

THE UNIVERSITY OF CHICAGO

OPTIMIZING COLLECTIONS OF BLOOM FILTERS WITHIN A SPACE BUDGET

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
MASTER OF SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

BY
GABRIEL MERSY

CHICAGO, ILLINOIS

APRIL 2024

Copyright © 2024 by Gabriel Mersy

All Rights Reserved

TABLE OF CONTENTS

LIST OF FIGURES	v
LIST OF TABLES	vi
ACKNOWLEDGMENTS	vii
ABSTRACT	viii
1 INTRODUCTION	1
2 BACKGROUND AND PROBLEM STATEMENT	4
2.1 Bloom Filters	4
2.1.1 Filter Collections	5
2.1.2 Utility-Weighted False Positive Rate	5
2.2 Problem Statement	5
3 TRUNCATED BLOOM FILTERS	7
3.1 Key Differences	7
3.2 False Positive Rate Analysis	9
4 OPTIMIZING FILTER COLLECTIONS	12
4.1 An Exact Optimization Problem	12
4.1.1 Formulation	12
4.1.2 Convexity	13
4.1.3 Complexity Analysis	16
4.2 A Faster Relaxation	16
4.2.1 The Lower Bound	17
4.2.2 Convexity	17
4.2.3 Complexity Analysis	18
4.3 Further Relaxations	18
5 EVALUATION	19
5.1 Implementation Details	19
5.2 Data Skipping	20
5.2.1 Experimental Setup	21
5.2.2 Results	23
5.3 Full-Text Search	26
5.3.1 Experimental Setup	28
5.3.2 Results	30
5.4 Microbenchmarks	32
5.4.1 Truncated Bloom Filters	32
5.4.2 Optimization Latency	34

5.4.3	Alternative Algorithms	36
5.4.4	Different Utility Distributions	37
6	RELATED WORK	39
6.1	Filters	39
6.2	Bloom Filters in Key-value Stores	39
6.3	Modular Bloom Filters	39
6.4	Other Bloom Filter Applications	40
7	CONCLUSION	41
	REFERENCES	42

LIST OF FIGURES

3.1	An overview of our method.	8
5.1	Average skip rate on RE, EV, and NASA at different size budgets (higher is better). We generally outperform the in-memory methods and are similar to the on-disk filter.	24
5.2	Average wasted time on RE, EV, and NASA at different index size budgets B (lower is better). We outperform the in-memory methods.	25
5.3	Average query latency on RE, EV, and NASA at different index size budgets B (lower is better). We are generally the lowest latency choice across all methods and budgets.	26
5.4	Increasing compute-storage disaggregation can greatly magnify the cost of false positives.	27
5.5	I&S (left) and MI (right). We have higher precision than the best-performing baseline (TMII) at larger budgets and rival its performance at smaller budgets.	31
5.6	I&S (left) and MI (right). We are slower than the baselines besides scan but are impeded by implementation.	32
5.7	FPR (left) and query latency (right) as a function of the truncation ratio p . Latency is given as 95% confidence intervals, + and - are positive and negative keys respectively.	33
5.8	Solving the relaxation scales to large inputs. The resulting throughput is approximately 20 thousand filters per second on a resource-constrained machine.	35
5.9	Our method harmonizes the benefits of Top Utility and Equal Truncation.	36
5.10	Skewed distributions imply better performance.	37

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank my advisor Professor Sanjay Krishnan for his ongoing creativity and guidance throughout the research process, my collaborator Professor Stavros Sintos who helped make this work possible, and the ChiData group. I also thank Professor Aaron Elmore and Professor Raul Castro Fernandez for joining my committee. Last but not least, I extend my heartfelt gratitude to Izabella and the rest of my family for their support.

ABSTRACT

With a single Bloom filter, one can approximately answer set membership queries within a space budget. Practical systems often use collections of Bloom filters to facilitate applications such as data skipping, sideways information passing, and network filtering. While the optimal space-to-accuracy allocation is well-understood for a single filter, jointly optimizing how space is used across a collection of filters is yet to be studied. We pose this problem in the following way: (1) let's assume that each Bloom filter has some likelihood of being queried, and (2) given knowledge of this likelihood, how do we allocate space to minimize the expected false positive rate? In other words, “hot” filters are allocated more space, and “cold” filters are allocated less space. In this paper, we show how to solve this optimization problem. We first develop the concept of a “truncated” Bloom filter and theoretically analyze its false positive rate. We then formulate an optimization problem for a collection of truncated Bloom filters that minimizes the false positive rate across a utility distribution while meeting a strict space budget. Next, we show that the problem is convex and find a fast relaxation. Lastly, we apply our method to data skipping and full-text search, demonstrating its effectiveness across the range of possible space budgets when compared to the state of the art.

CHAPTER 1

INTRODUCTION

A Bloom filter is a sketch that can approximately answer set membership queries within a space budget. Bloom filters have a crucial property that does not yield false negative results, which makes them highly valuable in database systems. The applications of Bloom filters span key-value stores Dayan et al. [2017], Redis [2024], data skipping Databricks [2023], Dageville et al. [2016], query optimization Gaffney et al. [2022], Zhu et al. [2017], and search engines Goodwin et al. [2017]. Filters are often employed in these applications to reduce storage reads by serving as an intermediary index structure that can short-circuit data loading.

Let's consider a typical use case for a Bloom filter. A data volume might be partitioned across several different files. A user who is interested in identifying a particular record that matches an equality condition (e.g., an employee identified by a social security number) wants to query this data. On one hand, we can simply scan the entire data volume each time we want to answer this query. This approach requires no indexing but is slow. Another approach is to build an index over the search keys, which can require substantial additional space and may be hard to operationalize in some environments (e.g., a write-heavy database or a distributed system). A Bloom filter allows one to span these two extremes by annotating each partition with a space-limited filter that determines if it needs to be processed or not – the larger the filter, the more precise the skipping will be. This filter is straightforward to maintain under insertions.

While the design space of a single Bloom filter is well-understood Mitzenmacher [2002], use cases like the one above show that we often consider collections of such filters (i.e., filters on each partition). The standard approach is to uniformly size each filter in a filter collection, which has a few key drawbacks. First, while a single Bloom filter has a fixed size regardless of the number of insertions, the total size of a collection of filters will grow as more

filters are utilized (e.g., over new partitions of data). If we wish to cache the filter collection in memory, the constant growth can be problematic. Next, not all partitions are equally likely to be queried, and uniform sizing allocates the same precision to frequently queried partitions as to rare ones. This paper studies the following problem: *given an anticipated query distribution and an overall space budget, can we optimally size each filter to maximize the expected precision?*

Let's contextualize this problem statement in the example above, and make it much more concrete. Consider a telemetry system, such as AWS CloudWatch, which can generate a vast amount of unstructured log data. In systems like CloudWatch, log data is broken up into temporal partitions of roughly the same size. Over this log data, users can run keyword search queries. Bloom filters constructed on each partition can help skip partitions that could not possibly contain the desired keyword. While in general, one needs to process all partitions, we know that more recent data is more likely to be relevant to a user (the anticipated query distribution). Often, users set a time interval restricting how far back the search should go. The solution to the optimization problem above can construct an in-memory filter collection with a fixed size that should automatically reduce the precision of filters for the older partitions.

Going beyond this example, there is a broader data retrieval architecture that this algorithmic problem enables. Data are stored in partitions on some form of slow storage, e.g., disk or a blob storage system. We wish to answer set membership queries over this data, e.g., keyword search. These partitions are indexed in memory on a query processing node with a guaranteed max size for the entire collection of Bloom filters. Since the total size of all Bloom filters is fixed via optimization, they can be predictably stored in memory. As more data are added, we show how this in-memory index can be re-optimized to maintain a fixed size (albeit with an increased error rate for some queries). The result is *a fixed-size index over growing data*. While this implies that overall filtering accuracy goes down over time,

if there is a sufficient, known skew in the query distribution (e.g., a bias towards querying more recent filters), then perhaps the query-weighted accuracy does not degrade as much.

We contribute the algorithms to make this architecture a reality. Assuming a query distribution as a proxy for *utility* is not a new idea. For example, caching heavily leverages this principle, as it persists certain data items that are accessed more frequently than others Berg et al. [2020]. In particular, we will show how to find an optimal index structure of Bloom filters that minimizes the false positive rate over a utility distribution, while adhering to a strict space budget. This can be thought of as a form of “partial” caching, where rather than evicting a low-utility filter completely from the cache, we simply reduce its precision. Interestingly enough, this is actually the optimal thing to do in the case of a Bloom filter.

The contributions of this paper are as follows.

- We define the *utility-compressed sketch problem*, where a sketch collection is compressed according to the utility values of its data objects (§2).
- We propose the *truncated Bloom filter* as a simple extension in which the filter is shortened after construction, enabling the trade-off between space and false positive rate to be traversed in a fine-grained fashion (§3).
- We formulate a constrained optimization problem for a collection of truncated Bloom filters that assigns higher false positive rates to lower-utility filters by modulating filter lengths, subject to a strict bit-level space budget (§4).
- We prove, at odds with intuition, that the resulting optimization problem is convex. We further derive a relaxation of the original problem that is also provably convex and features an objective function that runs in linear time (§4).
- We demonstrate the effectiveness of our method when compared to the state of the art on two representative Bloom filter applications: data skipping and full-text search (§5).

CHAPTER 2

BACKGROUND AND PROBLEM STATEMENT

We begin by covering some background and the problem statement.

2.1 Bloom Filters

A Bloom filter Bloom [1970] is a probabilistic data structure for representing a set that can answer approximate membership queries. For convenience, we write \mathcal{B} when we refer to a Bloom filter throughout the paper. An m -bit Bloom filter \mathcal{B} contains n elements and is associated with k independent hash functions h_1, h_2, \dots, h_k . We denote these parameters as the 3-tuple (m, k, n) . We will use the notation $|\cdot|$ to denote the size of a sketch. For a Bloom filter, this is initially given by the parameter m .

The filter is initially an array of all zeros. When adding an element b to \mathcal{B} a total of k hashed positions in the filter $\mathcal{B}[h_1(b)], \mathcal{B}[h_2(b)], \dots, \mathcal{B}[h_k(b)]$ are set to 1. For a query key q each of the k hashed positions in the filter is probed. If a position is not set to 1, then a negative ($-$) result is returned. If all of the positions are set to 1, then a positive ($+$) result is returned. If the key is in the filter, then it returns the correct answer. If the key is not in the filter then it returns the correct answer with probability $1 - \varepsilon$ for a small parameter $\varepsilon \in (0, 1)$. In other words, with probability ε the filter returns a positive result for a key that is not in the filter (i.e., a false positive).

Denote q^+ as a positive key that is in the filter and q^- as a negative key that is not in the filter. A Bloom filter by design has a false positive rate of $\mathbb{P}(+ | q^-) = \varepsilon$ and a zero false negative rate $\mathbb{P}(- | q^+) = 0$. The false positive rate ε is given as $\varepsilon \approx \left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^k$. It has been shown that this characterization is a very close approximation but is not quite correct due to a flawed bit independence assumption Bose et al. [2008]. Since the alternative is an impractical combinatorial function, we exclusively make use of the

closed-form approximation, as is standard Dayan et al. [2017].

2.1.1 Filter Collections

A filter collection is a set of N Bloom filters $\{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_N\}$ where each filter \mathcal{B}_i represents a data object \mathcal{D}_i . We further assume that there is a utility function $u : \mathcal{D} \rightarrow \mathbb{R}_{\geq 0}$ that maps each data object to a non-negative real number. Throughout the paper, we write u to refer to a utility value that has already been pre-computed by the function in question. A utility function represents some notion of usefulness with respect to a particular Bloom filter in the collection. For example, utility may represent the frequency at which a filter is queried when compared to other filters, or the probability that a filter is queried.

2.1.2 Utility-Weighted False Positive Rate

In the same way that we can calculate the false positive rate ε for a single Bloom filter, we can calculate the expected false positive rate for a filter collection using the individual false positive rates $\varepsilon_1, \dots, \varepsilon_N$ and utility values u_1, u_2, \dots, u_N . We arrive at a final figure of merit which is the utility-weighted false positive rate:

$$\mathcal{E} = \sum_{i=1}^N u_i \cdot \varepsilon_i \tag{2.1}$$

If one interprets the utility values as probabilities, this function measures the overall precision of the filter collection, e.g., how accurately the collection skips data given known partition query probabilities.

2.2 Problem Statement

We would like to produce a final index structure that allocates space to Bloom filters within a space budget B by minimizing the utility-weighted false positive rate \mathcal{E} .

Definition 1 (Problem Statement). We are given a space budget B , data objects $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_N$ and corresponding utility values u_1, u_2, \dots, u_N . The goal is to produce a filter collection such that:

$$\sum_{i=1}^N |\mathcal{B}_i| \leq B$$

by minimizing

$$\min \sum_{i=1}^N u_i \cdot \varepsilon_i$$

where $u_i = u(\mathcal{D}_i)$ and $\varepsilon_i = \phi(\mathcal{B}_i)$ with ϕ denoting the Bloom filter false positive rate formula.

This problem statement is very similar to a caching problem statement at first glance; namely, identify the data items with the highest utility. The key difference when compared to standard caching techniques is that items are not simply evicted or cached according to a space budget. We now have the additional option of degrading the accuracy of an item. It may very well be that the optimal solution in some cases is to allocate 0 bits to an item, but that should only be a consequence of the optimization algorithm.

This formulation is further an instance of a broader problem that we call a *utility-compressed sketch*. In the general case, the space for a particular data object’s sketch \mathcal{X}_i is allocated by minimizing a weighted sum across error functions of the form $\phi(|\mathcal{X}_i|)$ subject to a budget constraint $\sum_{i=1}^N |\mathcal{X}_i| \leq B$. In typical sketches, the error function ϕ is monotonic in space. We can think of a utility-compressed sketch as an optimization over lossy data structures that represent data items at several different “resolutions”, which is related to similar problems in data compression Barbarioli et al. [2023]. There are many possible definitions of $\phi(|\mathcal{X}_i|)$. In the case of a Bloom filter, $\phi(|\mathcal{X}_i|)$ is the false positive rate as a function of its size – we could apply the same reasoning to other sketches such as the count-min sketch Cormode and Muthukrishnan [2005].

CHAPTER 3

TRUNCATED BLOOM FILTERS

To solve this problem, we will introduce a new sketch called a truncated Bloom filter. It is based on the simple idea that previously-allocated space in a Bloom filter can later on be revoked. This concept is useful for a few key reasons. First, it allows us to formulate the optimization problem as a revocation of previously-allocated space; that is, we can start with a uniform allocation of space and remove bits based on the optimization objective. Second, it allows us to operate in an online setting where new data are added but we wish to keep a fixed storage size. Finally, it allows for consistent querying across the entire filter collection with a fixed set of hash functions.

The standard Bloom filter is inflexible when there are changing requirements for false positive rates. Both the target false positive rate and the expected number of elements in the filter are often specified prior to construction, which means that the Bloom filter must be reconstructed from scratch if less space is desired in the future. While there has been work on expanding filters Dayan et al. [2023], we study the opposite problem of contracting filters.

3.1 Key Differences

We begin by highlighting the key differences between a truncated Bloom filter and a standard Bloom filter.

Definition 2 (Truncated Bloom Filter). Let \mathcal{B} be a Bloom filter with m bits, k hash functions, and n added elements. A *truncated Bloom filter* is a new Bloom filter \mathcal{B}' formed by removing the rightmost $m - m'$ bits from \mathcal{B} . The truncated filter \mathcal{B}' satisfies $\mathcal{B}'[1, 2, \dots, m'] = \mathcal{B}[1, 2, \dots, m']$.




Figure 3.1: An overview of our method.

The only notational difference is that a new parameter m' is introduced to specify the truncated filter length. To operationalize this definition, we show that simply considering the subset of hash functions that fall before m' is sufficient to query the filter.

Definition 3 (Valid Hash Function). We say that a hash function $h_s : \mathbb{U} \rightarrow [m]$ taking an arbitrary object $o \in \mathbb{U}$ from universe \mathbb{U} as input is a *valid hash function* if $h_s(o) \leq m'$. We say that a hash function is an *invalid hash function* if $h_s(o) > m'$, i.e., a hash function is valid if it maps to the remaining region and is invalid if it maps to the truncated region.

These concepts are also highlighted visually in Figure 3.1. As we move to the right in the figure, the filters exhibit increasing levels of truncation (i.e., smaller values of m'). Given a negative query key q^- the number of valid hash functions (solid arrows) decreases, while the number of invalid hash functions (dotted arrows) increases.

Membership queries are issued to a truncated filter using the procedure described in Algorithm 1. The key algorithmic difference in querying a truncated Bloom filter is that there is an extra condition to check whether each hash function is valid or not.

There are also a few degenerate cases that are worth noting. First, if $m' = 0$ then the query always returns a positive result and has a 100% false positive rate. Second, if there are no valid hash functions, then the query always returns a positive result and also has a 100% false positive rate. Third, if $m' = m$ then the false positive rate is unchanged and is thus identical to the original filter.

Input: Query q , truncated Bloom filter \mathcal{B}' with $m' \leq m$ bits, k hash functions h_1, h_2, \dots, h_k where $h_s : \mathbb{U} \rightarrow [m]$

```

for  $s \in [k]$  do
  | if  $h_s(q) \leq m' \wedge \mathcal{B}'[h_s(q)] = 0$  then
  | | return -
  | end
end
return +

```

Algorithm 1: Querying a truncated Bloom filter.

3.2 False Positive Rate Analysis

Aside from the degenerate cases, the false positive rate of a truncated Bloom filter strictly falls between 100% and the original false positive rate. To build intuition concerning the false positive rate of a truncated Bloom filter, consider the first filter in Figure 3.1. Only 3 out of the 4 hash functions are valid, and the false positive rate in this example is $(1 - (1 - \frac{1}{m_1})^{4n_1})^3$ which is strictly higher than the false positive rate $(1 - (1 - \frac{1}{m_1})^{4n_1})^4$ of the original filter. In fact, the false positive rate is governed by the number of valid hash functions. Before a particular query, we do not know precisely how many hash functions will be valid. Instead, we define a random function (called the *false positive function*) $\phi(\cdot)$, that returns a random variable for the false positive rate, as shown in Proposition 1.

Definition 4 (False Positive Function). The false positive function $\phi(V)$ of a truncated Bloom filter is:

$$\phi(V) = \left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^V \quad (3.1)$$

where V is a random variable for the number of valid hash functions defined on the sample space $\Omega = \{0\} \cup [k]$.

Proposition 1 (False Positive Rate). *Given a collection of k independent and uniform hash functions h_1, \dots, h_k and the fraction of remaining bits $p = \frac{m'}{m}$ after truncation, the false*

positive rate ε of a truncated Bloom filter assuming bit independence is:

$$\varepsilon = \mathbb{E}[\phi(V)] = \sum_{v=0}^k \left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^v \cdot \binom{k}{v} \cdot p^v (1-p)^{k-v} \quad (3.2)$$

with $V \sim \text{Bin}(k, p)$ represented as a binomial random variable.

Proof. Let $\mathcal{B}[1, 2, \dots, m]$ be a Bloom filter and let $\mathcal{B}' = \mathcal{B}[1, 2, \dots, m']$ be its truncated filter. Let $H = \{h_1, \dots, h_k\}$ be the set of k hash functions used in filter \mathcal{B} . Let $\mathcal{B}'(q) \in \{-, +\}$ be the result of Algorithm 1. The false positive rate is defined as $\varepsilon = \Pr[\mathcal{B}'(q) = + \mid q \notin \mathcal{B}]$, where q is a key that is not stored in Bloom filter \mathcal{B} . From the query algorithm we have that $\mathcal{B}'(q^-) = +$ if for every $h \in H$ with $h(q) \leq m'$ it holds that $\mathcal{B}'[h(q)] = \mathcal{B}[h(q)] = 1$. Let $\mathcal{I} \subseteq H$ be a subset of the hash functions and let $q \notin \mathcal{B}$ be a key that is not stored in the filter. We define the random variables:

- $A_{\mathcal{I}}$: Given that $q \notin \mathcal{B}$, if $\bigwedge_{h \in \mathcal{I}} \mathcal{B}[h(q) = 1]$ then $A_{\mathcal{I}} = 1$, otherwise $A_{\mathcal{I}} = 0$.
- $C_{\mathcal{I}}$: if $h(q) \leq m'$ for every $h \in \mathcal{I}$ and $h(q) > m'$ for every $h \in H \setminus \mathcal{I}$ then $C_{\mathcal{I}} = 1$, otherwise $C_{\mathcal{I}} = 0$.

We have,

$$\begin{aligned} \varepsilon &= \Pr[\mathcal{B}'(q) = + \mid q \notin \mathcal{B}] = \Pr \left[\bigvee_{\mathcal{I} \subseteq H} (A_{\mathcal{I}} \wedge C_{\mathcal{I}}) \right] \\ &= \sum_{\mathcal{I} \subseteq H} \Pr[A_{\mathcal{I}} \wedge C_{\mathcal{I}}] = \sum_{v=0}^k \sum_{\mathcal{I} \subseteq H, |\mathcal{I}|=v} \Pr[A_{\mathcal{I}} \wedge C_{\mathcal{I}}] \\ &= \sum_{v=0}^k \sum_{\mathcal{I} \subseteq H, |\mathcal{I}|=v} \Pr[C_{\mathcal{I}}] \cdot \Pr[A_{\mathcal{I}} \mid C_{\mathcal{I}}] \\ &= \sum_{v=0}^k \sum_{\mathcal{I} \subseteq H, |\mathcal{I}|=v} p^v (1-p)^{k-v} \cdot \Pr \left[\bigwedge_{h \in \mathcal{I}} \mathcal{B}[h(q) = 1] \mid q \notin \mathcal{B} \right] \end{aligned}$$

$$\begin{aligned}
&= \sum_{v=0}^k \sum_{\mathcal{I} \subseteq H, |\mathcal{I}|=v} p^v (1-p)^{k-v} \cdot \prod_{h \in \mathcal{I}} \Pr[\mathcal{B}[h(q) = 1] \mid q \notin \mathcal{B}] \\
&= \sum_{v=0}^k \sum_{\mathcal{I} \subseteq H, |\mathcal{I}|=v} p^v (1-p)^{k-v} \left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^v \\
&= \sum_{v=0}^k \binom{k}{v} p^v (1-p)^{k-v} \left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^v \\
&= \mathbb{E}[\phi(V)]
\end{aligned}$$

□

CHAPTER 4

OPTIMIZING FILTER COLLECTIONS

In this section, we formalize a constrained optimization problem that produces a utility-compressed sketch of truncated Bloom filters. We first prove that the optimization problem is convex and can thus be solved using a standard algorithm for convex optimization. We then derive a relaxation that is also shown to be convex. Crucially, the relaxation can be solved faster than the original problem both in theory and in practice.

4.1 An Exact Optimization Problem

We are given a space budget B , utility values u_1, u_2, \dots, u_N and Bloom filters $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_N$ constructed with parameter tuples of the form (m_i, k_i, n_i) .

4.1.1 Formulation

We begin with a simple description of the optimization problem. The basic intuition is that a higher false positive rate is assigned to a filter that has a lower utility value, and a lower false positive rate is assigned to a filter with a higher utility value. To achieve this desired property, we write $u \cdot \mathbb{E}[\phi(V)]$. The false positive rate of a truncated Bloom filter is modulated by setting the number of remaining bits m' prior to the truncation operation. Therefore, our variables are simply the truncated filter lengths m'_1, m'_2, \dots, m'_N across our collection of Bloom filters. Referring back to Figure 3.1, we can see how the utility values affect the truncated filter lengths in our optimization routine.

Following this intuition, the objective is to minimize the dot product between the utility values and the false positive rate terms by finding an optimal assignment of truncated filter lengths, subject to a few constraints. First, the sum of the m'_i must not exceed the budget B . Second, each m'_i must fall between 0 and the original filter length m_i . Using the false

positive rate from Equation (3.2), we define our constrained optimization problem as follows:

$$\begin{aligned} \min_{m'_1, \dots, m'_N} \quad & \sum_{i=1}^N u_i \cdot \mathbb{E}[\phi(V_i)] \\ \text{s.t.} \quad & \sum_{i=1}^N m'_i \leq B \\ & 0 \leq m'_i \leq m_i \end{aligned}$$

where B is the target budget, and each variable m'_i corresponds to a truncated length. Note that each truncated filter length $m'_i \in \mathbb{R}_{\geq 0}$ is currently treated as a continuous optimization variable. Of course, a fractional truncated filter length is not possible, so we simply take the optimal solution and for each $i \in [N]$ round down to the nearest integer; $m'_i \leftarrow \lfloor m'_i \rfloor$. This ensures that the budget constraint is still satisfied. We can also substitute a budget equality constraint $\sum_{i=1}^N m'_i = B$ for the budget inequality constraint $\sum_{i=1}^N m'_i \leq B$.

4.1.2 Convexity

In order to establish convexity, we need to verify that the objective and constraint functions are all convex. By inspection, the constraint functions in the above formulation are linear and therefore convex. Hence, it suffices to establish the convexity of the objective function. In Lemma 2, we show that the false positive function $\phi(\cdot)$ as written in Equation (3.1) is convex on $\mathbb{R}_{\geq 0}$ which is a superset of the feasible region (i.e., polytope) defined by the constraints.

Lemma 2 (Convexity of ϕ). $\phi(x) = \left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^x$ is convex on $x \geq 0$ for $m, k, n > 0$.

Proof. By the definition of convexity, we want to show that the second derivative $\frac{d^2\phi(x)}{dx^2}$ is non-negative. Let $x \geq 0$ and $m, k, n > 0$. Define $w = 1 - \left(1 - \frac{1}{m}\right)^{kn}$ which is a strictly

positive value by construction. We differentiate twice to obtain:

$$\frac{d^2\phi(x)}{dx^2} = w^x \log^2(w) \geq 0.$$

□

It has now been shown that the truncated Bloom filter false positive function is convex under mild conditions. We now prove a more surprising result concerning the convexity of the objective function. Intuition derived from a plot of the binomial probability mass function may suggest that the objective function is unlikely to be convex in the optimization variable. On the contrary, we find that the function is indeed convex, as presented in Proposition 3. The structure of the proof relies the observation that the binomial distribution is an *exponential family*, which allows us to employ properties of such families to infer that our objective function is convex.

Proposition 3 (Objective Convexity). *The objective function $f(\mathbf{m}') = \sum_{i=1}^N u_i \cdot \mathbb{E}[\phi(V_i)]$ is convex in m'_i on $0 \leq m'_i \leq m_i$.*

Proof. We first note the symmetry of the objective function and reduce the problem to the univariate case. This is a valid argument because if $\frac{d^2}{dm'^2_i} \mathbb{E}[\phi(V_i)] \geq 0$ then, by linearity of differentiation and the non-negative utility values, the Hessian \mathbf{H}_f must be a diagonal matrix exclusively composed of non-negative entries (and f is convex). The core idea is to show that the base function $\mathbb{E}[\phi(V_i)]$ is convex in a single variable m'_i and then deduce that the objective function is composed of a convexity-preserving operation on multiple instances of the base function with distinct variables.

We borrow from the statistics literature the known fact that the binomial distribution is an exponential family for a fixed number of trials (the number of hash functions in our case). In particular, this means that the binomial probability mass function can be rewritten in a canonical form. Let X be an exponential family random variable $X : \Omega \rightarrow E$ with probability

density function $\mathbb{P}_X(x; \theta)$ that depends on the parameter θ . Assuming that $\phi(\cdot)$ is convex on the sample space Ω and the expectation $\mathbb{E}[X]$ is linear in the parameter θ , then it is known that $\mathbb{E}[\phi(X)]$ is convex in the parameter θ Shaked [1980], Schweder [1982].

Let $m_i, k_i, n_i > 0$ and $0 \leq m'_i \leq m_i$ for all $i \in [N]$. Each $V_i : \Omega \rightarrow E$ is a binomial random variable from an exponential family with probability mass function $\mathbb{P}_{V_i}(v_i; p_i)$. The sample space Ω of a binomial random variable is non-negative. By Lemma 2, it then follows that the function $\phi(\cdot)$ is convex on the sample space of V_i . The expectation $\mathbb{E}[V_i] = k_i p_i$ is linear in the parameter p_i Brown [1986]. Therefore, we have that the function $\mathbb{E}[\phi(V_i)]$ is convex in the success probability parameter p_i . This implies that the second derivative with respect to p_i is non-negative $\frac{d^2 \mathbb{E}[\phi(V_i)]}{dp_i^2} \geq 0$. We differentiate Equation (3.2) twice with respect to the success probability p_i and obtain:

$$\begin{aligned} \frac{d^2 \mathbb{E}[\phi(V_i)]}{dp_i^2} &= \sum_{v_i=0}^{k_i} \phi(v_i) \cdot \binom{k_i}{v_i} \left[v_i(v_i - 1) p_i^{v_i-2} (1 - p_i)^{k_i-v_i} \right. \\ &\quad \left. - 2v_i(k_i - v_i) p_i^{v_i-1} (1 - p_i)^{k_i-v_i-1} \right. \\ &\quad \left. + (k_i - v_i)(k_i - v_i - 1) p_i^{v_i} (1 - p_i)^{k_i-v_i-2} \right] \end{aligned}$$

It remains to show that $\mathbb{E}[\phi(V_i)]$ is also convex in the truncated filter length variable m'_i . Notice that $\frac{dp_i}{dm'_i} = \frac{1}{m_i} > 0$. We now differentiate $\mathbb{E}[\phi(V_i)]$ twice with respect to m'_i and find that the following holds:

$$\frac{d^2 \mathbb{E}[\phi(V_i)]}{dm_i'^2} = \frac{1}{m_i^2} \cdot \frac{d^2 \mathbb{E}[\phi(V_i)]}{dp_i^2} \geq 0$$

Specifically, the second derivative with respect to our desired variable m'_i can be rewritten as a positive constant multiplied by the second derivative with respect to p_i . Hence, $\mathbb{E}[\phi(V_i)]$ is convex in the variable m'_i since $\frac{d^2 \mathbb{E}[\phi(V_i)]}{dm_i'^2} \geq 0$. Lastly, a weighted sum with non-negative coefficients $u_i \geq 0$ is known to preserve convexity, so the objective function f is convex in

m'_i on the interval $0 \leq m'_i \leq m_i$ which concludes the proof. □

4.1.3 Complexity Analysis

The time complexity of computing the objective is quadratic $\mathcal{O}(N \max_{i \in [N]} k_i)$. To solve such convex optimization problems, interior-point methods convert constrained problems to unconstrained problems via the introduction of barrier functions Boyd and Vandenberghe [2004]. Additionally, these solvers often rely on Newton’s method as a crucial subroutine. When objective functions are computationally expensive, the evaluation of gradient vectors and Hessian matrices at high-dimensional points in space can become a critical bottleneck.

In our case, linearity of differentiation suggests that the run time of the underlying solver will be highly dependent on the time complexity of the objective function. For example, each entry of the gradient would contain a summation that mirrors the inner summation over k_i in the objective function, combined with the terms for the barrier functions. While a quadratic time objective is likely acceptable for many applications, in the next section we explore how to find a close approximation of the optimal solution with an objective function that can be computed in only linear time $\mathcal{O}(N)$. In doing so, we remove the dependency on the number of hash functions and eliminate the $\mathcal{O}(k_i)$ factor.

4.2 A Faster Relaxation

In this section, we describe a relaxation of the optimization problem that features an approximate objective function. In particular, the objective function is a lower bound of the original objective function from the previous section. When minimizing a lower bound, we are able to simplify and speed up the optimization problem by paying the price of an exact optimal solution. Such lower bound methods are often applied when the original problem is intractable. We find that, while our original problem is certainly tractable, minimizing a lower bound is an elegant solution that benefits both theory and practice. We will also

demonstrate these observations empirically in the experiments.

4.2.1 The Lower Bound

We now derive a lower bound for the objective function. Since Lemma 2 states that ϕ is a convex function, we invoke Jensen's inequality $\mathbb{E}[\phi(V)] \geq \phi(\mathbb{E}[V])$ and write the following lower bound for the false positive rate term using the fact that $\mathbb{E}[V] = kp$:

$$\phi(\mathbb{E}[V]) = \left(1 - \left(1 - \frac{1}{m}\right)^{kn}\right)^{kp} \quad (4.1)$$

Hence, we can reformulate the objective function to yield a relaxation by substituting the lower bound in Equation (4.1) for the original false positive rate term from Equation (3.2). We can verify that the new problem is a relaxation of original problem, since $\sum_{i=1}^N u_i \cdot \phi(\mathbb{E}[V_i]) \leq \sum_{i=1}^N u_i \cdot \mathbb{E}[\phi(V_i)]$ and the feasible region remains the same.

4.2.2 Convexity

In a similar fashion to the previous section, we now confirm that the approximate objective function is convex in our optimization variable. This brings us to Proposition 4.

Proposition 4 (Relaxation Convexity). *The objective function $g(\mathbf{m}') = \sum_{i=1}^N u_i \cdot \phi(\mathbb{E}[V_i])$ is convex in m'_i on $0 \leq m'_i \leq m_i$.*

Proof. We show that the desired function can be constructed from a base function $\phi(\cdot)$ that is known to be convex and a convexity-preserving operation. Let $m_i, k_i, n_i > 0$ and $0 \leq m'_i \leq m_i$ for all $i \in [N]$. Define $x_i = k_i p_i$ and compute $\frac{dx_i}{dm'_i} = \frac{k_i}{m_i} > 0$. By Lemma 2, the base function $\phi(x_i)$ is convex in x_i since $x_i \geq 0$. This means that its second derivative is non-negative. Then, we have:

$$\frac{d^2\phi(x_i)}{dm_i'^2} = \frac{k_i^2}{m_i^2} \cdot \frac{d^2\phi(x_i)}{dx_i^2} \geq 0$$

Therefore, the function $\phi(x_i)$ is convex in m'_i since the second derivative with respect to m'_i is non-negative. Finally, the weighted sum of the base function with non-negative coefficients $u_i \geq 0$ preserves convexity, so we conclude that g is convex in m'_i over the interval $0 \leq m'_i \leq m_i$. \square

4.2.3 Complexity Analysis

The key difference in complexity when comparing the lower bound to the original objective function is that the $\mathcal{O}(k_i)$ factor in the running time is replaced with $\mathcal{O}(1)$ time to compute the binomial expected value $\mathbb{E}[V_i] = k_i p_i$. The overall time complexity for computing the objective function is reduced to linear $\mathcal{O}(N)$ with respect to the number of Bloom filters. The run time of many interior-point solvers that rely on Newton's method would also be reduced, as was discussed previously.

4.3 Further Relaxations

We can further simplify the lower bound in Equation (4.1) observing that $1 - \left(1 - \frac{1}{m}\right)^{kn} \approx 1 - e^{-kn/m}$ and using the optimal value of k in the Bloom filter, $k = \frac{m}{n} \ln 2$. The objective function becomes $\min_{\mathbf{m}'} \sum_{i=1}^N u_i \cdot 0.618^{m'_i/n_i}$, which is also a convex function. We can further simplify the function observing that usually $m'_i \leq 4n_i$, so the function $0.618^{m'_i/n_i}$ can be approximated with the line $-0.213 \frac{m'_i}{n_i} + 1$. Hence, the objective function becomes $\min_{\mathbf{m}'} \sum_{i=1}^N -0.213 \cdot \frac{u_i}{n_i} m'_i$. Equivalently, the objective can be written as $\max_{\mathbf{m}'} \sum_{i=1}^N \frac{u_i}{n_i} m'_i$. Notice that it is trivial to optimize this function while satisfying the size constraints of the truncated filters. We sort and traverse all Bloom filters in descending order of $\frac{u_i}{n_i}$. For each Bloom filter \mathcal{B}_i , we assign the maximum value of $m'_i \leq m_i$ that does not violate the total size constraint. When the total size becomes B , then for all the remaining filters we set $m'_i = 0$. The overall running time of the algorithm is $\mathcal{O}(N \log N)$.

CHAPTER 5

EVALUATION

In this section, we evaluate our method on two applications and several microbenchmarks. We begin with the hardware and software configurations for the evaluation (Section 5.1). We then show that our method can be applied to data skipping Sudhir et al. [2023], Sun et al. [2014] in modern column-oriented query engines (Section 5.2), as well as to full-text queries Goodwin et al. [2017] that underlie search engines (Section 5.3). Conditioned on the assumption of skewed utility distributions, the results show that we outperform the state of the art on key metrics across the range of possible memory budgets. Additionally, we find that our method displays minimal performance degradation under very tight budgets. We end by characterizing the performance of the proposed data structure and optimization routine (Section 5.4).

5.1 Implementation Details

We first give some details about the hardware and software characteristics that form the basis of the evaluation. Experiments were run on a single machine that has a 1.4 GHz Quad-Core Intel i5 processor, 8 GB of RAM, and a 256 GB SSD.

The research prototype for our method was written in Python 3.10. The truncated Bloom filters were built using the `bitarray` package Schnell [2023]. Filters were constructed with the standard procedure that is commonly used in practice. Specifically, the construction algorithm takes in a target false positive rate $0 < \varepsilon < 1$ (in our case $\varepsilon = 1e - 4$) and the number of elements $n > 0$ that will be added to the filter. It then calculates an estimate for the optimal filter length m and number of hash functions k given these parameters. We employed the efficient `murmurhash3` hash function family Appleby [2015].

The optimization code was implemented in CVXPY Diamond and Boyd [2016], and the

optimal solution was found by an embedded conic solver (ECOS) Domahidi et al. [2013]. ECOS is an interior-point method for second-order cone programs. The relaxation formulation was implemented, since the original objective function contains the non-convex binomial probability mass function, which is outside the scope of the solver. The implementation additionally featured a budget equality constraint, as well as floor function rounding to map the optimal solution onto the natural numbers.

5.2 Data Skipping

The first application that we evaluate our method on is data skipping Sun et al. [2014], Kipf et al. [2020], Sudhir et al. [2023], Tong et al. [2023]. Data skipping is a technique employed by several query engines Databricks [2023], Clickhouse [2024], Armbrust et al. [2020], Dageville et al. [2016], PostgreSQL [2024] to reduce read operations, often when tables are stored in a horizontally-partitioned columnar file format Stonebraker et al. [2005] such as Apache Parquet. In the conventional case, a partition (i.e., row group) is associated with attribute metadata that enables a query to skip certain partitions when the metadata index reveals that there are no matching tuples. By filtering out irrelevant data with a lightweight index, data skipping can significantly accelerate queries that have highly-selective predicates Ta-Shma et al. [2020].

A Bloom filter is often the core data structure underlying point queries over categorical attributes. This is because we can make a determination whether at least one tuple that satisfies the predicates is either certainly not in the partition or is likely in the partition with a small probability of being wrong (i.e., false positive). In our evaluation, we associate a Bloom filter with each categorical attribute in a row group. This yields a total of $N \cdot p$ Bloom filters in the full index. Each attribute filter collection is given a space budget that is a fraction of the original collection size (e.g., 20% of the original). The filters are then truncated according to the mean utility of the tuples contained in a row group.

The metric of success for data skipping with Bloom filters is minimizing the time that is wasted due to false positives. Therefore, the goal in this setting is to produce a truncated Bloom filter index that has only a negligible impact on wasted time and falls within a space budget.

5.2.1 *Experimental Setup*

The setup for the data skipping experiments is now described in detail.

Datasets. We evaluate our approach on 3 datasets. First, Connecticut real estate sales (RE) contains 997,162 tuples and 2 categorical attributes Connecticut [2023]. Each tuple corresponds to a real estate sale in the State of Connecticut over a 20-year period. Second, Washington electric vehicle registrations (EV) contains 162,637 tuples and 6 categorical attributes Washington [2023]. Each tuple is a current electric vehicle registration. Third, NASA HTTP web server logs (NASA) contains 3,452,337 tuples and 1 categorical attribute Laboratory [1995]. Tuples consist of log entries from a web server that hosted the NASA website.

Each dataset has an additional temporal attribute that is used to compute utility values. These attributes are List Year (RE), Model Year (EV), and Timestamp (NASA) respectively. Utility values are calculated as a deterministic exponentially-decaying function of time, with a small amount of additional stochastic noise added. We sort the tuples by utility value before writing the result to a Parquet file. This assumption serves the purpose of simulating a realistic pattern of data arrival. Row group size is systematically determined according to the total number of tuples in a table.

Query Model. Every query on a table \mathcal{T} is associated with a set of p equality predicates formed over categorical attributes A_1, \dots, A_p that have the structure $A_i = a_i$. We further restrict our evaluation to the class of conjunctive queries that limit their result set to the cardinality k . The form of the queries is as follows:

```

SELECT * FROM  $\mathcal{T}$ 
WHERE  $A_1 = a_1$  AND ... AND  $A_p = a_p$ 
LIMIT  $k$ ;

```

We assume that during the execution of a query, we visit partitions in decreasing order of utility until k tuples have been added to the result set, or we reach the lowest utility partition. In the present case, utility is correlated with time, and the likelihood that a particular tuple is added to the result set decreases with tuple age.

Query Generation. We employ a statistical algorithm to generate 250 queries that conform to the structure given above. We compute the joint probability mass function over the p predicates. Our query generation algorithm then builds a set of queries by choosing the most common predicate combinations. We globally set the parameter k for the number of query results to ensure that the distribution of the k -th tuple’s position closely mirrors the utility distribution.

Metrics. Three key evaluation metrics are defined.

- **Skip rate:** The ratio $\frac{\# \text{skipped}}{\# \text{visited}}$ between the number of partitions that are skipped by the indexes and the total number of partitions that are visited by a query.
- **Wasted time:** The total amount of time that is wasted due to index false positives. Wasted time captures when the index returns a positive result for a partition, and the query result set is then found to be empty. This metric is a proxy for false positive rate but is also correlated with predicate selectivity.
- **Query latency:** The execution time of a query. Query latency includes the time from reading the indexes from disk (if applicable), checking the query against the indexes, reading partitions from disk, and running queries on partitions.

Baselines. We define the baselines below.

- **Alphabetical Range (Range):** In each partition, the values for an attribute are first sorted in lexicographical order $y_1 \leq y_2 \leq \dots \leq y_l$. The query attribute value y_q is tested in the predicate $y_1 \leq y_q \leq y_l$. If the predicate evaluates to true, the partition is read from disk and is skipped otherwise.
- **On-disk Filters (Disk):** Each filter is stored at full resolution (i.e., no truncation) on disk Databricks [2023]. When a partition is visited, the filter is read into main memory, and the query is checked against the index.
- **Equal Truncation:** Each filter is truncated according to the policy $m'_i = \lfloor B/F \cdot m_i \rfloor$ where F is the full-resolution (original) size of the index. This approach is an analog of a modular Bloom filter Mun et al. [2022], since we are effectively constructing a collection of smaller in-memory modules.
- **Top Utility:** A policy is applied that closely resembles those commonly used for cache eviction (e.g., least recently used or least frequently used) Chou and DeWitt [1985]. Full-resolution filters are incrementally added to the cache C in decreasing order of utility until the budget is reached, and then the remaining filters are allocated 0 bits. Formally, this greedy algorithm finds the subset of filters $C \subseteq \{\mathcal{B}_1, \dots, \mathcal{B}_N\}$ that maximizes $\sum_{c \in C} u_c$ subject to $\sum_{c \in C} m_c \leq B$.

5.2.2 Results

We measure the performance of each method across the three metrics. For each dataset, we evaluate several budgets that fall in regular increments between 10% and 90% of the original index size. The baselines that guarantee the desired budgets are plotted as curves, while the baselines that have a fixed size are represented as points. We report the median metric for each budget across 5 independent trials.

Equal Truncation and Top Utility have distinct scenarios in which they produce increased




Figure 5.1: Average skip rate on RE, EV, and NASA at different size budgets (higher is better). We generally outperform the in-memory methods and are similar to the on-disk filter.

false positives. On the one hand, Equal Truncation is utility-oblivious in that it has a similar increase in false positive rate across all of the filters in the index. On the other hand, Top Utility is utility-aware but struggles in the setting of long-tailed queries, since it exclusively produces false positives (i.e., cache misses) past a certain point. The goal of our method is to mitigate these two drawbacks and thereby maintain satisfactory performance across the range of possible budgets.

Skip Rate. In Figure 5.1, we see that our method achieves a similar skip rate to the full-resolution filters stored on disk, while generally improving upon the in-memory baselines. Alphabetical range has a near-zero skip rate, which shows that such a strategy is a poor choice for point queries. For this reason, we will omit further discussion about alphabetical range. Notably, the skip rate of Equal Truncation falls sharply at budgets of less than 30%, while our method maintains only a modest reduction in skip rate. Additionally, at a 10% budget, Top Utility falls sharply on two datasets, while our method again maintains a modest reduction.

Wasted Time. Most importantly, we find in Figure 5.2 that our method wastes less time on false positives than the other in-memory baselines. In particular, the distinct points of




Figure 5.2: Average wasted time on RE, EV, and NASA at different index size budgets B (lower is better). We outperform the in-memory methods.

failure for Equal Truncation and Top Utility result in significant wasted time in scenarios with tight budgets. Consider the NASA dataset as a representative example. It is observed that at a 10% budget, our method wastes less time than both Equal Truncation and Top Utility by a factor of 3. The performance advantage holds until Top Utility catches up to our method at a budget of 40% of the original index size. Equal Truncation also eventually catches up at a budget of 60% of the original index size.

Furthermore, our method performs better than the alternatives under varying attribute counts and predicate selectivity. Indeed, low predicate selectivity can lead to increased wasted time, a conclusion that can be reached by examining the adversarial EV dataset (middle). The reason for the increased wasted time is because there are many partitions that contain all of the attribute values, even though few tuples jointly satisfy the predicates. The performance gap between our method and the in-memory baselines is smaller in this specific instance.

Query Latency. We show in Figure 5.3 that our method is generally the fastest overall when compared to each of the baselines. The reason for this is two-fold. First, by minimizing the false positive rate in an optimal fashion, we reduce the impact that the failure modes of Equal Truncation and Top Utility have on wasted time, as aforementioned. Second, our




Figure 5.3: Average query latency on RE, EV, and NASA at different index size budgets B (lower is better). We are generally the lowest latency choice across all methods and budgets.

index sits in memory, which accelerates our queries in most instances when compared to disk.

Disaggregation. Although our experiments were run on a machine with physically-allocated compute and storage, we run an additional experiment on the NASA dataset to simulate the effect of increasing resource disaggregation. It is assumed that disaggregation applies a linear multiplicative cost factor δ to the wasted time. Each factor gets its own curve. In Figure 5.4, it is shown that when storage and compute are increasingly separated, the cost of time wasted on false positives can rise, leading our method to outperform the other baselines by larger margins. We reach this conclusion by examining the shaded region between the pairs of cost factor curves.

5.3 Full-Text Search

Full-text search is a fundamental problem in information retrieval. In the context of search engines, the task is to find the documents that match the search terms (i.e., tokens), and then rank the matching documents according to some notion of relevance. If we represent each document as a set of tokens, we find that the matching problem is no different than issuing membership queries to the document. Following this intuition, a document can

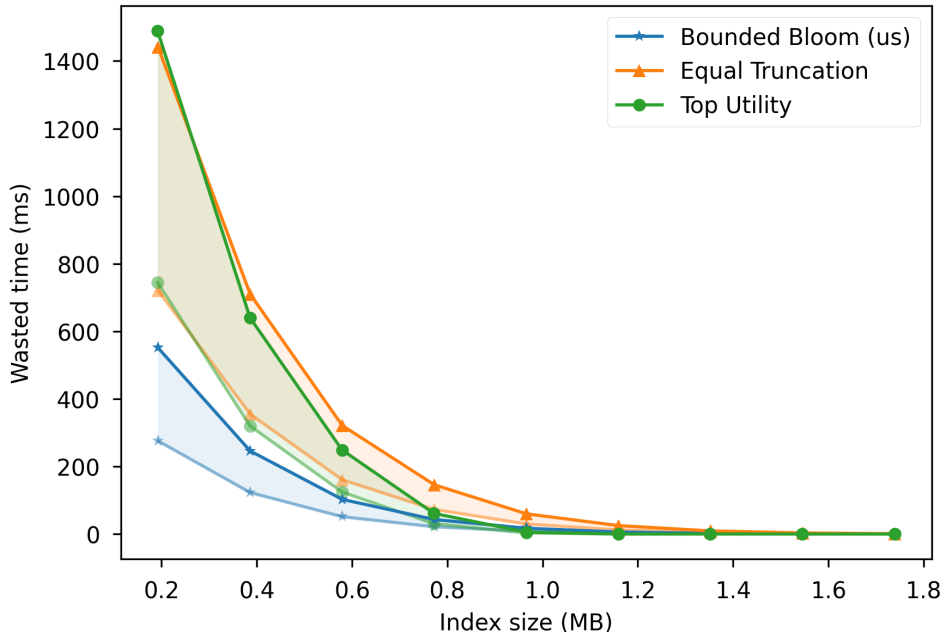


Figure 5.4: Increasing compute-storage disaggregation can greatly magnify the cost of false positives.

be represented as a Bloom filter (i.e., a signature file). Provided that the Bloom filters are constructed with the same hash functions and filter lengths, we can compute the set intersection between a query filter and a document filter in $\mathcal{O}(m)$ time by simply applying a bit-wise AND.

The state-of-the-art approach for full-text search is the inverted index. In its most basic form, an inverted index maps a search term to a list of documents that contain the term. Since documents are represented as fixed-length integer identifiers, the total size of an inverted index grows quickly, even with documents that are rarely returned in a search result. The prevailing space-reduction technique for inverted indexes is lossless compression (e.g., Ottaviano and Venturini [2014]), which does not exploit document utility skew. In addition, caching often loads the full document lists into memory for frequent search tokens Zhang et al. [2008]. This means that in the absence of careful optimization, cache space is wasted on low-utility documents.

If instead each document is represented as a Bloom filter, we can derive a lossy variable-

length signature for each document by directly applying our method. Taking inspiration from BitFunnel Goodwin et al. [2017], a commercial search engine that employs a Bloom filter index in place of an inverted index, we will demonstrate how our method can be used to effectively optimize a collection of Bloom filter signatures according to utility.

5.3.1 *Experimental Setup*

Our setup for full-text search is different from data skipping, so we will again walk through each of the points individually.

Datasets. We use Amazon review datasets that are from two product categories Ni et al. [2019], Ni [2018]. The first category is industrial and scientific (I&S) with 49,595 cleaned documents, and the second product category is musical instruments (MI) with 160,523 cleaned documents. We remove punctuation and stopwords before tokenizing the documents. We remove the minority of documents that do not contain between 5 and 100 tokens to produce a realistic document shard Goodwin et al. [2017].

Query Model. We assume that each document in the corpus has a utility derived from a PageRank-like algorithm Brin and Page [1998]. In this setting, we adopt a top- k query model. In particular, the probability that a given document appears in the top- k query results is proportional to its utility. The underlying skewness assumption is that only a small number of documents from the corpus end up appearing in the top- k search results. Similar to data skipping, our method enumerates documents in descending order of utility.

Query Generation. We consider the case of $k = 1$ and generate 250 queries using the following procedure. Given the utility values u_1, \dots, u_N we first normalize to a probability distribution like so: $\mathbb{P}_i = u_i / \sum_{j \in [N]} u_j$. To generate a query, we sample a document from the probability distribution and choose n terms from the selected document. The specific terms are chosen so that they are only jointly present in a few documents on average.

Metrics. We begin the discussion of metrics with some notation. Let D be the document

corpus. \mathcal{T}_q is defined as the set of tokens in query q and \mathcal{T}_d is the set of tokens in document d . The set R contains the documents retrieved by a particular query. The match set of a query $M = \{d : \mathcal{T}_d \cap \mathcal{T}_q = \mathcal{T}_q\}$ captures the documents that contain every token in the query Goodwin et al. [2017]. Following this definition, the set $G \subseteq M$ represents the ground truth top- k matching documents in terms of utility. Finally, we denote the utility value of document d by u_d . We are now ready to define the metrics.

- **Precision@ k** : $\frac{|R \cap G|}{k}$. For a Bloom filter index, precision is a salient metric for false positives. For other indexes, precision is higher when there are fewer false negatives.
- **Query latency**: The execution time of a query.
- **Average utility**: $\frac{1}{k} \sum_{d \in R} u_d$.
- **Average overlap coefficient**: $\frac{1}{k} \sum_{d \in R} \frac{|\mathcal{T}_d \cap \mathcal{T}_q|}{\min(|\mathcal{T}_d|, |\mathcal{T}_q|)}$. For a Bloom filter, the overlap coefficient measures the average token-wise effect of false positives.

Baselines. The baselines are defined below.

- **Inverted Index (II)**: The standard approach Brin and Page [1998] where the index maps from a term t to the documents that contain the term such that $t \rightarrow [d_1, d_2, \dots, d_r]$. Documents are represented as integer identifiers. The size of each document d is $\mathcal{O}(|\mathcal{T}_d|)$. The size of the index grows quickly, even with low utility documents.
- **Top- M Inverted Index (TMII)**: Documents are added to the inverted index until the budget is reached using the same general approach as Top Utility. The greedy algorithm finds the highest-utility subset of documents with cardinality M . Note that a document either exists in the index with all of its terms, or does not exist at all.
- **Top- k Inverted Index (TKII)**: The value k is first specified. A full inverted index is constructed using the procedure described above. Each term’s document identifier

list l is sorted in descending order of utility and then truncated to contain only the top- k highest-utility documents. The key difference from the Top- M Inverted Index is that token false negatives are possible, since only a (potentially empty) subset of the terms for each document is now stored.

- **Top- M Document Set (TMDS):** A forward index is maintained over the M highest-utility documents, where every document is stored as the set of terms \mathcal{T}_d . A query enumerates the documents from high to low utility and directly checks if $\mathcal{T}_q \cap \mathcal{T}_d = \mathcal{T}_q$.
- **Scan (SCN):** Each document is stored in a forward index on disk and is sequentially read at query time. Upon completion of the scan, the matching documents are sorted by utility, and then the top- k are returned.

5.3.2 Results

We select a range of space budgets by first calculating the smallest index size between the inverted index, forward index, and Bloom filter index. The budget is then varied from 10% to 90% of the smallest index size in regular increments, just like in the previous application. We find that the smallest index size is either the inverted index or the Bloom filter index. We evaluate the approaches on in-memory performance (solid lines), but we also include a version of our method that checks an on-disk forward index when the filter returns a positive result (dotted line). We also add a basic Bloom filter index that is not truncated as an additional point of comparison. The results for precision and latency across the two datasets are highlighted in Figures 5.5 and 5.6 respectively. We defer the rest of the metrics to the microbenchmarks.

In Figure 5.5, we observe that our method has higher precision than the other in-memory baselines at larger budgets and rivals the Top- M Inverted Index at smaller budgets. It is clear that the Top- M Document Set (forward) index is highly space-inefficient, which results

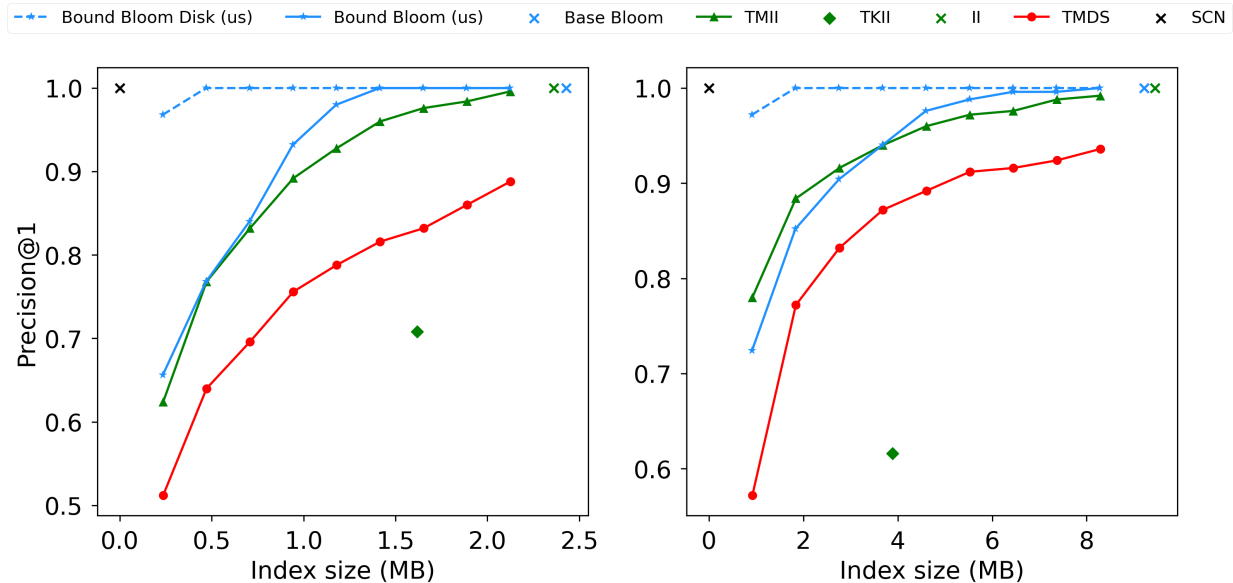


Figure 5.5: I&S (left) and MI (right). We have higher precision than the best-performing baseline (TMII) at larger budgets and rival its performance at smaller budgets.

in reduced performance. When comparing the inverted index approaches, it is evident that the Top- M Inverted Index is strictly better than the naive strategy undertaken by the Top- k Inverted Index where the document identifier lists are truncated. At the lone space budget of the Top- k Inverted Index, the Top- M variant is up to 60% higher in terms of precision.

Next, in Figure 5.6 our method is found to be slower than the other approaches, besides the full scan. The reason for this result is two-fold. First, inverted indexes are able to reduce the search space rather significantly, which avoids the need for sequential iteration through the document corpus. This is not so much a surprising result, but it is worth noting that a combination of algorithmic optimizations and careful implementation can be undertaken to make Bloom signatures faster than conventional inverted indexes Goodwin et al. [2017]. Second, the implementation style that we ultimately chose likely had a sizable impact on latency.

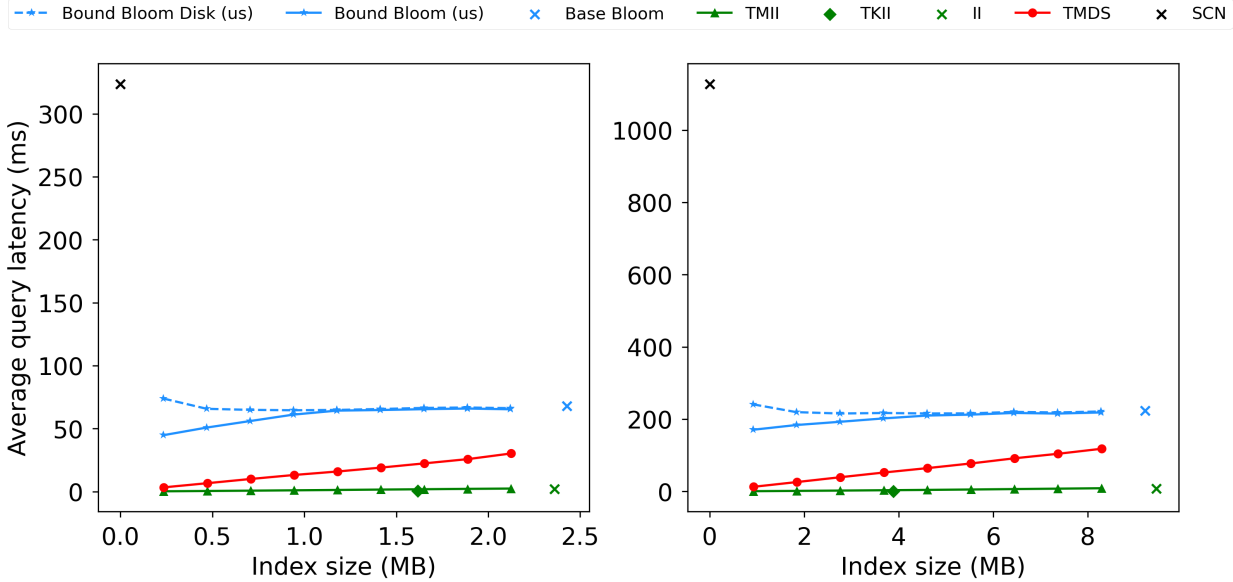


Figure 5.6: I&S (left) and MI (right). We are slower than the baselines besides scan but are impeded by implementation.

5.4 Microbenchmarks

Lastly, we evaluate the main characteristics of our method. In particular, we highlight the truncated Bloom filter, our optimization routine, as well as several alternative algorithms for constructing a Bloom filter collection within a space budget.

5.4.1 Truncated Bloom Filters

As a reminder, the parameter $p = \frac{m'}{m}$ is the ratio between the number of bits that are kept after truncation (m') and original filter size (m). In Figure 5.7, we show the effect of truncation on both the false positive rate and filter query latency. We generate an equal number of positive q^+ and negative q^- keys (20,000 each). The keys are generated in such a way that the intersection between the positive and negative sets is empty.

After adding the positive keys, we query the filter with the negative keys at each truncation ratio, and the false positive rate is measured empirically. We compare the empirical false positive rate to the expected false positive rate from Equation (3.2) and the lower

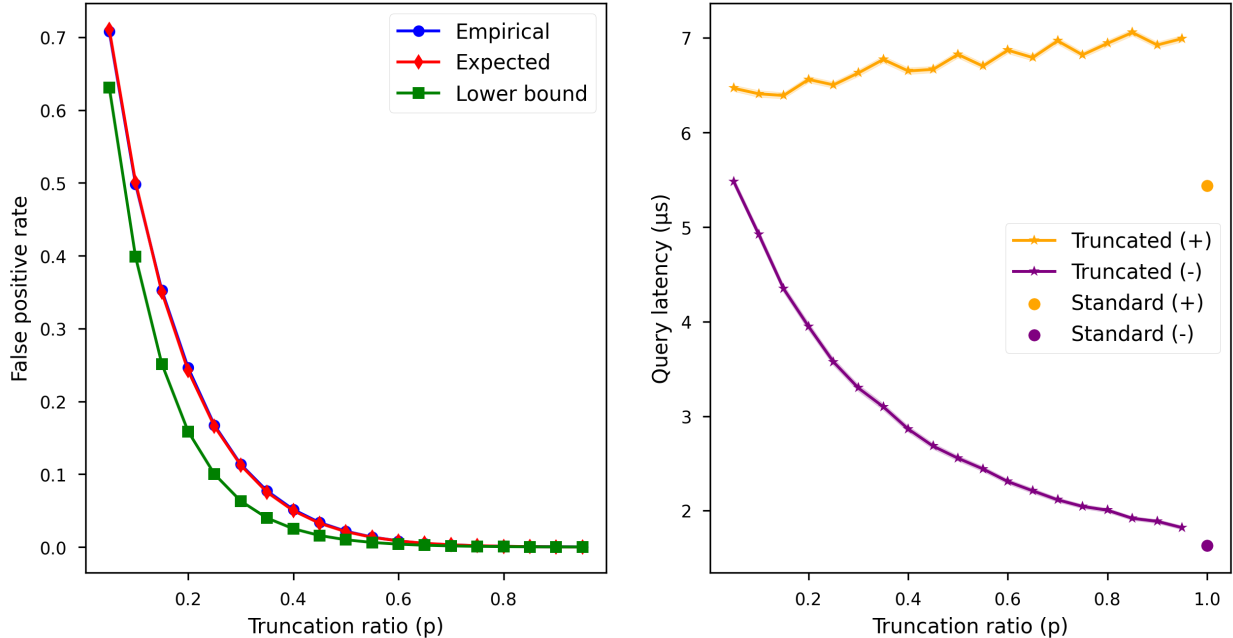


Figure 5.7: FPR (left) and query latency (right) as a function of the truncation ratio p . Latency is given as 95% confidence intervals, + and - are positive and negative keys respectively.

bound from Equation (4.1). We observe that given a sufficiently large number of negative queries, the empirical false positive rate (blue) is practically identical to the expected false positive rate (red), which is consistent with the law of large numbers. It is also evident that the gap between the expected false positive rate (red) and the lower bound (green) as given by the function $\mathbb{E}[\phi(V)] - \phi(\mathbb{E}[V])$ converges to 0 as $p \rightarrow 1$. This indicates that the lower bound approximation is better when fewer bits are truncated. Equivalently, the optimality gap between the relaxation and original formulation would generally be smaller under larger budgets.

Next, the truncated filter is queried with both negative and positive keys to measure query latency. The result is compared to a standard filter. We see that there are two sources of overhead. First, the extra condition in Algorithm 1 that checks for valid hash functions introduces a small amount of latency overhead when compared to the standard filter, which is seen in the vertical distance between the rightmost truncated measurements and each

standard measurement. Second, as the filter is increasingly truncated, the probability p of observing a valid hash function decreases, leading to additional latency overhead in the negative case only. We also see that the latency for the positive case decreases slightly, since fewer filter positions are probed in total.

These results suggest that our proposed data structure is particularly suitable in the high truncation regime $m' \ll m$ for applications where downstream costs far exceed the cost of querying the filter. For example, in one of the applications that we presented, the cost of checking a partition is several orders of magnitude higher than querying the filter. However, given a collection of truncated Bloom filters with varying utility values, we would expect that the latency overhead is minimal in expectation. This is because a truncated filter that operates in the low truncation regime $m' \approx m$ exhibits only a slight increase in latency on negative queries when compared to the standard filter.

We can also confirm the latency result with a theoretical argument. Let $X \sim G(p)$ be a geometric random variable for the number of hash functions that are checked to observe a valid hash function with mean $\mathbb{E}[X] = \frac{1}{p}$. We find that, in expectation, the number of required probes increases for negatives and stays the same for positives, since only $\Omega(1)$ hash functions are required for a negative while $\Theta(k)$ hash functions are required for a positive.

However, if we treat a truncated Bloom filter as a set signature like in full-text search, then the intersection and union operations are computed in *reduced* time, provided that the filters were constructed with shared hash functions. For example, given two sets S_1 and S_2 represented as truncated Bloom filters, an approximation of $S_1 \cap S_2$ can be found in $\mathcal{O}(\min(m'_1, m'_2))$ time via bitwise operations.

5.4.2 Optimization Latency

We now evaluate our lower bound optimization approach by measuring the solver latency at a budget of half the original collection size and different quantities of Bloom filters. We gen-




Figure 5.8: Solving the relaxation scales to large inputs. The resulting throughput is approximately 20 thousand filters per second on a resource-constrained machine.

erate N synthetic parameter tuples of the form (m_i, k_i, n_i) by sampling from a multivariate probability distribution. We then measure the solver latency as the number of Bloom filters N is varied from 1 thousand to 1 million in increments of 50 thousand. We report the solver latency as the mean across 5 runs with the same input parameters.

As seen in Figure 5.8, the solver run time scales linearly with N and the optimization problem can be solved quickly even on a resource-constrained machine. For example, an (approximate) optimal solution can be found for a collection of half a million Bloom filters in around 30 seconds. This observation highlights the work that was done in Section 4.2 to reduce the time complexity of the objective function.

We reach the conclusion that the optimization overhead of our method is small, and we can quickly adapt the collection under changing utility values, or as new filters are added. In other words, this is a desirable property in the online case. Crucially, if we store a copy of the original filters on disk, we can re-optimize without having to reconstruct the Bloom filters from scratch. This is one of the key advantages of our method.




Figure 5.9: Our method harmonizes the benefits of Top Utility and Equal Truncation.

5.4.3 Alternative Algorithms

We now explore several alternative algorithms for constructing a filter collection that respects a space budget. These techniques broadly fall into two categories: sampling and truncation. A sampling algorithm finds a smaller subset of full-resolution filters, whereas a truncation algorithm constructs a lower-resolution version of each filter. The difference between these two categories is exemplified by Equal Truncation and Top Utility, as defined in Section 5.2.1. The two remaining sampling methods are Random Uniform, which samples filters with equal probability, and Random Utility, which samples filters according to the utility distribution.

We evaluate the in-memory performance of these alternative algorithms on the I&S dataset. The reader is referred to Section 5.3.1 for a review of the evaluation metrics for full-text search. Note that the utility distribution is highly skewed. The results are presented in Figure 5.9. In general, we find that our method effectively harmonizes the benefits of Top Utility and Equal Truncation, while simultaneously mitigating the drawbacks of both approaches. Like Top Utility, our method maintains satisfactory performance at smaller budgets, precisely where Equal Truncation fails. Like Equal Truncation, we achieve near-perfect performance at larger budgets, while Top Utility misses the documents in long-tailed queries.

It is evident that our method generally has a higher overlap coefficient than precision. This means that when false positives do occur, they are often with respect to only a few




Figure 5.10: Skewed distributions imply better performance.

tokens, as opposed to many tokens. Additionally, we find that with our method, the average utility value quickly converges to the true value. Notice that the uptick in average utility at smaller budgets for truncation stems from false positives. On the contrary, the decline in average utility at smaller budgets for sampling is a consequence of false negatives. Finally, it is observed that truncation slightly reduces the in-memory query latency when compared to the alternative of full-resolution filters.

5.4.4 Different Utility Distributions

Our method is sensitive to the specific utility distribution. To evaluate the performance in this setting, we carefully construct a set of utility distributions that differ from one another. We begin with a uniform and normal distribution that have a shared mean and variance. We then create a sequence of right-skewed utility distributions by generating mixtures of normal distributions that exhibit increasing distance between the means of the constituent density functions (governed by the parameter ρ).

In Figure 5.10, we see the effect of the utility distribution on precision (again the I&S dataset). Shape is the first factor, as seen by contrasting the uniform distribution against the normal distribution with the exact same mean and variance. Referring back to Figure 5.9, when utility values are concentrated around the mean and the budget is tight, our method reduces to truncating the filters equally. As the parameter governing the skewness ρ increases, we observe that the performance in the tight budget regime gradually improves until we mirror Top Utility.

CHAPTER 6

RELATED WORK

6.1 Filters

There are several filter variants such as learned filters Kraska et al. [2018], Deeds et al. [2020], expanding filters Dayan et al. [2023], and compressed filters Mitzenmacher [2002]. Other examples include the cuckoo filter Fan et al. [2014] and quotient filter Pandey et al. [2017]. While truncation may be a valid operation on some of these filter variants, we leave that possibility to future work.

6.2 Bloom Filters in Key-value Stores

The most similar body of work to our method is the application of Bloom filters in key-value stores. Key-value stores are represented as log-structured merge (LSM) trees, where each level in the LSM tree is associated with Bloom filters to prevent unnecessary reads from slower storage. Previous work such as Monkey Dayan et al. [2017] and ElasticBF Zhang et al. [2018] has shown that it is suboptimal to assign filter false positive rates uniformly across the different LSM tree levels. This is a restricted case of the framework that we present in this work.

6.3 Modular Bloom Filters

The closest data structure to the truncated Bloom filter is the modular Bloom filter Mun et al. [2022] which in turn is similar to the blocked Bloom filter Putze et al. [2010]. Like a modular Bloom filter, we employ a smaller in-memory filter at the expense of a higher false positive rate and do not require re-indexing. Unlike a modular Bloom filter, we are better able to navigate the space versus accuracy trade-off, since truncation removes one bit at a

time in a fine-grained fashion. A modular Bloom filter is limited to the product of false positive rates across several coarse-grained modules.

6.4 Other Bloom Filter Applications

There are many applications of Bloom filters besides those that we presented in this paper. As one example, Bloom filters are often used to optimize joins Gaffney et al. [2022], Zhu et al. [2017], Services [2020], Ting and Cole [2021]. As another example, Bloom filters are a vital data structure for in-memory data stores such as Redis Redis [2024].

CHAPTER 7

CONCLUSION

We presented an approach that optimizes a collection of Bloom filters across a utility distribution to satisfy a space budget. It was shown that truncating a Bloom filter after construction can lead to a subtly different trade-off between accuracy and space that benefits optimization. Fittingly, we developed a convex optimization problem that admits a fast relaxation to navigate this trade-off. Moving from theory to practice, we demonstrated on two separate applications that our method effectively mitigates the limitations of existing approaches.

REFERENCES

- Austin Appleby. murmurhash3, 2015. URL <https://github.com/aappleby/smhasher/blob/master/src/MurmurHash3.cpp>.
- Michael Armbrust, Tathagata Das, Liwen Sun, Burak Yavuz, Shixiong Zhu, Mukul Murthy, Joseph Torres, Herman van Hovell, Adrian Ionescu, Alicja Łuszczak, Michał Świtakowski, Michał Szafranski, Xiao Li, Takuya Ueshin, Mostafa Mokhtar, Peter Boncz, Ali Ghodsi, Sameer Paranjpye, Pieter Senster, Reynold Xin, and Matei Zaharia. Delta lake: high-performance acid table storage over cloud object stores. *Proc. VLDB Endow.*, 13(12): 3411–3424, aug 2020. ISSN 2150-8097. doi:10.14778/3415478.3415560. URL <https://doi.org/10.14778/3415478.3415560>.
- Bruno Barbarioli, Gabriel Mersy, Stavros Sintos, and Sanjay Krishnan. Hierarchical residual encoding for multiresolution time series compression. *Proc. ACM Manag. Data*, 1(1), may 2023. doi:10.1145/3588953. URL <https://doi.org/10.1145/3588953>.
- Benjamin Berg, Daniel S. Berger, Sara McAllister, Isaac Grosf, Sathya Gunasekar, Jimmy Lu, Michael Uhlar, Jim Carrig, Nathan Beckmann, Mor Harchol-Balter, and Gregory R. Ganger. The CacheLib caching engine: Design and experiences at scale. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 753–768. USENIX Association, November 2020. ISBN 978-1-939133-19-9. URL <https://www.usenix.org/conference/osdi20/presentation/berg>.
- Burton H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM*, 13(7):422–426, jul 1970. ISSN 0001-0782. doi:10.1145/362686.362692. URL <https://doi.org/10.1145/362686.362692>.
- Prosenjit Bose, Hua Guo, Evangelos Kranakis, Anil Maheshwari, Pat Morin, Jason Morrison, Michiel Smid, and Yihui Tang. On the false-positive rate of bloom filters. *Information Processing Letters*, 108(4):210–213, 2008. ISSN 0020-0190. doi:<https://doi.org/10.1016/j.ipl.2008.05.018>. URL <https://www.sciencedirect.com/science/article/pii/S0020019008001579>.
- Stephen P Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30(1):107–117, 1998. ISSN 0169-7552. doi:[https://doi.org/10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X). URL <https://www.sciencedirect.com/science/article/pii/S016975529800110X>. Proceedings of the Seventh International World Wide Web Conference.
- Lawrence D. Brown. Fundamentals of statistical exponential families with applications in statistical decision theory. *Lecture Notes-Monograph Series*, 9:i–279, 1986. ISSN 07492170. URL <http://www.jstor.org/stable/4355554>.

- Hong-Tai Chou and David J. DeWitt. An evaluation of buffer management strategies for relational database systems. In *Proceedings of the 11th International Conference on Very Large Data Bases - Volume 11, VLDB '85*, page 127–141. VLDB Endowment, 1985.
- Clickhouse. Understanding clickhouse data skipping indexes, 2024. URL <https://clickhouse.com/docs/en/optimize/skipping-indexes>.
- State of Connecticut. Real estate sales 2001-2020 gl, 2023. URL <https://catalog.data.gov/dataset/real-estate-sales-2001-2018>.
- Graham Cormode and S. Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005. ISSN 0196-6774. doi:<https://doi.org/10.1016/j.jalgor.2003.12.001>. URL <https://www.sciencedirect.com/science/article/pii/S0196677403001913>.
- Benoit Dageville, Thierry Cruanes, Marcin Zukowski, Vadim Antonov, Artin Avanes, Jon Bock, Jonathan Claybaugh, Daniel Engovatov, Martin Hentschel, Jiansheng Huang, Allison W. Lee, Ashish Motivala, Abdul Q. Munir, Steven Pelley, Peter Povinec, Greg Rahn, Spyridon Triantafyllis, and Philipp Unterbrunner. The snowflake elastic data warehouse. In *Proceedings of the 2016 International Conference on Management of Data, SIGMOD '16*, page 215–226, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450335317. doi:10.1145/2882903.2903741. URL <https://doi.org/10.1145/2882903.2903741>.
- Databricks. Bloom filter indexes, 2023. URL <https://docs.databricks.com/en/optimizations/bloom-filters.html>.
- Niv Dayan, Manos Athanassoulis, and Stratos Idreos. Monkey: Optimal navigable key-value store. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD '17*, page 79–94, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450341974. doi:10.1145/3035918.3064054. URL <https://doi.org/10.1145/3035918.3064054>.
- Niv Dayan, Ioana Bercea, Pedro Reviriego, and Rasmus Pagh. Infinifilter: Expanding filters to infinity and beyond. *Proc. ACM Manag. Data*, 1(2), jun 2023. doi:10.1145/3589285. URL <https://doi.org/10.1145/3589285>.
- Kyle Deeds, Brian Hentschel, and Stratos Idreos. Stacked filters: Learning to filter by structure. *Proc. VLDB Endow.*, 14(4):600–612, dec 2020. ISSN 2150-8097. doi:10.14778/3436905.3436919. URL <https://doi.org/10.14778/3436905.3436919>.
- Steven Diamond and Stephen Boyd. Cvxpy: a python-embedded modeling language for convex optimization. *J. Mach. Learn. Res.*, 17(1):2909–2913, jan 2016. ISSN 1532-4435.
- Alexander Domahidi, Eric Chu, and Stephen Boyd. Ecos: An socp solver for embedded systems. In *2013 European Control Conference (ECC)*, pages 3071–3076, 2013. doi:10.23919/ECC.2013.6669541.

- Bin Fan, Dave G. Andersen, Michael Kaminsky, and Michael D. Mitzenmacher. Cuckoo filter: Practically better than bloom. In *Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies*, CoNEXT '14, page 75–88, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450332798. doi:10.1145/2674005.2674994. URL <https://doi.org/10.1145/2674005.2674994>.
- Kevin P. Gaffney, Martin Prammer, Larry Brasfield, D. Richard Hipp, Dan Kennedy, and Jignesh M. Patel. Sqlite: past, present, and future. *Proc. VLDB Endow.*, 15(12): 3535–3547, aug 2022. ISSN 2150-8097. doi:10.14778/3554821.3554842. URL <https://doi.org/10.14778/3554821.3554842>.
- Bob Goodwin, Michael Hopcroft, Dan Luu, Alex Clemmer, Mihaela Curmei, Sameh El-nikety, and Yuxiong He. Bitfunnel: Revisiting signatures for search. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '17, page 605–614, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450350228. doi:10.1145/3077136.3080789. URL <https://doi.org/10.1145/3077136.3080789>.
- Andreas Kipf, Damian Chromejko, Alexander Hall, Peter Boncz, and David G. Andersen. Cuckoo index: a lightweight secondary index structure. *Proc. VLDB Endow.*, 13(13): 3559–3572, sep 2020. ISSN 2150-8097. doi:10.14778/3424573.3424577. URL <https://doi.org/10.14778/3424573.3424577>.
- Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*, SIGMOD '18, page 489–504, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450347037. doi:10.1145/3183713.3196909. URL <https://doi.org/10.1145/3183713.3196909>.
- Lawrence Berkeley National Laboratory. Nasa-http, 1995. URL <https://ita.ee.lbl.gov/html/contrib/NASA-HTTP.html>.
- Michael Mitzenmacher. Compressed bloom filters. *IEEE/ACM Transactions on Networking*, 10(5):604–612, 2002. doi:10.1109/TNET.2002.803864.
- Ju Hyoung Mun, Zichen Zhu, Aneesh Raman, and Manos Athanassoulis. Lsm-trees under (memory) pressure. In Rajesh Bordawekar and Tirthankar Lahiri, editors, *International Workshop on Accelerating Analytics and Data Management Systems Using Modern Processor and Storage Architectures, ADMS@VLDB 2022, Sydney, Australia, September 5, 2022*, pages 23–35, 2022. URL <https://cs-people.bu.edu/mathan/publications/adms22-mun.pdf>.
- Jianmo Ni. Amazon review data (2018), 2018. URL https://cseweb.ucsd.edu/~jmcauley/datasets/amazon_v2/.
- Jianmo Ni, Jiacheng Li, and Julian McAuley. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In Kentaro Inui, Jing Jiang, Vincent Ng,

- and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 188–197, Hong Kong, China, November 2019. Association for Computational Linguistics. doi:10.18653/v1/D19-1018. URL <https://aclanthology.org/D19-1018>.
- Giuseppe Ottaviano and Rossano Venturini. Partitioned elias-fano indexes. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR '14, page 273–282, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450322577. doi:10.1145/2600428.2609615. URL <https://doi.org/10.1145/2600428.2609615>.
- Prashant Pandey, Michael A. Bender, Rob Johnson, and Rob Patro. A general-purpose counting filter: Making every bit count. In *Proceedings of the 2017 ACM International Conference on Management of Data*, SIGMOD '17, page 775–787, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450341974. doi:10.1145/3035918.3035963. URL <https://doi.org/10.1145/3035918.3035963>.
- PostgreSQL. F.7. bloom — bloom filter index access method, 2024. URL <https://www.postgresql.org/docs/current/bloom.html>.
- Felix Putze, Peter Sanders, and Johannes Singler. Cache-, hash-, and space-efficient bloom filters. *ACM J. Exp. Algorithmics*, 14, jan 2010. ISSN 1084-6654. doi:10.1145/1498698.1594230. URL <https://doi.org/10.1145/1498698.1594230>.
- Redis. Bloom filter, 2024. URL <https://redis.io/docs/data-types/probabilistic/bloom-filter/>.
- Ilan Schnell. bitarray, 2023. URL <https://pypi.org/project/bitarray/>.
- Tore Schweder. On the dispersion of mixtures. *Scandinavian Journal of Statistics*, 9(3):165–169, 1982. ISSN 03036898, 14679469. URL <http://www.jstor.org/stable/4615873>.
- Amazon Web Services. Amazon redshift now leverages bloom filters to improve data lake query performance by up to 2x, 2020. URL <https://aws.amazon.com/about-aws/whats-new/2020/05/amazon-redshift-now-leverages-bloom-filters-to-improve-data-lake-query-performance/>.
- Moshe Shaked. On mixtures from exponential families. *Journal of the Royal Statistical Society. Series B (Methodological)*, 42(2):192–198, 1980. ISSN 00359246. URL <http://www.jstor.org/stable/2984960>.
- Mike Stonebraker, Daniel J. Abadi, Adam Batkin, Xuedong Chen, Mitch Cherniack, Miguel Ferreira, Edmond Lau, Amerson Lin, Sam Madden, Elizabeth O’Neil, Pat O’Neil, Alex Rasin, Nga Tran, and Stan Zdonik. C-store: a column-oriented dbms. In *Proceedings of the 31st International Conference on Very Large Data Bases*, VLDB '05, page 553–564. VLDB Endowment, 2005. ISBN 1595931546.

- Sivaprasad Sudhir, Wenbo Tao, Nikolay Laptev, Cyrille Habis, Michael Cafarella, and Samuel Madden. Pando: Enhanced data skipping with logical data partitioning. *Proc. VLDB Endow.*, 16(9):2316–2329, may 2023. ISSN 2150-8097. doi:10.14778/3598581.3598601. URL <https://doi.org/10.14778/3598581.3598601>.
- Liwen Sun, Michael J. Franklin, Sanjay Krishnan, and Reynold S. Xin. Fine-grained partitioning for aggressive data skipping. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, page 1115–1126, New York, NY, USA, 2014. Association for Computing Machinery. ISBN 9781450323765. doi:10.1145/2588555.2610515. URL <https://doi.org/10.1145/2588555.2610515>.
- Paula Ta-Shma, Guy Khazma, Gal Lushi, and Oshrit Feder. Extensible data skipping. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 372–382, Los Alamitos, CA, USA, dec 2020. IEEE Computer Society. doi:10.1109/BigData50022.2020.9377740. URL <https://doi.ieeecomputersociety.org/10.1109/BigData50022.2020.9377740>.
- Daniel Ting and Rick Cole. Conditional cuckoo filters. In *Proceedings of the 2021 International Conference on Management of Data*, SIGMOD '21, page 1838–1850, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450383431. doi:10.1145/3448016.3452811. URL <https://doi.org/10.1145/3448016.3452811>.
- Yulai Tong, Jiazhen Liu, Hua Wang, Ke Zhou, Rongfeng He, Qin Zhang, and Cheng Wang. Sieve: A learned data-skipping index for data analytics. *Proc. VLDB Endow.*, 16(11):3214–3226, jul 2023. ISSN 2150-8097. doi:10.14778/3611479.3611520. URL <https://doi.org/10.14778/3611479.3611520>.
- State of Washington. Electric vehicle population data, 2023. URL <https://catalog.data.gov/dataset/electric-vehicle-population-data>.
- Jiangong Zhang, Xiaohui Long, and Torsten Suel. Performance of compressed inverted list caching in search engines. In *Proceedings of the 17th International Conference on World Wide Web*, WWW '08, page 387–396, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605580852. doi:10.1145/1367497.1367550. URL <https://doi.org/10.1145/1367497.1367550>.
- Yueming Zhang, Yongkun Li, Fan Guo, Cheng Li, and Yinlong Xu. Elasticbf: Fine-grained and elastic bloom filter towards efficient read for lsm-tree-based kv stores. In *Proceedings of the 10th USENIX Conference on Hot Topics in Storage and File Systems*, HotStorage'18, page 11, USA, 2018. USENIX Association.
- Jianqiao Zhu, Navneet Potti, Saket Saurabh, and Jignesh M. Patel. Looking ahead makes query plans robust: making the initial case with in-memory star schema data warehouse workloads. *Proc. VLDB Endow.*, 10(8):889–900, apr 2017. ISSN 2150-8097. doi:10.14778/3090163.3090167. URL <https://doi.org/10.14778/3090163.3090167>.