(x)$ to start with 1.[11]

The main difference lies in how we solve the linear system. The *insertion phase* maintains a reduced system M of linear equations using on-the-fly Gaussian elimination [8]. This system is of the form shown in Figure 3 and has m rows. Each row i is represented by a w-bit vector $c_i \in \{0, 1\}^w$ and $b_i \in \{0, 1\}^r$. Logically, the i-th row is either empty ($c_i = 0^w$) or specifies a linear equation $c_i \cdot Z_{[i,i+w)} = b_i$ where c_i starts with a 1. With $Z_{[i,i+w)} \in \{0, 1\}^{w \times r}$ we refer to rows $i, \dots, i + w - 1$ of Z. We ensure $c_i \cdot Z_{[i,i+w)}$ is well-defined even when $i + w - 1 > m$ with the invariant that c_i values never select "out of bounds" rows of Z.

We consider the equations $(\vec{h}(x) \cdot Z = b(x))_{x \in S}$ one by one, in arbitrary order, and try to integrate each into M using Algorithm 1, which we explain now. A key's equation may be modified several

---

[11]In asymptotic considerations this change is inconsequential (and mildly annoying). For better alignment with [22] our theorems still assume that $c(x)$ is uniformly distributed in $\{0, 1\}^w$.
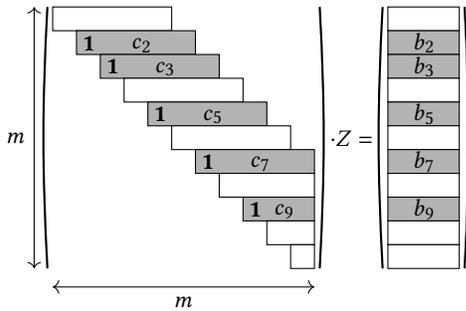
Figure 3: Shape of the linear system M central to Boolean banding on the fly.

Algorithm 1: Adding a key's equation to the linear system M.

```
1  i ← s(x)
2  c ← c(x)
3  b ← b(x)
4  loop
5  |   if M.c[i] = 0 then  // row i of M is empty
6  |   |   M.c[i] ← c
7  |   |   M.b[i] ← b
8  |   |   return SUCCESS (inserted)
9  |   c ← c ⊕ M.c[i]
10 |   b ← b ⊕ M.b[i]
11 |   if c = 0 then
12 |   |   if b = 0 then  return SUCCESS (redundant)
13 |   |   else  return FAILURE (inconsistent)
14 |   j ← findFirstSet(c)  // a.k.a. BitScanForward
15 |   i ← i + j
16 |   c ← c >> j  // logical shift last toward first
```

---

times before it can be added to M, but a loop invariant is that its form is

$$c \cdot Z_{[i,i+w)} = b \text{ for } i \in [m], c \in 1 \circ \{0,1\}^{w-1}, b \in \{0,1\}^r. \quad (2)$$

The initial equation $\vec{h}(x) \cdot Z = b(x)$ of key $x \in S$ has this form with $i = s(x)$, $c = c(x)$ and $b = b(x)$. We proceed it as follows.

Case 1: In the simplest case, row i of M is empty and we can incorporate Equation (2) as the new i-th row of M.

Case 2: Otherwise row i of M is already occupied by an equation $c_i \cdot Z_{[i,i+w)} = b_i$. We compute the summed equation

$$c' \cdot Z_{[i,i+w)} = b' \text{ with } c' = c \oplus c_i \text{ and } b' = b \oplus b_i, \quad (3)$$

which, in the presence of row i of M, puts the same constraint on Z as Equation (2). Both c and c_i start with 1, so c' starts with 0. We consider the following sub-cases.

Case 2.1: $c' = 0^w$ and $b' = 0^r$. The equation is void and can be ignored. This case is reached when the key's original equation is implied by equations previously added to M.

Case 2.2: $c' = 0^w$ and $b' \neq 0^r$. The equation is unsatisfiable. This case is reached when the key's original equation is inconsistent with equations previously added to M.

Case 2.3: c' starts with $j > 0$ zeroes followed by a 1. Then Equation (3) can be rewritten as $c'' \cdot Z_{[i',i'+w)} = b'$ where $i' = i + j$ and c'' is obtained from c' by discarding the j leading zeroes of c and appending j trailing zeroes.

Note that in the bit-shift of Algorithm 1 the roles of "leading" and "trailing" may seem reversed because the least-significant "first" bit of a word is conventionally thought of as the "right-most" bit.

Termination is guaranteed since i increases with each loop iteration.

Once equations for all keys are successfully inserted, we obtain a solution Z to M in the *back substitution phase*. The rows of Z are obtained from bottom to top. If row i of M contains an equation then this equation uniquely determines row i of Z in terms of later rows of Z. If row i of M is empty, then row i of Z can be initialized arbitrarily.

"On-the-fly" and "incremental." The insertion phase of ribbon is on-the-fly in the sense that for a sequence $S = (x_1, x_2, x_3, \dots)$ of keys we can easily determine the longest prefix $(x_1, \dots, x_n)$ of S for which construction succeeds: Simply insert keys until the first failure. The insertion phase is incremental because we can easily

undo a set of most recent successful insertions: Simply remove the rows from M that were added. These properties are not shared by SGAUSS and will be exploited in Section 6.

Analysis. All performance guarantees for the construction algorithm carry over from SGAUSS as follows.

THEOREM 3.2. *Let $S \subseteq \mathcal{U}$ be an arbitrary key set.*

(i) *If an SGAUSS construction succeeds for S then so does the Ribbon construction.*

(ii) *If both constructions succeed on S, expected running times coincide up to constant factors.*

PROOF.

(i) This is unsurprising as both approaches attempt to solve the same linear system.[12] A superficial difference concerns redundant equations. In Algorithm 1 it is natural to ignore them. SGAUSS treats them as failures, to avoid special cases during back-substitution.

(ii) Consider a set S on which SGAUSS succeeds, i.e. S gives rise to a solvable system without redundant equations. The back-substitution phases are identical in both algorithms. The analysis of the Ribbon insertion phase hinges on counting row additions. Each key $x \in S$ has a starting position $s(x)$ and causes some row $i(x)$ of M to be filled. The number of row additions for the insertion is clearly at most the *displacement* $i(x) - s(x)$ of x. Summing over all keys yields

$$D = \sum_{x \in S} i(x) - s(x) = \sum_{i \in P} i - \sum_{x \in S} s(x)$$

where $P = \{i(x) \mid x \in S\}$ is the set of row-indices of M that end up being occupied. A crucial observation is that even though the values $i(x)$ depend on the insertion order of the keys, the set P does not. Indeed, for any $j \in [m]$ the value $|P \cap \{1, \dots j\}|$ is the rank of the sub-matrix of M formed by its first j columns, which is invariant under row operations. So

---

[12]See Footnote 11.

no matter in what order the keys of S are inserted, we always observe the same sets $P \cap \{1, \dots j\}$ and hence the same set P and the same value D. The number of row additions of SGAUSS is bounded by D by a similar argument (see [22, Lemma 3]). The analysis in [22] proceeds by bounding D in expectation and hence carries over to our case. □

Efficiency. While SGAUSS and Ribbon are tied in O-notation, Ribbon improves upon SGAUSS in constant factors for the following reasons:

- There is no need to pre-sort the keys by $s(x)$.
- SGAUSS requires explicitly storing a pivot position for each row. This is because SGAUSS does not compute an echelon form but only ensures that in each row the left-most 1-entry — the pivot — is the bottom-most 1-entry of its column.
- SGAUSS performs roughly D elimination steps that, depending on some bit, turn out to be xor-operations or no-ops. Ribbon on the other hand performs roughly $D/2$ bit shifts and $D/2$ (unconditional) xor operations. Though the details are complicated, intuition on branching complexity seems to favour ribbon.

3.1 Ribbon with Smash

When aiming for high space efficiency, there is an issue with early and late columns of the linear system. We shall describe the problem and its solution in an extreme but simple case where perfect space efficiency, i.e. $m = n$ is desired.

In the absence of redundant equations, the construction of the linear system succeeds only if all slots of M can be filled. For the first i slots to be filled, it is necessary that $|\{x \in S \mid s(x) \leq i\}| \geq i$. Figure 4 **(a)** illustrates that random fluctuations make this unlikely. There is a similar problem relating to the last i columns[13].
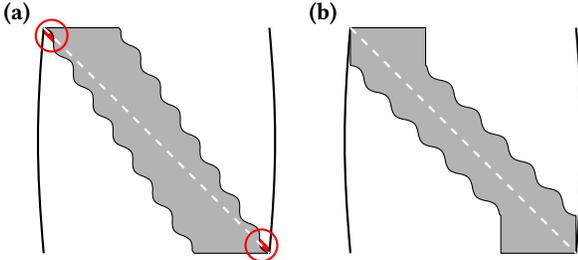
Figure 4: (a) Consider the diagonal (dashed) in a square Ribbon system. Its beginning and end may lie outside of the of the (shaded) ribbon area.
(b) Our "smash" variant solves this problem.

We can address this issue by artificially inflating the probabilities $\Pr[s(x) = 1]$ and $\Pr[s(x) = m - w + 1]$ of the first and last starting position by a factor of ℓ we call the *smash value*. Such a distribution for s is easy to implement using a uniform distribution on $[-\ell + 2, m - w + \ell]$ and "clamping" the sampled value to $[1, m - w + 1]$ using min and max functions. Micro-benchmarks show roughly 3ns overhead per query for smash, on an Intel Skylake CPU.

For a smash value of $\ell = w/2$ and a ribbon width of $w = \omega(\sqrt{n})$ the matrix diagonal is firmly within the ribbon, see Figure 4 **(b)**. It is not hard to prove that such a matrix is asymptotically as likely to be regular as a fully random $n \times n$ matrix. That probability is $c_2 \approx 0.289$; see [16].

In Section 5.4 we present empirical findings showing that non-zero smash values also benefit success probabilities in practically more relevant cases with $\varepsilon > 0$ and $w = O(\frac{\log n}{\varepsilon})$.

4 HOMOGENEOUS RIBBON FILTERS

Recall that the idea underlying Ribbon filters is to pick hash functions $\vec{h} : \mathcal{U} \rightarrow \{0, 1\}^m$, $b : \mathcal{U} \rightarrow \{0, 1\}^r$ and find $Z \in \{0, 1\}^{m \times r}$ such that all $x \in S$ satisfy $\vec{h}(x) \cdot Z = b(x)$, while most $x \in \mathcal{U} \setminus S$ will not.

We now examine what happens when we get rid of the fingerprint function b effectively setting $b(x) = 0$ for all $x \in \mathcal{U}$. A filter is then given by a solution Z to the *homogeneous* system $(\vec{h}(x) \cdot Z = 0^r)_{x \in S}$. The FP rate for Z is $f_Z = \Pr_{a \sim H}[a \cdot Z = 0^r]$ where H is the distribution of $\vec{h}(x)$ for $x \in \mathcal{U}$. An immediate issue with the idea is that $Z = 0^{m \times r}$ is a solution giving $f_Z = 1$. A solution Z chosen *uniformly at random* from all solutions fares better, however. To obtain one, all free variables, i.e. the variables corresponding to empty rows of M, are initialized randomly during back substitution.[14] The overall FP rate is then $f = \mathbb{E}[f_Z]$ where Z depends on the randomness in $(\vec{h}(x))_{x \in S}$ and the free variables.

We call the resulting construction *Homogeneous Ribbon filter*. It has two obvious advantages over Standard Ribbon filters:

- Constructions can never fail, regardless of n, ε and w. This is simply because a homogeneous linear system always has at least the trivial solution.
- The absence of fingerprints slightly improves time and space in construction. (In optimized implementations, query times are essentially the same.)

A complication is that the FP rate f can be higher than 2^{-r}. Intuitively, if too many equations constrain some part of Z then that part will be insufficiently random. For $\varepsilon w > Cr$ (for some constant $C > 0$) and large n, this effect is negligible as we shall argue in Theorem 4.1. This still leaves us with two disadvantages, especially for small n and high r:

- The product εw must be proportional to r (for acceptable f) whereas for Standard Ribbon εw need only be proportional to $\log n$ (for acceptable success probability). We should therefore not expect an improvement over Standard Ribbon for $r = \Omega(\log n)$.
- For small n, the variance of f_Z is quite high, meaning small filters will occasionally have significantly more false positives (e.g. due to random skew in $(s(x))_{x \in S}$) with no obvious way to detect this during construction. This could be dangerous for some applications.

4.1 Analysis

We can make a strong case for Homogeneous Ribbon filters by showing that arbitrarily small space overhead at arbitrarily large

---

[13]To see the symmetry, we could have argued about the rank of the first i columns of M which is *at most* $|\{x \in S \mid s(x) \leq i\}|$.

[14]Our implementation uses trivial pseudo-random assignments: a free variable in row i is assigned $pi \bmod 2^r$ for some fixed large odd number p.

size n is achievable. This neither requires the ribbon width w to scale with n, nor a deviation from the pure construction (e.g. by partitioning the key set into small shards).

By space overhead we mean $\frac{\text{SPACE}}{\text{OPT}} - 1$ where SPACE is the space usage of the filter in bits per key and $\text{OPT} = -\log_2(f)$ is the information-theoretic lower bound for filters that achieve the same FP rate.

THEOREM 4.1. *There exists $C \in \mathbb{R}^+$ such that for any $\varepsilon < 1/2$, any desired FP rate $\varphi > 0$ and r the closest integer to $-\log_2(\phi)$ the following holds. For any $w \in \mathbb{N}$ with $\varepsilon w > C \max(\log w, r)$ and $n \in \mathbb{N}$ the Homogeneous Ribbon filter with n keys and parameters ε, w, r has $f \in [\varphi/2, 2\varphi]$ and space overhead at most 2ε.*

Our argument starts with the following simple observation.

LEMMA 4.2. *In the context of a Homogeneous Ribbon filters let p be the probability that for $y \in \mathcal{U} \setminus S$ the vector $\vec{h}(y)$ is in the span of $(\vec{h}(x))_{x \in S}$. Then we have*

$$f = p + (1-p)2^{-r}.$$

PROOF. First assume there exists $S' \subseteq S$ with $\vec{h}(y) = \sum_{x \in S'} \vec{h}(x)$ which happens with probability p. In that case

$$\vec{h}(y) \cdot Z = \left(\sum_{x \in S'} \vec{h}(x) \right) \cdot Z = \sum_{x \in S'} (\vec{h}(x) \cdot Z) = 0$$

and y is a false positive. Otherwise, i.e. with probability $1 - p$, an attempt to add $\vec{h}(y) \cdot Z = 0$ to M after all equations for S were added would have resulted in a (non-redundant) insertion in some row i. During back substitution, only one choice for the i-th row of Z satisfies $\vec{h}(y) \cdot Z = 0$. Since the i-th row was initialized randomly we have $\Pr[\vec{h}(y) \cdot Z = 0 \mid \vec{h}(y) \notin \text{span}((\vec{h}(x))_{x \in S})] = 2^{-r}$. □

We shall now derive an asymptotic bound on p in terms of large w and small ε (recall $\varepsilon = \frac{m-n}{n}$). It is too imprecise to estimate p and f in practical settings, which we do empirically in Section 4.2. The main takeaway is rather that p does not depend on n. We may therefore expect Homogeneous Ribbon filters to scale to arbitrary sizes n with no increase in f even when w and ε are constants.

LEMMA 4.3. *There exists a constant C such that for any $w \in \mathbb{N}$ and $C\frac{\log w}{w} \leq \varepsilon \leq \frac{1}{2}$ we have $p = \exp(-\Omega(\varepsilon w))$.*

The main ingredient in the proof is that $\exp(-\Omega(\varepsilon w))$ bounds the number of keys that cannot be (non-redundantly) inserted, which follows from [22].[15]

PROOF SKETCH. We may imagine that $S \subseteq \mathcal{U}$ and $y \in \mathcal{U} \setminus S$ are obtained from a set $S^+ \subseteq \mathcal{U}$ of size $n + 1$ by picking $y \in S^+$ at random and setting $S = S^+ \setminus \{y\}$. Then p is simply the expected fraction of keys in S^+ that are contained in some *dependent set*, i.e. in some $S' \subseteq S^+$ with $\sum_{x \in S'} \vec{h}(x) = 0^m$. Clearly, x is contained in a dependent set if and only if it is contained in a *minimal* dependent set. Such a set S' always touches a consecutive set of positions, i.e. $\text{pos}(S') := \bigcup_{x \in S'}[s(x), s(x) + w - 1]$ is an interval.

We call an interval $I \subseteq [m]$ *long* if $|I| \geq w^2$ and *short* otherwise. We call it *overloaded* if $S_I := \{x \in S^+ \mid s(x) \in I\}$ has size $|S_I| \geq |I| \cdot (1 - \varepsilon/2)$. Finally, we call a position $i \in [m]$ *bad* if one of the following is the case:

(B1) i is contained in a long overloaded interval.
(B2) $i \in \text{pos}(S')$ for a minimal dependent set S' with long non-overloaded interval $\text{pos}(S')$.
(B3) $i \in \text{pos}(S')$ for a minimal dependent set S' with short interval $\text{pos}(S')$.

We shall now establish the following

Claim: $\forall i \in [m] : \Pr[i \text{ is bad}] = \exp(-\Omega(\varepsilon w))$.

For each $i \in [m]$ the contributions from each of the badness conditions (B1,B2,B3) can be bounded separately. In all cases we use our assumption $\varepsilon \geq C\frac{\log w}{w}$. It ensures that $\exp(-\Omega(\varepsilon w))$ is at most $\exp(-\Omega(\log w)) = w^{-\Omega(1)}$ and can "absorb" factors of w in the sense that by adapting the constant hidden in Ω we have $w \exp(-\Omega(\varepsilon w)) = \exp(-\Omega(\varepsilon w))$.

(B1) A Chernoff bound for sums $X = \sum_j X_j$ of i.i.d. indicator random variables with $\mu = \mathbb{E}[X]$ is

$$\Pr[X \geq (1+\delta)\mu] \leq \exp(-\delta^2\mu/3). \qquad (4)$$

We use it in a case where I is an interval and $X_1, \ldots, X_{n+1}$ indicate which of the keys in S^+ have a starting position within I. For $n \gg w$ we have

$$\mu = \mathbb{E}[X] \leq \frac{(n+1)|I|}{m-w+1} \approx \frac{n|I|}{m} = |I|/(1+\varepsilon).$$

Skipping over some uninteresting details, the probability for I to be overloaded is (for $n \gg w$)

$$\Pr[X \geq (1 - \varepsilon/2)|I|] \leq \Pr[X \geq (1 + \underbrace{\varepsilon/6}_{\delta}) \underbrace{|I|/(1+\varepsilon)}_{\geq \mu}]$$
$$\overset{\text{Eq. 4}}{\leq} \exp\left(\frac{-\varepsilon^2|I|}{108(1-\varepsilon)}\right). \qquad (5)$$

The probability for $i \in [m]$ to be contained in a long overloaded interval is bounded by the sum of Equation (5) over all lengths $|I| \geq w^2$ and all $|I|$ offsets that I can have relative to i. The result is of order $\exp(-\Omega(\varepsilon^2 w^2))$ and hence small enough.

(B2) Consider a long interval I that is not overloaded, i.e. $|I| \geq w^2$ and $|S_I| \leq (1 - \varepsilon/2)|I|$. There are at most $2^{|S_I|}$ sets S' of keys with $\text{pos}(S') = I$ and each is a dependent set with probability $2^{-|I|}$ because each of the $|I|$ positions of I that S' touches imposes one parity condition.
The probability for I to support at least one dependent set is therefore at most $2^{-|I|} \cdot 2^{|S_I|} = 2^{-\frac{\varepsilon}{2}|I|} = \exp(-\Omega(\varepsilon|I|))$.
Similar as in (B1) for $i \in [m]$ we can sum this probability over all admissible lengths $|I| \geq w^2$ and all offsets that i can have in I to bound the probability that i is bad due to (B2).

(B3) Let $S_{\text{red}} \subseteq S$ be the set of *redundant* keys, i.e. keys for which Algorithm 1 returns "SUCCESS (redundant)". While S_{red} depends on the insertion order, the *rank defect* $|S_{\text{red}}| = n - \text{rank}((\vec{h}(x))_{x \in S})$ does not. A central step in [22] implies that $\mathbb{E}[|S_{\text{red}}|] = m \cdot \exp(-\Omega(\varepsilon w))$.[16]
Now if i is bad due to (B3) then $i \in \text{pos}(S')$ for some minimal dependent set S' with short $\text{pos}(S')$. At least one key from S' is redundant (even if all keys from $\text{pos}(S')$ are inserted first). In particular, i is within short distance ($< w^2$) of the starting position of a redundant key x. Therefore at most $|S_{\text{red}}| \cdot$
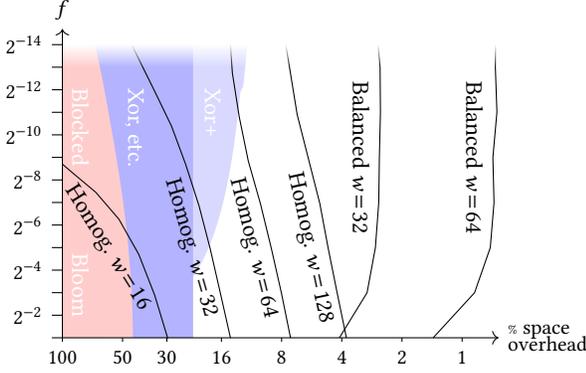
Figure 5: Combinations of space overhead and false positive rate achievable with various ribbon widths w and large n. (Experimental data for Ribbon configurations use $n = 10^6$ but generalize.)

$2w^2$ positions are bad due to (B3), which is an $\exp(-\Omega(\varepsilon w))$-fraction of all positions as desired.

With a concentration argument the following variant of the claim can be proved. We omit the details.

Claim': $\forall x \in S^+ : \Pr[s(x) \text{ is bad}] = \exp(-\Omega(\varepsilon w))$.

Now assume that a key $x \in S^+$ is contained in a minimal dependent set S'. It follows that all of $pos(S')$ is bad. Indeed, either $pos(S')$ is a short interval ($\rightarrow$ B3) or it is long. If it is long, then it is overloaded ($\rightarrow$ B1) or not overloaded ($\rightarrow$ B2). In any case $s(x) \in pos(S')$ is bad.

Therefore, the probability p for $y \in S^+$ to be contained in a dependent set is at most the probability for $s(y)$ to be bad. This is upper-bounded by $\exp(-\Omega(\varepsilon w))$ using Claim'. $\qquad\square$

We are now ready to prove Theorem 4.1.

PROOF OF THEOREM 4.1. We shall find that our filter has $f \leq 2^{-r}(1 + \varepsilon^2)$ and hence $f \in [\varphi/2, 2\varphi]$ with high probability.

The space consumption of $Z \in \{0, 1\}^{m \times r}$ is $\text{SPACE} = \frac{mr}{n} = r(1+\varepsilon)$ bits per key. To relate this to OPT, we need bounds on p.

By assumption $\varepsilon w > C \log w$, so a large enough choice for C permits the use of Lemma 4.3, which guarantees $p = \exp(-\Omega(\varepsilon w))$. Again using $\varepsilon w > C \max(\log w, r)$ for large enough C gives

$$p \leq \exp(-2\log(w) - r) \leq \frac{1}{w^2} e^{-r} \leq \varepsilon^2 2^{-r}. \qquad (6)$$

Together with Lemma 4.2 we get

$$\text{OPT} = -\log_2(f) \overset{\text{Lem 4.2}}{=} -\log_2(p + (1-p)2^{-r})$$
$$\geq -\log_2(p + 2^{-r}) \overset{\text{Eq. 6}}{\geq} -\log_2(2^{-r}(1 + \varepsilon^2))$$
$$= r - \log_2(1 + \varepsilon^2) \geq r - \varepsilon^2.$$

Putting everything together yields

$$\frac{\text{SPACE}}{\text{OPT}} = \frac{r(1 + \varepsilon)}{r - \varepsilon^2} \leq \frac{(1 + \varepsilon)}{(1 - \varepsilon^2)} \leq 1 + 2\varepsilon.$$

The last step makes use of $\varepsilon < \frac{1}{2}$. $\qquad\square$

4.2 Space efficiency in practice

To minimize the space overhead $r(1 + \varepsilon)/\log f^{-1} - 1$ for chosen values of r and w in a Homogeneous Ribbon filter, ε must be neither too large or too small. (Small ε causes f to explode due to densely-packed constraints.) To choose ε, we turn to simulations on random data, building large structures and testing the FP rate. Using $m = 3 \cdot 10^7$ (among others), $w \in \{16, 32, 64, 128\}$, and $r \in [1, 16]$, we see a pretty clear pattern in where the space overhead is minimized for any large n:

$$\varepsilon \approx \frac{4 + r/4}{w}. \qquad (7)$$

Note how the recommendation $\varepsilon > C \max(\log w, r)/w$ one might derive from Theorem 4.1 (vaguely) agrees. We use (7) in all experiments. In Figure 5 we show combinations of space overhead and FP rate to expect from Homogeneous Ribbon filters for $w \in \{16, 32, 64, 128\}$ and large n.[17] Balanced Ribbon filters, also in Figure 5, are discussed in Section 6.

For example, consider using $r = 7$ for roughly 1% FP rate and $w = 64$ for reasonable space-vs.-time trade-off. Using $\varepsilon \approx 0.09$ from Equation (7) we observe $f \approx 0.81\% > 0.78\% \approx 2^{-7}$, so actual space overhead is closer to 10% than the 9% allocated with ε.

5 MAKING RIBBON PRACTICAL

5.1 Configurability and elasticity

A useful and perhaps previously unreplicated feature of Bloom filters is the ability to efficiently utilize any amount of space for minimizing the FP rate in representing any number of keys. We call this *configurability* and suggest it is practically important for space efficiency. Consider an application with little control over the number of keys going into a filter. Even if we use a perfectly space efficient filter for a specific FP rate, we could be wasting significant space due to internal fragmentation from an allocator. A memory allocator like jemalloc [33] averages about 10% internal fragmentation on arbitrarily sized allocations, space that should ideally be used by the filter to reduce its FP rate [18].

More specifically, Bloom alternatives such as Cuckoo, Quotient, and Xor conventionally use cells of some whole number of bits, as Ribbon does with r. Whole number r limits space-efficient choices of FP rates and bits per key. For example, using only 5-bit cells when 5.5 bits per cell is available adds roughly 10% space overhead to our filter. Other than Ribbon, which is tied to the two-element field, these same Bloom alternatives can use fractional-bit cell sizes. Some configurations can even be made efficient, such as $64/i$ bits per cell for whole i (an existing Xor10.666 implementation is tested in Section 7), but fine granularity would surely be more CPU intensive.

An alternative way of generalizing to *effectively* fractional r is to split available space between two structures: one using $\lceil r \rceil$ solution columns (or bits per cell) and the other using $\lfloor r \rfloor$ for a weighted average of r, chosen to fit available space. This only slightly increases the overall space overhead. For example, using $r = 5.5$ yields (non-homogeneous) FP rate of $3/128$. The lower bound for this rate is $\lambda = 5.415$ bits per key, so the approach adds 1.57% to

---

[17] Standard Ribbon filters are not shown because the achievable overhead depends on n. We do recommend Standard Ribbon (with smash) for small $n < 10^4$ where Homogeneous Ribbon filters have high variance in FP rate.

[18] See RocksDB's `optimize_filters_for_memory` option [27].

overall space overhead. This addition is a larger 6.00% for $r = 1.5$, or smaller 0.82% for $r = 10.5$ bits per key. Practical concerns with splitting into two structures includes (a) essentially doubling many of the space usage penalties associated with small structures (when applicable), and (b) independently seeding hashes or accepting joint construction success probability (when applicable).

For Ribbon we recommend a variant of that approach within a single structure: using only $\lfloor r \rfloor$ solution bits per row for some prefix of rows, and $\lceil r \rceil$ bits per row for the rest of rows. The banding process (Algorithm 1) is unchanged, but small changes are needed to back-substitution and query (more details in Section 5.2). Because of Ribbon's locality of probes in queries, unlike standard Xor filters, a diminishingly small w/m portion of queries cross the boundary between $\lfloor r \rfloor$ and $\lceil r \rceil$ columns (utilizing only $\lfloor r \rfloor$ in such cases), so space efficiency is very close to the idealized split approach, and probably better in practice: around 1% additional space overhead for common configurations (see e.g. $r = 7.7$ in Section 7).

Although the split approach enables near-continuous configurability for many kinds of filters, the single-structure approach for Ribbon filters has an advantage we call *elasticity*, for applications like ElasticBF [46]. Like an Xor filter, one can drop entire columns from a Ribbon filter, for a corresponding higher FP rate. With Ribbon filters, we also have the ability to drop part of the last column, so a finished filter can be trimmed down with bit granularity. Generalizing further, a finished Ribbon filter can be split to several smaller structures with independent FP rates[19], with product as small as the FP rate of the starting structure. Similarly, a finished Ribbon filter could be physically (re-)partitioned at arbitrary boundaries by duplicating as little as $(w - 1)r$ bits at each partition boundary such that each query accesses only one partition.

5.2 Solution structure layout

Here we examine memory layouts for the solution matrix $Z \in \{0, 1\}^{m \times r}$, which is critical for fast Ribbon filter queries. Prior work [22] only evaluated the $r = 1$ case, where Z is a Boolean (bit) vector.

Xor filters conventionally use **row-major layout** of the solution structure, wherein the whole row i of Z immediately precedes the whole row $i + 1$ in memory. A $w = 64$ Ribbon filter combines roughly an order of magnitude more rows, conditionally, than a standard Xor filter combines unconditionally (three rows). In fact, with number of solution columns commonly $5 \leq r \leq 15$, standard Xor filters typically access more columns than rows, while Ribbon filters typically access more rows than columns. Although row-major layout could likely be made efficient for Ribbon in some special cases using SIMD, we do not find it generally workable for a fast and highly-configurable filter.

The opposite is **column-major layout**, in which the entire column i of Z precedes column $i+1$ in memory. Column-major is essentially ideal for querying a single result bit for a key, as we only have to access a "contiguous" (usually unaligned) w bits, bitwise-AND with $c(x)$, and get the bit parity. The problem with column-major is that accessing more result bits is not an adjacent memory access. Although the several memory addresses are easily computed and

---

[19]For Homogeneous Ribbon, the portion of the FP rate from degradation is not independent.
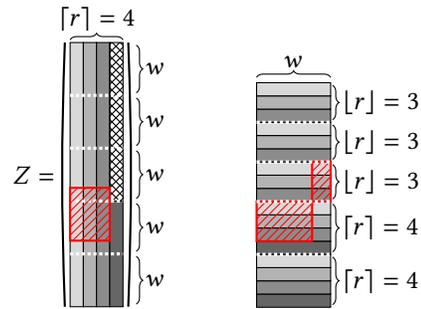
Figure 6: A solution matrix (left) and w-bit interleaved column-major layout (ICML, right) of that matrix, mixing $\lfloor r \rfloor$ and $\lceil r \rceil$ columns as in Section 5.1. The shaded region shows the bits used in a query crossing the boundary between $\lfloor r \rfloor$ and $\lceil r \rceil$.

can be fetched in parallel, some testing shows this to be relatively expensive for $r > 2$ for a filter that is not hot in memory cache.

Our preferred solution layout for Ribbon filters is **interleaved column-major layout** (ICML), because it has locality very close to row-major and decoding efficiency very close to column-major. The memory space is divided into conveniently sized ICML words, and grouped into blocks of r words. Each block is the column-major layout of some contiguous rows of Z. See Figure 6, which generalizes this layout to a mixed number of columns for fractional r.

For ICML word size equal to w, at most and almost always two words are combined for reconstructing each result bit. This means that the amount of adjacent memory accessed for a full query is $2rw$ bits, while the ideal minimum is rw bits. For example, with $r = 6$ and $w = 64$, ICML accesses 768 bits per query, with 384-bit alignment, which translates to an average of accessing 2.25 Intel cache lines (512 bits) and essentially 1 page (4KB) per query; with only 6-bit "alignment," row-major would access 1.5 cache lines per query. A standard Xor filter accesses essentially 3 cache lines and nearly as many pages per query.

Ribbon back-substitution is an especially fast, streaming operation for layouts based on column major. We can buffer w rows of Z in r temporary values of width w, likely fitting in CPU registers, and use those buffers for (a) computing the logical previous bit for each column, and (b) flushing to our solution structure for each w rows (rw bits).

As is well known for Bloom filters, queries can potentially be optimized with short-circuiting: returning from a "negative" query as soon as a probed bit is zero, ensuring the query must return FALSE. A similar approach works for Ribbon filters using layouts based on column-major, returning as soon as a result bit does not match expectation. Although cache-local Bloom filters are so optimized that this approach rarely pays off any longer [44], our Ribbon implementation uses short-circuiting except for compile-time fixed $r \leq 4$. The distinction is visible in observed query time ranges in Section 7.

We also like the clean configurability of layouts based on column-major. Parameter r should be freely chosen to balance FP rate vs.

space usage, and that choice is much more free when it does not affect instruction-level data alignment, only alignment in CPU caches and pages, which we consider a relatively minor concern. Ribbon-width ICML is good for using $\lfloor r \rfloor$ columns before $\lceil r \rceil$ columns (Section 5.1), because upon determining the starting memory location and (smallest) applicable number of columns, which can be done without conditional branches, the remaining query code does not have to be aware of mixed numbers of columns; see Figure 6. ($\lceil r \rceil$ before $\lfloor r \rfloor$ is better for pure column major.)

A minor disadvantage of ICML is that the number m of solution rows must be a multiple of the number of bits in an ICML word, which can present a conflict between configurability (accommodating any number of keys) and space efficiency for small n.

5.3 Practical hashing for Ribbon

Hash expansion. A filter structure "consuming" some quantity of hash information can operate from a smaller hash $\mathcal{H}(x)$ of the original key x [28, 29]. The practical requirements for Ribbon filters are these:

- $\mathcal{H}(x)$ values must be large enough to have an insignificant baseline FP rate due to full hash collisions, i.e. $\mathcal{H}(x_1) = \mathcal{H}(x_2)$ for $x_1 \neq x_2$. A 64-bit hash for $\mathcal{H}$ should suffice for almost all non-cryptographic applications, as having 2^{32} keys in a single filter incurs a baseline FP rate of just 2^{-32}.
- $\mathcal{H}(x)$ is effectively extended / expanded / remixed to what is consumed[20]. For Ribbon, it is most important to minimize correlations between the starting location and other hash consumers. A starting location computed with fastrange [45] on $\mathcal{H}(x)$ relies primarily on upper bits, so multiplying $\mathcal{H}(x)$ by a large odd constant (as with Knuth multiplicative hash [42]) seems to suffice for removing correlation. Details are in the reference implementation of Ribbon [25].

Re-seeding. Some Ribbon designs need the ability to retry construction with sufficiently independent hashing to have an independent probability of construction success. Observe that for Ribbon filters a full hash collision does not interfere with construction success (it only produces a redundant equation). We find in significant testing that modifying an unseeded stock hash value with simple XOR with a pseudorandom seed then multiplication by a large odd constant suffices for independent probability of construction success. See [25] for details. Assuming uniform hashes, an effective alternative to re-seeding on failed construction is simply to increase m by a factor of $\frac{w+1}{w}$.

5.4 Standard Ribbon scalability

We refer to the non-homogeneous Ribbon construction, including *smash* when appropriate, as *Standard Ribbon*. Construction fails with some probability depending on m, n and w, though we have no formula. Table 1 provides some empirical data points for how much configured space overhead, $\varepsilon = (m - n)/n$, is required for several failure probabilities that represent different construction time vs. solution space trade-offs. Observe that unlike standard Xor filters, Ribbon does not exhibit sharp threshold behavior in construction

success; almost sure construction success with Standard Ribbon requires significantly more space overhead than 5% failure chance, a good space-time trade-off in our judgment.

Although Standard Ribbon does not scale infinitely for fixed ribbon width and space overhead, it is more space-efficient than standard Xor (23% overhead) for most practical n, which can be seen in Table 1 and Section 7. Unlike many other Gaussian structures, construction speed is not a significant concern for scaling Ribbon to large n.

Scaling with sharding. There are many standard or obvious ways to construct a large, space-efficient filter from many smaller space-efficient filters [37, 54]. Two ways of leveraging Ribbon features are notable, but not evaluated in detail:

- If uniformly hash-partitioning keys into fixed-size data structure shards, the *fractional r* feature of Ribbon (Section 5.1) can be used to accommodate variance in the number of keys mapped to each shard.
- If determining hash ranges to assign to each shard, on-the-fly banding allows adding entries or buckets (in sharding hash order, independent of start location order) until one fails and starts the next shard. The "⟨add till failure⟩" rows in Table 1 correspond to this strategy of adding entries until one fails, so yields good average space efficiency without construction retries, such as < 1% overhead with $m = 2^{10}$ per shard, not including sharding metadata.

Shard sizes. Section 3.1 describes an inherent unlikelihood of filling all slots in a Ribbon system, even if $m = w$, and how the likelihood is similar with $m = \Theta(w^2)$. Because the expected number of empty slots at first failure to add remains constant even for small m, the median proportion of unoccupied slots at failure decreases with m before increasing with m, for a fixed ribbon width w. For common ribbon widths, the minimum appears to be around $w^2/4$, which we suggest is the *natural* shard size, subject to practical adjustment for the application.

Soft sharding. With Ribbon we can apply sharding at a higher abstraction layer than memory space, for potentially better space efficiency. In a typical *hard sharding*, construction optimizes for each shard either (a) a set of keys, (b) a memory size, (c) a hash seed, or (d) some other configuration parameters, based on the others, and records the optimized configuration in some metadata. The change with *soft sharding* is that each shard is assigned a contiguous range of Ribbon start locations (from a single Ribbon system) rather than a contiguous memory space (containing an independent Ribbon system). This should be a pure win for space efficiency, because the overlap of $w - 1$ probing rows between adjacent soft shards allows them to, in effect, borrow some space from each other without expending metadata. (A Standard or Homogeneous Ribbon filter is a naive soft sharding with no metadata guiding the shard assignments.) With some ordering constraints and temporary tracking data, the Ribbon algorithm allows us to backtrack, such as for changing the hash seed within a shard or key assignments to shards. We do not analyze the "soft sharding" design space in detail, but use the idea for Balanced Ribbon.

---

[20]In at least two cases [1, 23], implementations citing an asymptotic result for efficient hashing in Bloom filters [40] had practical flaws that previous work [28] warned about.

Table 1: Standard Ribbon space overhead ε from empirical data

Ribbon width	Failure probability	$w/2$ smash $m = 2^{10}$	$w/4$ smash $m = 2^{10}$	0 smash $m = 2^{10}$	$m = 2^{14}$	$m = 2^{17}$	$m = 2^{24}$	Each additional doubling of m
$w = 128$	0.5	0.2%	0.1%	1.0%	1.1%	2.2%	4.7%	+0.38%
$w = 128$	⟨add till failure⟩	0.2%	0.2%	1.1%	1.2%	2.3%	4.8%	+0.38%
$w = 128$	0.05	0.5%	0.5%	2.2%	2.6%	3.7%	5.9%	+0.38%
$w = 128$	0.001	1.1%	1.2%	4.1%	4.6%	5.8%	> 7%	+0.38%
$w = 64$	0.5	0.3%	0.4%	2.0%	3.8%	6.3%	11.7%	+0.83%
$w = 64$	⟨add till failure⟩	0.8%	0.8%	2.2%	4.1%	6.5%	12.1%	+0.83%
$w = 64$	0.05	3.7%	2.9%	4.8%	7.0%	9.4%	15.0%	+0.83%
$w = 64$	0.001	11.4%	7.1%	9.2%	11.5%	13.8%	> 19%	+0.83%
$w = 32$	⟨add till failure⟩	5.9%	5.2%	6.3%	13.2%	19.2%	35.2%	+2%
$w = 16$	⟨add till failure⟩	28.9%	27.0%	27.5%	67.3%	> 100%	≫ 100%	+??%

6 BALANCED RIBBON

Balanced Ribbon is an experimental design for scaling and space-optimizing Standard Ribbon; for implementation details see [26]. We intend Balanced Ribbon as an example in the design space opened up by the new on-the-fly and incremental Gaussian elimination algorithm, and encourage follow-up work to explore, optimize, and analyze this design space.

Balanced Ribbon extends Standard Ribbon with soft sharding and a new balanced allocation scheme tailored to this domain (related: [4, 6, 7, 17, 49, 58]). Like many other hashing schemes, we start with the idea that each entry has two possible locations in the Ribbon, given by two hashes: an earlier primary location and a later secondary location in a distinct shard. The shard with the primary location is constructed before the shard with the secondary location and accommodates the key if possible. If not, we say the key is "bumped" and must be accommodated in its secondary shard, so shards add "bumped" entries first for best chance of success. Metadata is constructed to indicate which keys are bumped[21]. The construction is greedy in that a shard tries to accommodate as many keys as possible, without considering where keys will be bumped to, and constructed shards are never revisited. Difficulty arises if using two uniform hashes, the smaller primary and the larger secondary: later shards are dominated by entries in their secondary location.

Organizing shards for bumping. To make this work we organize the soft shards into levels $1..\ell$, with level i containing exactly $\lceil 2^{\ell-i-1} \rceil$ shards, so we assume a power of two number of shards overall, $s = 2^{\ell-1}$. Unlike some multi-level hashing schemes [12, 41, 43], an entry's primary location can be on any level i, with its secondary location uniformly on level $\min(i + 1, \ell)$. Because no secondary locations are in level 1, we overload it with primary locations; level 1 shards have relative weight $1 + \alpha$ for primary locations and other level shards have weight $1 - \alpha$. See Figure 7. With average n/s keys per shard, we choose $\alpha \approx 3.5/\sqrt{n/s}$ to ensure a sufficient supply of entries even for shards with three Poisson standard deviations below the mean number of entries. (We use $\alpha = 1/8$ for $n/s \approx 1000$.) Assuming we allocate our space overhead

---

[21]There is no such disambiguation in [21] and a query would combine information from *both* locations. This significantly complicates the linear system, however.
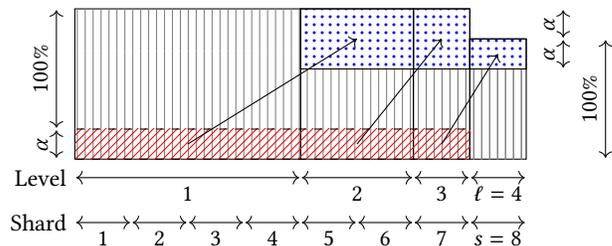
Figure 7: Bumping behavior between levels of Balanced Ribbon. The expected relative quantity of entries with primary locations in each level is indicated by ▤ where ▨ are bumped and added in their secondary location ▦. Not shown: allocation overheads (ε), Poisson variances, bucket boundaries, and dispersion within each shard.

perfectly, the average number of entries bumped from each level 1 shard will be $\alpha n/s$. Because level 2 shards are half as many, they receive $2\alpha n/s$ entries on average for adding in secondary location. With those bumped entries, level 2 shards are now overloaded to relative weight $1 + \alpha$ compared to $1 - \alpha$ for later shards. With this, we have a recursive structure to ensure a continuous supply of entries eligible for bumping down to the last shard. (Like spatial coupling [57] and "always go left" [56], we are making productive use of *less* randomness.)

The last shard (level ℓ) is different but does not need to be complicated. If we configure our allocated space overhead assuming Standard Ribbon overheads for the last shard, along with tighter overheads for the other shards, it seems to work (single shard Balanced Ribbon ≡ Standard Ribbon). For large number of shards, we observe the Balanced Ribbon final shards either completely overwhelmed with bumped entries (construction failure) or receiving almost no bumped entries. We believe this is because the overall variance in utilization of slots in all prior shards is concentrated into the last shards, and that variance is large relative to a single shard. Because the variance is small overall with a large number of

shards, construction success is more predictable at scale (apparent threshold behavior).

Buckets for bumping. We use another (semi-)independent hash to order or group keys within a shard strictly for bumping. Selecting a threshold on that hash works well for both Ribbon and metadata space efficiency. First, we have chosen our shard size such that controlling the number of entries going into a shard is much more important than the particular set of entries. Second, a $\Theta(\log(n/s))$ bit threshold value is small metadata per shard. However, we do not want to incur the CPU time for sorting entries into so many buckets per shard.

Instead of a threshold we use $\Theta(\log(n/s))$ buckets of geometrically distributed sizes that can be selected independently, using the incremental feature of Ribbon to backtrack on failed buckets. Using some bit tricks to approximate a geometric distribution with $p = 2^{-0.5}$ seems better than $p = 2^{-1}$, perhaps due to variance in actual bucket sizes. We prefer 8 buckets per shard for 8 bits of metadata per shard, 0.008 bits per key for common shard size. A subtle part of maximizing space efficiency with independent buckets and soft sharding is to attempt adding a larger (in expectation) bucket in shard $i + 1$ before attempting to add a smaller (in expectation) bucket in shard i, if the two shards are in the same level. Successfully adding the smaller could overflow enough to make it impossible to add the larger; on average, the better greedy choice is trying to add the larger bucket first. For CPU efficiency, we skip attempts to add a bucket that is very likely to fail given the number of successful additions to the shard.

Overall. Balanced Ribbon construction resembles an external sort between CPU cache and main memory. Queries depend on just one bit of sharding metadata, out of a typical 8 bits per shard. Space usage for $w = 64$ Balanced Ribbon is $1.005\lambda + 0.008$ bits per key, relative to the information-theoretic lower bound λ, as shown in Figure 5. We have tested Balanced Ribbon with these space efficiencies up to 4 billion keys; arbitrary scaling might require natural increases in w ($w \sim \log n$ and thus $n/s \sim \log^2 n$) or might benefit from design changes.

7 EXPERIMENTS

Setup. For validation we extend the experimental setup used for Cuckoo and Xor filters [35, 38, 39], with our code available in our fork on GitHub [26]. Timings are performed on a single-socket Intel® Xeon® D-2191 (Skylake DE) with 64GB RAM. Tests are compiled with GCC 8.4.1, using `g++ -O3 -DNDEBUG -march=skylake-avx512`. The Ribbon code is portable C++ using no processor intrinsics but using compiler built-ins for prefetch, count leading/trailing zero bits, and bit parity.

Results. First, Figure 1 shows which approach is fastest for various space overheads and various FP rates, and query and construction times for the corresponding overall fastest approach. It uses a mix of $n \in \{10^6, 10^7, 10^8\}$; more details in Section 1. To maximize the coverage of "BlockedBloom" in the figure, we include an AVX2 SIMD-optimized implementation from RocksDB [24] that sacrifices SIMD-optimized construction for enhanced configurability (any number of probes) and minimized FP rate compared to [44]; all

Table 2: Experimental performance comparisons

Configuration	Space ovr %	ns/key, $n = 10^6$ con	query	ns/key, $n = 10^8$ con	query
↓ FP rate around 1%, **Ribbons** using $r = 7$ ↓					
BlockedBloom[44]	52.0	11	14 ± 0	32	37 ± 0
BlockedBloom[24]	49.8	21	10 ± 0	72	36 ± 0
Cuckoo12†	46.2	69	21 ± 0	147	58 ± 0
Cuckoo12	40.3	91	20 ± 0	205	58 ± 0
Morton12	40.5	87	55 ± 3	106	125 ± 13
Xor, $r = 7$‡	23.0	195	21 ± 0	264	66 ± 1
Xor, $r = 8$	23.0	148	15 ± 0	211	50 ± 0
Xor+, $r = 8$	14.5	171	35 ± 1	299	104 ± 10
Homog., $w = 16$	52.0	56	40 ± 6	101	88 ± 6
Homog., $w = 32$	20.6	58	38 ± 6	116	85 ± 5
Standard, $w = 64$	14;20	71	42 ± 5	130	94 ± 7
Homog., $w = 64$	10.1	83	39 ± 7	160	90 ± 6
Standard, $w = 128$	6;8	166	58 ± 2	235	140 ± 26
Homog., $w = 128$	5.1	164	53 ± 3	270	145 ± 25
Balanced, $w = 32$†	15.3	84	47 ± 6	278	104 ± 5
Balanced, $w = 32$	2.5	162	48 ± 5	372	107 ± 5
Balanced, $w = 64$	0.7	292	49 ± 5	516	111 ± 5
Balanced, $w = 128$	0.3	985	66 ± 1	1335	162 ± 29
Homog.32, $r = 7$	20.6	58	38 ± 6	116	85 ± 5
Homog.32, $r = 7.7$	22.7	61	43 ± 4	120	96 ± 7
Homog.32, $r = 8$	22.1	61	39 ± 5	116	86 ± 9
↓ FP rate around 10%, **Ribbons** using $r = 3$ ↓					
Xor, $r = 3$‡	23.1	194	16 ± 0	264	55 ± 0
Homog., $w = 16$	34.6	46	23 ± 0	91	66 ± 1
Homog., $w = 32$	16.1	52	21 ± 0	108	55 ± 1
Standard, $w = 64$	14;20	66	22 ± 0	124	62 ± 0
Homog., $w = 64$	8.0	71	21 ± 0	152	72 ± 4
Standard, $w = 128$	6;8	145	35 ± 0	213	79 ± 1
Homog., $w = 128$	4.2	156	33 ± 0	266	79 ± 1
Balanced, $w = 32$	2.8	155	29 ± 0	377	68 ± 1
Balanced, $w = 64$	0.8	288	29 ± 0	517	84 ± 11
↓ FP rate around 2^{-11}, **Ribbons** using $r = 11$ ↓					
Cuckoo16†	36.0	68	20 ± 0	147	59 ± 0
Cuckoo16	30.2	84	20 ± 0	201	59 ± 1
CuckooSemiSort	26.6	146	50 ± 0	323	143 ± 4
Xor, $r = 12$	23.1	175	23 ± 0	251	74 ± 4
Xor, $r = 10.666$‡	23.0	191	20 ± 0	284	61 ± 0
Xor+, $r = 11$‡	12.8	185	38 ± 0	317	107 ± 2
Homog., $w = 32$	28.0	74	51 ± 2	129	118 ± 32
Standard, $w = 64$	14;20	78	54 ± 3	139	138 ± 41
Homog., $w = 64$	12.7	89	51 ± 3	163	118 ± 26
Standard, $w = 128$	6;9	173	71 ± 9	247	164 ± 39
Homog., $w = 128$	7.3	189	72 ± 9	292	166 ± 43
Balanced, $w = 32$	2.3	169	60 ± 2	394	142 ± 32
Balanced, $w = 64$	0.5	301	61 ± 3	536	147 ± 28

† Larger space allocated to improve construction time.
‡ Potentially unfavorable bit alignment.
; Standard Ribbon space overhead depends on n.

"BlockedBloom" in this paper use aligned 512-bit blocks. To maximize coverage of "Bloom" we include an obvious variant splitting probes between two independent blocks.

Second, Table 2 shows detailed timings for some specific configurations with approximate $f \in \{2^{-3}, 2^{-7}, 2^{-11}\}$. Construction and query times are given in nanoseconds per key over at least $5 \cdot 10^7$ keys, with a range of query times for different ratios of positive vs. negative queries. The timing data for $n = 10^6$ represents the core CPU time of each approach with negligible memory access overheads, while $n = 10^8$ includes memory access overheads. The reported space overhead is $\frac{r(1+\varepsilon)}{\log_2(f^{-1})} - 1$ where the FP rate f is measured by sampling. Standard Ribbon has distinct overheads for the different n. Cuckoo12 and Morton12 use 12 bits per cell [11, 35], for similar FP rate as $r = 9$ Xor or Ribbon.

Third, Table 1 shows many more space overheads for Standard Ribbon, which mostly improve for smaller n. While we do not include timings for these smaller n, we can infer that Standard Ribbon, with or without smash, takes territory from other Ribbon approaches in Figure 1 **(a)** for smaller n.

Observations. When saving space compared to Bloom and Cuckoo (incl. Morton), we see in Table 2,

- Ribbon can achieve much lower space overheads than the practical alternatives. $w = 64$ Balanced Ribbon has less than 1% overhead except when $r < 3$ (bottom right of Figure 1).
- Ribbon mostly wins in construction time, sometimes by a large factor. Ribbon construction times are generally only higher than Xor when saving space compared to Xor.
- Xor mostly wins in query time. For smaller rw, which is proportional to the contiguous memory loaded per query, Ribbon query times are similar to Xor. Ribbon query times increase with r and/or w.
- Ribbon has a clear advantage for configurability (aside from Bloom). Xor and Cuckoo incur a measurable penalty when cells are not aligned to a friendly size, such as 4, 8, 12, or 16 bits, and have limited options for partial-bit granularity. Ribbon performance is continuous for any integer r, and the penalty for arbitrary fractional average r is small .
- Ribbon can easily trade space efficiency for construction time efficiency (like Cuckoo), by increasing allocated space overhead; see † vs. corresponding non-† configurations. Xor construction time is more fixed.

A somewhat surprising result is that $w = 32$ Balanced Ribbon is sometimes faster *and* more space efficient than other Ribbon variants with $w = 128$. We can reason that the benefits of smaller ribbon width sometimes exceed the costs associated with balancing. We consider $w = 128$ Balanced Ribbon "impractical" because the execution time is much higher for a tiny benefit in space usage.

Limitations. A notable limitation of this test is that it tests query *throughput* more than query *latency*. Like executing a batch of filter queries, the test does not depend on the result of each query for what to do next. For applications executing a single filter query with unpredictable outcome, main memory latency can be 200-300ns. However, we expect latency to be similar between the competing approaches, because aside from Balanced Ribbon sharding metadata (very small, cachable) and Xor+ compression metadata (not

as small), only constant size metadata is read to determine which memory to fetch to complete a query. Although the test relies on out-of-order execution between queries for memory prefetching, applications can use explicit prefetching.[22]

To minimize sampling noise, the test machine was otherwise idle. This does not match a production environment, but we have not seen a significant difference in relative results when running the tests under load.

To match existing test code, hash seeds are not configurable. All structures are configured for high chance of construction success (roughly 99% or more). In practice, approaches other than BlockedBloom and Homogeneous would incur some small additional overheads associated with seeding and retries.

8 CONCLUSION

Our result changes the narrative around data structures constructed with Gaussian elimination vs. with peeling. With our new algorithm and the query structure from [22], Gaussian elimination can be faster than peeling while also opening up better space efficiency.

On this foundation we have built Ribbon filters with the following efficiently scaling variants.

- The Homogeneous Ribbon filter is simple and has a construction algorithm that never fails. We provide a full analysis.
- The Balanced Ribbon filter leverages the on-the-fly and incremental construction algorithm in an experimental load-balancing scheme that further reduces space overhead.

As with Bloom filters, the true practicality of Ribbon also comes from being able to configure it for or adapt it to the application, including dynamic conditions. This includes using any amount of space to efficiently represent any number of keys, even by shrinking the structure after construction. Ribbon filters replicate the smooth FP-rate-for-space configurability of Bloom filters and extend that with configurability between space efficiency and time efficiency.

Future work could pursue the following goals.

- Deepen the theoretical understanding of Ribbon filters to better guide parameter choices in practice.
- Manage the variance of FP rate of Homogeneous Ribbon filters at small scales.
- Further explore the design space surrounding Balanced Ribbon for approaches that are better and/or easier to analyze.
- Create a SIMD-optimized implementation of Ribbon queries, perhaps using AVX512 POPCNT and/or 8-bit ICML rather than w-bit ICML.
- Explore how competitive Ribbon can be as a data structure for static functions, including cases where the number r of solution columns is quite high, e.g. $r \geq 32$.

ACKNOWLEDGMENTS

---

[22]RocksDB MultiGet uses batched filter queries with explicit prefetching; RocksDB Get uses single filter queries.

REFERENCES

[1] Paulo Sérgio Almeida. 2020. A Case for Partitioned Bloom Filters. arXiv:2009.11789 [cs.DS]

[2] Raja Appuswamy, Goetz Graefe, Renata Borovica-Gajic, and Anastasia Ailamaki. 2019. The five-minute rule 30 years later and its impact on the storage hierarchy. *Commun. ACM* 62, 11 (2019), 114–120. https://doi.org/10.1145/3318163

[3] Martin Aumüller, Martin Dietzfelbinger, and Michael Rink. 2009. Experimental Variations of a Theoretically Good Retrieval Data Structure. In *Proc. 17th ESA*. 742–751. https://doi.org/10.1007/978-3-642-04128-0_66

[4] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. 1999. Balanced Allocations. *SIAM J. Comput.* 29, 1 (1999), 180–200. https://doi.org/10.1137/S0097539795288490

[5] Michael A. Bender, Martin Farach-Colton, Rob Johnson, Russell Kraner, Bradley C. Kuszmaul, Dzejla Medjedovic, Pablo Montes, Pradeep Shetty, Richard P. Spillane, and Erez Zadok. 2012. Don't Thrash: How to Cache Your Hash on Flash. *Proc. VLDB Endow.* 5, 11 (2012), 1627–1637. https://doi.org/10.14778/2350229.2350275

[6] Petra Berenbrink, Tom Friedetzky, Zengjian Hu, and Russell Martin. 2008. On Weighted Balls-into-Bins Games. *Theor. Comput. Sci.* 409, 3 (Dec. 2008), 511–520. https://doi.org/10.1016/j.tcs.2008.09.023

[7] Petra Berenbrink, Kamyar Khodamoradi, Thomas Sauerwald, and Alexandre Stauffer. 2013. Balls-into-Bins with Nearly Optimal Load Distribution. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Parallelism in Algorithms and Architectures* (Montréal, Québec, Canada) *(SPAA '13)*. Association for Computing Machinery, New York, NY, USA, 326–335. https://doi.org/10.1145/2486159.2486191

[8] Valerio Bioglio, Marco Grangetto, Rossano Gaeta, and Matteo Sereno. 2010. On the Fly Gaussian Elimination for LT Codes. *Communications Letters, IEEE* 13 (01 2010), 953 – 955. https://doi.org/10.1109/LCOMM.2009.12.091824

[9] Burton H. Bloom. 1970. Space/Time Trade-offs in Hash Coding with Allowable Errors. *Commun. ACM* 13, 7 (1970), 422–426. https://doi.org/10.1145/362686.362692

[10] Fabiano Cupertino Botelho, Rasmus Pagh, and Nivio Ziviani. 2013. Practical Perfect Hashing in Nearly Optimal Space. *Inf. Syst.* 38, 1 (2013), 108–131. https://doi.org/10.1016/j.is.2012.06.002

[11] Alex D. Breslow and Nuwan Jayasena. 2020. Morton filters: fast, compressed sparse cuckoo filters. *VLDB J.* 29, 2-3 (2020), 731–754. https://doi.org/10.1007/s00778-019-00561-0

[12] Andrei Z. Broder and Anna R. Karlin. 1990. Multilevel Adaptive Hashing. In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms, 22-24 January 1990, San Francisco, California, USA*, David S. Johnson (Ed.). SIAM, 43–53. http://dl.acm.org/citation.cfm?id=320176.320181

[13] Andrei Z. Broder and Michael Mitzenmacher. 2003. Network Applications of Bloom Filters: A Survey. *Internet Mathematics* 1, 4 (2003), 485–509. https://doi.org/10.1080/15427951.2004.10129096

[14] Bernard Chazelle, Joe Kilian, Ronitt Rubinfeld, and Ayellet Tal. 2004. The Bloomier Filter: An Efficient Data Structure for Static Support Lookup Tables. In *Proc. 15th SODA*. 30–39. http://dl.acm.org/citation.cfm?id=982792.982797

[15] John G. Cleary. 1984. Compact Hash Tables Using Bidirectional Linear Probing. *IEEE Trans. Computers* 33, 9 (1984), 828–834. https://doi.org/10.1109/TC.1984.1676499

[16] Colin Cooper. 2000. On the rank of random matrices. *Random Structures & Algorithms* 16, 2 (2000), 209–232. https://doi.org/10.1002/(SICI)1098-2418(200003)16:2<209::AID-RSA6>3.0.CO;2-1

[17] Artur Czumaj, Chris Riley, and Christian Scheideler. 2003. Perfectly Balanced Allocation. In *Proc. 6th APPROX - 7th RANDOM*. 240–251. https://doi.org/10.1007/978-3-540-45198-3_21

[18] Niv Dayan, Manos Athanassoulis, and Stratos Idreos. 2017. Monkey: Optimal Navigable Key-Value Store. In *Proceedings of the 2017 ACM International Conference on Management of Data* (Chicago, Illinois, USA) *(SIGMOD '17)*. Association for Computing Machinery, New York, NY, USA, 79–94. https://doi.org/10.1145/3035918.3064054

[19] Martin Dietzfelbinger, Andreas Goerdt, Michael Mitzenmacher, Andrea Montanari, Rasmus Pagh, and Michael Rink. 2010. Tight Thresholds for Cuckoo Hashing via XORSAT. In *Proc. 37th ICALP (1)*. 213–225. https://doi.org/10.1007/978-3-642-14165-2_19

[20] Martin Dietzfelbinger and Rasmus Pagh. 2008. Succinct Data Structures for Retrieval and Approximate Membership (Extended Abstract). In *Proc. 35th ICALP (1)*. 385–396. https://doi.org/10.1007/978-3-540-70575-8_32

[21] Martin Dietzfelbinger and Stefan Walzer. 2019. Constant-Time Retrieval with $O(\log m)$ Extra Bits. In *Proc. 36th STACS*. 24:1–24:16. https://doi.org/10.4230/LIPIcs.STACS.2019.24

[22] Martin Dietzfelbinger and Stefan Walzer. 2019. Efficient Gauss Elimination for Near-Quadratic Matrices with One Short Random Block per Row, with Applications. In *Proc. 27th ESA*. 39:1–39:18. https://doi.org/10.4230/LIPIcs.ESA.2019.39

[23] Peter C. Dillinger. 2018. *RocksDB Issue #4120*. https://github.com/facebook/rocksdb/issues/4120

[24] Peter C. Dillinger. 2019. *RocksDB FastLocalBloom*. https://github.com/facebook/rocksdb/blob/master/util/bloom_impl.h

[25] Peter C. Dillinger. 2020. ribbon_*.{cc,h} *in RocksDB*. https://github.com/facebook/rocksdb/

[26] Peter C. Dillinger. 2021. *Fork of fastfilter_cpp*. https://github.com/pdillinger/fastfilter_cpp

[27] Peter C. Dillinger et al. 2021. *Bloom filter on RocksDB Wiki*. https://github.com/facebook/rocksdb/wiki/RocksDB-Bloom-Filter

[28] Peter C. Dillinger and Panagiotis Manolios. 2004. Bloom Filters in Probabilistic Verification. In *Formal Methods in Computer-Aided Design, 5th International Conference, FMCAD 2004, Austin, Texas, USA, November 15-17, 2004, Proceedings (Lecture Notes in Computer Science)*, Alan J. Hu and Andrew K. Martin (Eds.), Vol. 3312. Springer, 367–381. https://doi.org/10.1007/978-3-540-30494-4_26

[29] Peter C. Dillinger and Panagiotis Manolios. 2004. Fast and Accurate Bitstate Verification for SPIN. In *Model Checking Software, 11th International SPIN Workshop, Barcelona, Spain, April 1-3, 2004, Proceedings (Lecture Notes in Computer Science)*, Susanne Graf and Laurent Mounier (Eds.), Vol. 2989. Springer, 57–75. https://doi.org/10.1007/978-3-540-24732-6_5

[30] Peter C. Dillinger and Panagiotis Manolios. 2009. Fast, All-Purpose State Storage. In *Model Checking Software, 16th International SPIN Workshop, Grenoble, France, June 26-28, 2009. Proceedings (Lecture Notes in Computer Science)*, Corina S. Pasareanu (Ed.), Vol. 5578. Springer, 12–31. https://doi.org/10.1007/978-3-642-02652-2_6

[31] Siying Dong, Andrew Kryczka, Yanqin Jin, and Michael Stumm. 2021. Evolution of Development Priorities in Key-value Stores Serving Large-scale Applications: The RocksDB Experience. In *19th USENIX Conference on File and Storage Technologies (FAST 21)*. USENIX Association. https://www.usenix.org/conference/fast21/presentation/dong

[32] David Eppstein. 2016. Cuckoo Filter: Simplification and Analysis. In *Proc. 15th SWAT*. 8:1–8:12. https://doi.org/10.4230/LIPIcs.SWAT.2016.8

[33] Jason Evans et al. 2019. *jemalloc manual*. http://jemalloc.net/jemalloc.3.html Version 5.2.1.

[34] Bin Fan, David G. Andersen, and Michael Kaminsky. 2013. Cuckoo Filter: Better Than Bloom. *;login:* 38, 4 (2013), 5. https://www.usenix.org/publications/login/august-2013-volume-38-number-4/cuckoo-filter-better-bloom

[35] Bin Fan, David G. Andersen, Michael Kaminsky, and Michael Mitzenmacher. 2014. Cuckoo Filter: Practically Better Than Bloom. In *Proc. 10th CoNEXT*. 75–88. https://doi.org/10.1145/2674005.2674994

[36] Michael L. Fredman and János Komlós. 1984. On the Size of Separating Systems and Families of Perfect Hash Functions. *SIAM Journal on Algebraic Discrete Methods* 5, 1 (1984), 61–68. https://doi.org/10.1137/0605009

[37] Marco Genuzio, Giuseppe Ottaviano, and Sebastiano Vigna. 2016. Fast Scalable Construction of (Minimal Perfect Hash) Functions. In *Proc. 15th SEA*. 339–352. https://doi.org/10.1007/978-3-319-38851-9_23

[38] Thomas Mueller Graf and Daniel Lemire. 2019. *fastfilter_cpp*. https://github.com/FastFilter/fastfilter_cpp

[39] Thomas Mueller Graf and Daniel Lemire. 2020. Xor Filters: Faster and Smaller Than Bloom and Cuckoo Filters. *ACM J. Exp. Algorithmics* 25 (2020), 1–16. https://doi.org/10.1145/3376122

[40] Adam Kirsch and Michael Mitzenmacher. 2008. Less hashing, same performance: Building a better Bloom filter. *Random Struct. Algorithms* 33, 2 (2008), 187–218.

[41] A. Kirsch and M. Mitzenmacher. 2010. The Power of One Move: Hashing Schemes for Hardware. *IEEE/ACM Trans. Netw.* 18 (2010), 1752–1765.

[42] Donald Ervin Knuth. 1998. *Sorting and Searching* (2nd ed.). Addison-Wesley. http://dl.acm.org/citation.cfm?id=280635

[43] Sailesh Kumar, Jonathan S. Turner, and Patrick Crowley. 2008. Peacock Hashing: Deterministic and Updatable Hashing for High Performance Networking. In *INFOCOM 2008. 27th IEEE International Conference on Computer Communications, Joint Conference of the IEEE Computer and Communications Societies, 13-18 April 2008, Phoenix, AZ, USA*. IEEE, 101–105. https://doi.org/10.1109/INFOCOM.2008.29

[44] Harald Lang, Thomas Neumann, Alfons Kemper, and Peter A. Boncz. 2019. Performance-Optimal Filtering: Bloom overtakes Cuckoo at High-Throughput. *Proc. VLDB Endow.* 12, 5 (2019), 502–515. https://doi.org/10.14778/3303753.3303757

[45] Daniel Lemire. 2019. Fast Random Integer Generation in an Interval. *ACM Trans. Model. Comput. Simul.* 29, 1 (2019), 3:1–3:12. https://doi.org/10.1145/3230636

[46] Yongkun Li, Chengjin Tian, Fan Guo, Cheng Li, and Yinlong Xu. 2019. ElasticBF: Elastic Bloom Filter with Hotness Awareness for Boosting Read Performance in Large Key-Value Stores. In *28th USENIX*, Dahlia Malkhi and Dan Tsafrir (Eds.). USENIX Association, 739–752. https://www.usenix.org/conference/atc19/presentation/li-yongkun

[47] Siqiang Luo, Subarna Chatterjee, Rafael Ketsetsidis, Niv Dayan, Wilson Qin, and Stratos Idreos. 2020. Rosetta: A Robust Space-Time Optimized Range Filter for Key-Value Stores. In *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*, David Maier, Rachel Pottinger, AnHai Doan, Wang-Chiew Tan, Abdussalam Alawini, and Hung Q. Ngo (Eds.). ACM, 2071–2086. https://doi.org/10.1145/3318464.3389731

[48] Yoshinori Matsunobu, Siying Dong, and Herman Lee. 2020. MyRocks: LSM-Tree Database Storage Engine Serving Facebook's Social Graph. *Proc. VLDB Endow.* 13, 12 (Aug. 2020), 3217–3230. https://doi.org/10.14778/3415478.3415546

[49] Michael Mitzenmacher. 2001. The Power of Two Choices in Randomized Load Balancing. *IEEE Trans. Parallel Distributed Syst.* 12, 10 (2001), 1094–1104. https://doi.org/10.1109/71.963420

[50] Patrick E. O'Neil, Edward Cheng, Dieter Gawlick, and Elizabeth J. O'Neil. 1996. The Log-Structured Merge-Tree (LSM-Tree). *Acta Inf.* 33, 4 (1996), 351–385. http://dblp.uni-trier.de/db/journals/acta/acta33.html#ONeilCGO96

[51] Anna Pagh, Rasmus Pagh, and S. Srinivasa Rao. 2005. An Optimal Bloom Filter Replacement. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (Vancouver, British Columbia) *(SODA '05)*. Society for Industrial and Applied Mathematics, USA, 823–829.

[52] Boris Pittel and Gregory B. Sorkin. 2016. The Satisfiability Threshold for k-XORSAT. *Combinatorics, Probability & Computing* 25, 2 (2016), 236–268. https://doi.org/10.1017/S0963548315000097

[53] Ely Porat. 2009. An Optimal Bloom Filter Replacement Based on Matrix Solving. In *Proc. 4th CSR*. 263–273. https://doi.org/10.1007/978-3-642-03351-3_25

[54] Felix Putze, Peter Sanders, and Johannes Singler. 2009. Cache-, hash-, and space-efficient Bloom filters. *ACM Journal of Experimental Algorithmics* 14 (2009), 18. https://doi.org/10.1145/1498698.1594230

[55] Kapil Vaidya, Eric Knorr, Tim Kraska, and Michael Mitzenmacher. 2020. Partitioned Learned Bloom Filter. *CoRR* abs/2006.03176 (2020), 13. arXiv:2006.03176 https://arxiv.org/abs/2006.03176

[56] Berthold Vöcking. 2003. How Asymmetry Helps Load Balancing. *J. ACM* 50, 4 (2003), 568–589. https://doi.org/10.1145/792538.792546

[57] Stefan Walzer. 2021. Peeling Close to the Orientability Threshold: Spatial Coupling in Hashing-Based Data Structures. In *Proc. 32nd SODA*.

[58] Udi Wieder. 2017. Hashing, Load Balancing and Multiple Choice. *Foundations and Trends® in Theoretical Computer Science* 12, 3–4 (2017), 275–379. https://doi.org/10.1561/0400000070

[59] Huanchen Zhang, Hyeontaek Lim, Viktor Leis, David G. Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. 2018. SuRF: Practical Range Query Filtering with Fast Succinct Tries. In *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, Gautam Das, Christopher M. Jermaine, and Philip A. Bernstein (Eds.). ACM, 323–336. https://doi.org/10.1145/3183713.3196931