# Support Optimality and Adaptive Cuckoo Filters* 

Tsvi Kopelowitz ${ }^{\dagger} \quad$ Samuel McCauley ${ }^{\ddagger} \quad$ Ely Porat ${ }^{\S}$


#### Abstract

Filters (such as Bloom Filters) are a fundamental data structure that speed up network routing and measurement operations by storing a compressed representation of a set. Filters are very space efficient, but can make bounded one-sided errors: with tunable probability $\epsilon$, they may report that a query element is stored in the filter when it is not. This is called a false positive. Recent research has focused on designing methods for dynamically adapting filters to false positives, thereby reducing the number of false positives when some elements are queried repeatedly.

Ideally, an adaptive filter would incur a false positive with bounded probability $\epsilon$ for each new query element, and would incur $o(\epsilon)$ total false positives over all repeated queries to that element. We call such a filter support optimal.

In this paper we design a new Adaptive Cuckoo Filter, and show that it is support optimal (up to additive logarithmic terms) over any $n$ queries when storing a set of size $n$. Our filter is very simple: fixing previous false positives requires a simple cuckoo-like operation, and the filter does not need to store any additional metadata. This data structure is the first practical data structure that is support optimal, and the first support optimal filter that does not require additional space beyond a normal cuckoo filter.

We complement these bounds with experiments that show that our data structure is effective at fixing false positives on network trace datasets, outperforming previous Adaptive Cuckoo Filters.

Finally, we investigate adversarial adaptivity, a stronger notion of adaptivity in which an adaptive adversary repeatedly queries the filter, using the result of previous queries to drive the false positive rate as high as possible. We prove a lower bound showing that a broad family of filters, including all known Adaptive Cuckoo Filters, can be forced by such an adversary to incur a large number of false positives.


## 1 Introduction

A filter is a data structure that supports membership queries for a set of elements $S=x_{1}, \ldots x_{n}$ from a universe $U$. The answer to each filter query is present or absent. Typically, a filter has a correctness guarantee: if an element $q \in S$, the filter must return present to the query with probability 1 . There is also a performance guarantee: if an element $q \notin S$, the filter must return present with tunable probability at most $\epsilon$. If a query on an element $q \notin S$ returns present then $q$ is called a false positive. Typically, filters use a small amount of space.

A filter's small size means that the filter can be stored in an efficiently accessible location. Meanwhile, the no-false-negative guarantee implies that if the filter returns $q \notin S$ for a query $q$, then there is no need for accessing the actual data, which is typically stored in a medium with expensive access cost. This ability to filter out queries to items not in $S$ in a small-size structure has found a wide variety of network applications such as collaboration in peer-to-peer networks, resource routing, packet routing, and measurement infrastructures 6] as well as many areas of network security 14 .

[^0]There are several different kinds of filters. The Bloom filter 5] was the first filter data structure to be designed; it is still very popular due to its simplicity and efficiency. Later filters were designed to provide better worst-case lookup times and space guarantees [3,20,23], improved practical performance [11, 24], and improved cache performance [4].

In this paper, we focus on filters that achieve space very close to the optimal lower bound of $n \log 1 / \epsilon$ bits [8, 16] and that store elements from a large universe $|U| \gg n$.

Fixing False Positives A well-known issue with many existing filters is that they cannot adapt to queries: if a query $q \notin S$ is a false positive, all subsequent queries $q^{\prime}=q$ will be false positives. The focus of this paper is designing filters that do adapt to false positive queries, so that if a query $q$ is a false positive, the filter undergoes structural changes so that a later query to $q$ is unlikely to be a false positive. An element $q$ is said to be fixed if $q$ was previously a false positive, but is no longer a false positive. Similarly, $q$ is broken if $q$ was previously fixed, but is now a false positive.

Related Work. Bender et al. [3] analyzed how to fix false positives against an adversary. They give a data structure such that if queries are generated by an adversary trying to maximize the false positive rate, each query to a filter is a false positive with probability at most $\epsilon$, even if the query element was queried before. This requirement essentially provides concentration bounds: over $n$ queries, their filter incurs $\epsilon n$ false positives, even if the queries are maliciously chosen based on previous false positives.

However, the benefit of adaptivity goes beyond resisting an adversary. As shown experimentally by Mitzenmacher et al. 17, adapting to queries can significantly decrease the number of false positives-in fact, if queries are repeated sufficiently frequently, the performance can be much better than $O(\epsilon n)$. In particular, network trace data consists of a structured sequence of queries - can we give a data structure that performs particularly well on this kind of data?

Most recently, Bender et al. [2] compared adaptivity to cache-based strategies, finding that adaptivity leads to significantly better practical performance.

Support Optimality Ideally, an adaptive filter would incur a false positive with probability $\epsilon$ for each new query, and incur no further queries asymptotically. Thus, every new false positive is fixed, and this fixing is unlikely to break previously-fixed false positives. In particular, let $q_{1}, \ldots q_{n}$ be a predetermined sequence of queries $\}^{1}$ to a filter $\mathcal{F}$, and let $Q=\bigcup_{i=1}^{n}\left\{q_{i}\right\}$ be the set of unique queries in the sequence. We say that $\mathcal{F}$ is support optimal if the expected number of false positives when querying $q_{1}, \ldots q_{n}$ is $\epsilon|Q|(1+o(1))$. In this paper we give a support-optimal filter up to additive polylogarithmic terms, and show that it significantly improves practical performance.

Avoiding Remote Memory Accesses Adaptive filters focus on the case where a filter is used to avoid remote memory accesses. Oftentimes, network data is stored on large, slow storage. In this case, if the filter is small enough to fit in faster storage, it can be used to rule out queries to elements not in the set. Thus, when the filter answers absent, a slow query to the large storage is avoided, greatly improving throughput.

This use case is important because it allows us to assume (costly) access to the original set of items being held in the filter. Bender et al. [3] gave a lower bound showing that this access is necessary to achieve adaptivity, and indeed this assumption was used both in their data structre as well as in the Adaptive Cuckoo Filter 17.

In particular, we assume that we can make $O(1)$ expected accesses to $S$ on each insert, as well as each time the filter answers present. This does not asymptotically increase the number of accesses to slow storage ${ }^{2}$ We

[^1]use this assumption to "cuckoo" elements on a false positive, rehashing them and swapping them to another slot. All previous adaptive cuckoo filters also require rehashing elements of $S$ on a false positive [3, 17.

### 1.1 Results

We discuss three data structures in this paper: two versions of the Adaptive Cuckoo Filter originally presented in 17 (which we call the $\boldsymbol{C y c l i c} \boldsymbol{A C F}$ and $\boldsymbol{S w a p p i n g} \boldsymbol{A C F}$ ), and a Cuckoo Filter augmented with a new method of achieving adaptivity, which we call the Cuckooing $\boldsymbol{A C F}$.

The first contribution of this paper is the Cuckooing ACF, a support-optimal filter which can be implemented using almost-trivial changes to current Cuckoo Filter implementations.

In Section 3, we analyze the Cuckooing ACF and prove that it is support optimal over any $n$ queries, up to additive polylogarithmic terms. This gives a significant performance improvement over previous filters even for large $Q$, and the difference becomes more dramatic for small $Q$. For example, for the case of a repeated single query $(|Q|=1)$, static filters incur $\epsilon n$ false positives in expectation, whereas we show that the Cuckooing ACF incurs $O\left(\log ^{4} n\right)$ expected false positives.

We show that despite their strong practical performance, the Cyclic ACF and Swapping ACF are not support optimal-even if there are a constant number of queries $(|Q|=O(1))$, they may incur $\Omega(n)$ false positives, whereas the Cuckooing ACF incurs at most $O\left(\log ^{4} n\right)$. Thus, from the standpoint of support optimality, cuckooing is a better method for achieving adaptivity.

In Section 4, we provide experimental results that show that the theory bears out in practice: the Cuckooing ACF attains a low false positive rate on network trace datasets, which contain many repeated queries. The performance is not only stronger than a vanilla Cuckoo Filter, but also improves upon the performance of a Cyclic ACF and a Swapping ACF of the same size. This shows that the Cuckooing ACF is effective at fixing false positives in a practical sense. These results also emphasize the benefit of a simple adaptive filter: not only is the resulting data structure easier to implement, the simplicity entails less space usage compared to previous Adaptive Cuckoo Filters, leading to a significant performance improvement.

Finally, in Section 5, we prove lower bounds that demonstrate that a broad family of filters cannot be adaptive in the adversarial sense of Bender et al. $\sqrt{3}]$; this includes the Cyclic ACF, the Swapping ACF, and the Cuckooing ACF. This lower bound motivates the concept of support optimality: a support optimal filter achieves strong performance on real datasets without achieving adversarial adaptivity. Our proof also gives insight into the structure of adaptive filters-specifically, it shows that a space-efficient filter must have variable-sized fingerprints in order to be adversarially adaptive.

### 1.2 Related Work

Mitzenmacher et al. [17] were the first to describe the notion of adaptivity. They show that Adaptive Cuckoo Filters can significantly improve the number of false positives incurred on real-world network trace datasets. We replicate their experiments in Section 4, and show that the Cuckooing ACF improves upon their performance. They also give theoretical bounds that guarantee that a small number of false positives are incurred by their data structures if queries are selected at random. Their work briefly mentions the swapping strategy used in this paper (in the Cuckooing ACF), but does not analyze it.

Bender et al. [3] give a filter (achieving optimal space and optimal query time) that is adaptive against an adversary. Their filter works by adding a constant number of bits to the filter on average for each false positive encountered. This data structure gives strong bounds, but dynamically changing the number of bits for each stored element is likely to be an obstacle to a practical implementation. Bender et al. also give a lower bound showing that access to an external dictionary storing $S$ is necessary to achieve adaptivity; this helps motivate such accesses in our paper.

The Bloomier filter of Chazelle et al. [9] and the Weighted Bloom Filter of Bruck et al. [7] also aim to reduce the false positive rate; however, both assume that frequent queries are known to the filter in advance.

## 2 Three Adaptive Cuckoo Filters

In this section describe a new kind of filter, the Cuckooing ACF. We then discuss the Cyclic ACF and the Swapping ACF, both originally introduced in 17 .

We include a symbol table for reference in Appendix A.

### 2.1 ACF Parameters and Internal State

We begin by defining a more general data structure which we call the adaptive cuckoo filter (ACF). As the name suggests, the Cyclic ACF, the Swapping ACF, and the Cuckooing ACF are adaptive cuckoo filters.

An ACF $\mathcal{F}$ has integer parameters $f, k, b>0$, an additional parameter $\gamma>1$, and supports storing $n$ elements from a universe $U$ with a false positive rate $\epsilon$. The internal representation of a filter $\mathcal{F}$ consists of $k$ hash tables, each of $N=\gamma n / b k$ bins ${ }^{3}$ where each bin consists of $b$ slots of $f$ bits; thus, the space usage of $\mathcal{F}$ is $N \cdot b \cdot f \cdot k$ bits. The parameter $\gamma$ determines how densely elements are packed, trading off between insert time and space; often $\gamma \approx 1.05$ is used.

The hash tables are accessed using $k+1$ hash functions: $k$ location hash functions $h_{1}^{\ell}, \ldots, h_{k}^{\ell}: U \rightarrow$ $\{0, \ldots, N-1\}$ that hash from $U$ to a bit string of length $\log N \|^{4}$ and a single fingerprint hash $h^{f}$ mapping each $x \in U$ to an $f$-bit fingerprint. The range and domain of $h^{f}$ depend on which ACF is used and may depend on the internal state of $\mathcal{F}$; we provide details below. Following previous results on filters $3,5,10,11,17$, this paper assumes free access to uniform random hash functions ${ }^{5}$

When a set $S$ is stored in $\mathcal{F}$, for each element $x_{i} \in S$, the fingerprint of $x_{i}$ is stored in one of the slots of bin $B\left(x_{i}\right)$ in the $\beta_{i}$ th hash table; this bin is defined using a location hash: $B\left(x_{i}\right)=h_{\beta_{i}}^{\ell}(x)$ for some integer $0 \leq \beta_{i}<k$. We say a slot $\sigma$ is occupied if the fingerprint of some $x_{i} \in S$ is stored in $\sigma$; otherwise $\sigma$ is empty. We call $\beta_{i}$ the hash index of $x_{i}$.

Since an ACF stores each element using a hash index, we can keep track of the internal state of a filter using the hash index of each element. Thus, we use $C=(C[1], C[2], \ldots, C[n])=\left(\beta_{1}, \ldots, \beta_{n}\right)$ to define the configuration of $\mathcal{F}$. This fully defines the internal representation of a Cuckooing ACF. The internal representation of a Cyclic ACF also depends on $s$ metadata bits stored for each element, and the internal representation of a Swapping ACF also depends on which slot within the bin is used to store each element.

Suppose $S$ is stored using hash indices $\beta_{1}, \ldots \beta_{n}$. Then query $q \notin S$ collides with an element $x_{i} \in S$ under configuration $C$ when $h_{\beta_{i}}^{\ell}\left(x_{i}\right)=h_{\beta_{i}}^{\ell}(q)$ and $q$ and $x_{i}$ have the same fingerprint.

### 2.2 Cuckoo Filter Operations

We begin by describing how inserts and queries work for an ACF. The Cuckooing ACF and Swapping ACF insert and query using these methods; the Cyclic ACF uses a generalization of these methods.

Insert. Suppose an element $x_{i}$ is inserted into a set $S$ of size $i-1$ currently stored with filter $\mathcal{F}$ in configuration $C$, where elements $S=x_{1}, \ldots x_{i-1}$ have hash indices $\beta_{1}, \ldots \beta_{i-1}$. Assume that $\mathcal{F}$ can store up to $n \geq i$ elements. The insertion algorithm finds a valid configuration $C^{\prime}$ of $\mathcal{F}$ on $S$ such that there exists a hash index $\beta_{i}^{\prime} \in\{0, \ldots k-1\}$ for which bin $h_{\beta_{i}^{\prime}}^{\ell}$ has an empty slot. This may involve updating the hash indices of other elements; for $1 \leq j<i$ let $\beta_{j}^{\prime}$ be the hash index of $x_{j}^{\prime}$ under $C^{\prime}$. We describe how to determine $C^{\prime}$ below.

If there is already an available empty slot, the filter stores the element immediately in that slot. Specifically, if there exists a $\beta \in\{0, \ldots, k-1\}$ where $\operatorname{bin} h_{\beta}^{\ell}\left(x_{i}\right)$ in hash table $\beta$ contains an empty slot, the filter sets $\beta_{i}^{\prime}=\beta$, and stores the fingerprint of $x_{i}$ in the empty slot. All other slots remain unchanged: $\beta_{j}^{\prime}=\beta_{j}$ for all $1 \leq j<i$.

[^2]Now, consider the case where there is no available empty slot. Then the ACF makes room by shifting elements as one would in cuckoo hashing [21]. The filter selects a hash index $\beta_{i}$ arbitrarily from $\{0, \ldots, k-1\}$. Since all slots in bin $h_{\beta_{i}^{\prime}}^{\ell}\left(x_{i}\right)$ are occupied in $C$, the filter moves the fingerprint of some element $x_{j}$ stored in a slot in $h_{\beta_{i}^{\prime}}^{\ell}\left(x_{j}\right)=h_{\beta_{i}^{\prime}}^{\ell}\left(x_{i}\right)$, leaving an empty slot in which $x_{i}$ can be stored. If $h_{\beta}^{\ell}\left(x_{j}\right)$ contains an empty slot for some $\beta \in\{0, \ldots, k-1\}$ (i.e. if $x_{j}$ can be stored in an empty slot), one such empty slot is arbitrarily selected to store $x_{j}$. Otherwise, the filter increments $\beta_{j}^{\prime}=\beta_{j}+1(\bmod k)$ and recurses, moving an element stored in $h_{\beta_{j}^{\prime}}^{\ell}\left(x_{j}\right)$ as necessary.

The move the elements as described above, the ACF must be able to access the set $S$ during an insert in order to rehash each $x_{j}$. We follow all past work on adaptive filters $[2,3,17]$ in assuming that an external dictionary can be accessed, enabling an element to be rehashed while inserting or fixing.

If this recursive process takes too many steps (more than $\Theta(\log n)$ elements are moved), the filter chooses new hashes and is rebuilt from scratch. If $\mathcal{F}$ uses $N=O(n)$ hash slots, then over $n$ inserts, the probability of a rebuild is $O(1 / n)$ 21.

There has been a great deal of previous work on improving the above cuckoo hashing procedure with various methods for finding the best slots to store each element, see i.e. $[1,12,13,15,19]$.

Notice that the procedure we describe does not use "partial-key" cuckoo hashing, as in the original cuckoo filter of Fan et al. [11]. The ability to access $S$ (necessary for adaptivity [3]) means that elements can be moved based on an entirely new hash.

Query. On a query $q$, a filter $\mathcal{F}$ in configuration $C$ returns present if there exists a $\beta$ and a slot index $\sigma \in\{1, \ldots b\}$ such that slot $\sigma$ in bin $h_{\beta}^{\ell}(q)$ of table $\beta$ is occupied and stores the fingerprint of $q$. This immediately guarantees correctness of the filter (queries to $x_{i} \in S$ always return present) and, via a union bound over the elements of $S$, a false positive rate of at most $n /\left(N 2^{f}\right)$. The filter achieves a desired false positive rate $\epsilon$ by setting $f=\log (n /(N \epsilon))=\log (b k / \epsilon \gamma)$.

Fixing false positives. If an ACF returns present on a false positive query $q$ (the filter knows that $q \notin S$ from the external dictionary storing $S$ ), the ACF modifies its configuration to attempt to fix $q$, so that subsequent queries to $q$ return absent. Each type of ACF has its own method for fixing false positives, which we describe below. Notice that the process of modifying the configuration may cause some query $q^{\prime} \notin S$ to become a false positive, even if $q^{\prime}$ was fixed some time in the past.

### 2.3 Cuckooing ACF

The primary data structure contribution of this paper is the $\boldsymbol{C u c k o o i n g} \boldsymbol{A C F}$. This data structure is a standard Cuckoo Filter 11 with an added operation to fix false positives; inserts and queries work exactly as described in Section 2.2

Let $q$ be a false positive under configuration $C$; we define how the Cuckooing ACF finds a new configuration $C^{\prime}$ with hash indices $\beta_{1}^{\prime}, \ldots \beta_{n}^{\prime}$ to attempt to fix $q$. For each $x_{i} \in S$ that collides with $q$ under $C$, the filter moves $x_{i}$ recursively as it would during an insert. Specifically, the filter sets the new hash index $\beta_{i}^{\prime}=\beta_{i}+1$ $(\bmod k)$; if bin $h_{\beta_{i}^{\prime}}^{\ell}\left(x_{i}\right)$ in table $\beta_{i}^{\prime}$ does not contain an empty slot, an element $x_{j}$ stored in $h_{\beta_{i}^{\prime}}^{\ell}\left(x_{i}\right)$ under $C$ is moved recursively. If more than $\Omega(\log n)$ steps are taken, the filter is rebuilt. Standard Cuckoo Hashing analysis shows that for any insert on a Cuckooing ACF with $\gamma=1+\Omega(1)$ a rebuild occurs with probability $O\left(1 / n^{2}\right) 21$.

### 2.4 Cyclic ACF

The cyclic $\boldsymbol{A C F}$ of Mitzenmacher et al. 17 is an ACF where each slot contains $s$ additional hash selector bits. The cyclic ACF generally has $b=1$; thus, the total space used by a Cyclic ACF is $k N(f+s)$. Usually, $s$ is a small constant.

In the Cyclic ACF, the fingerprint hash maps $U \times\left\{0, \ldots, 2^{s}-1\right\} \rightarrow\left\{0, \ldots, 2^{f}-1\right\}$. In particular, the hash selector bits are used to determine the fingerprint of an element stored in a given slot.

When an element $x_{i}$ is initially inserted, the insertion process continues as in Section 2.2 , with fingerprint $h^{f}\left(x_{i}, 0\right)$. When an empty slot $\sigma$ is found that can store $x_{i}$, the hash selector bits of $\sigma$ are set to 0 , and $h^{f}\left(x_{i}, 0\right)$ is stored in $\sigma$.

To query an element $q$, for each location hash $h_{\beta}^{\ell}$, with $\beta \in\{0, \ldots, k-1\}$, the filter looks at the slot $h_{\beta}^{\ell}$ of table $\beta$. The $s$ hash selector bits stored in the slot contain a value $0 \leq \alpha \leq 2^{s}-1$. The filter compares $h^{f}(q, \alpha)$ with the fingerprint stored in the slot; the filter returns present if they are equal. Otherwise the filter increments $\beta$ and repeats. If no collisions are found for all $0 \leq \beta \leq k-1$, the filter returns absent.

If a query $q$ is a false positive, the Cyclic ACF fixes the query as follows. Let $x_{i}$ be the element that collides with $q$, let $\sigma$ be the slot storing $x_{i}$, and let $\alpha$ be the value of the $s$ hash selector bits stored in $\sigma$. Then the filter sets the hash selector bits of $\sigma$ to store value $\alpha+1$, and stores $h^{f}\left(x_{i}, \alpha+1\right)$ in $\sigma$. If multiple $x \in S$ collide with $q$, this procedure is repeated for each such $x$.

### 2.5 Swapping ACF

The idea of the Swapping ACF (of Mitzenmacher et al. 17 ) is to have elements hash to bins with $b>1$ slots, and to have the fingerprint of an item depend on the slot it is stored in. In this way, elements can be (potentially) fixed by moving them to a different slot.

Inserts proceed as described in Section 2.2. However, in the Swapping ACF, the fingerprint hash maps $U \times\{0, \ldots, b-1\} \rightarrow\left\{0, \ldots, 2^{f}-1\right\}$. During an insert, the slot storing an element must be determined before its fingerprint can be calculated.

If a query $q$ is a false positive under configuration $C$, the filter can fix the query as follows. Let $x_{i}$ be the element that collides with $q$ and let $b\left(x_{i}\right)=h_{\beta_{i}}^{\ell}\left(x_{i}\right)$ be the bin currently storing $x_{i}$. Let $\sigma_{i} \in\{0, \ldots, b-1\}$ be the index of the slot in $b\left(x_{i}\right)$ currently storing $x_{i}$; thus $x_{i}$ is stored in slot $h_{\beta_{i}}^{\ell}\left(x_{i}\right) \cdot b+\sigma_{x}$.

The filter picks a slot index $\sigma^{\prime} \in\left\{0, \ldots, \sigma_{i}-1, \sigma_{i}+1, \ldots b-1\right\}$, selected at random from the slots in $b\left(x_{i}\right)$, excluding the slot currently storing $x_{i}$. Let $x_{j}$ be the element currently stored in that slot if it exists. The filter then swaps the elements: it stores fingerprint $h^{f}\left(x_{i}, \sigma^{\prime}\right)$ in slot $h_{\beta_{i}}^{\ell}\left(x_{i}\right) \cdot b+\sigma^{\prime}$, and fingerprint $h^{f}\left(x_{j}, \sigma_{i}\right)$ in slot $h_{\beta_{i}}^{\ell}\left(x_{i}\right) \cdot b+\sigma_{i}$ (if $x_{j}$ does not exist, $\sigma_{i}$ becomes unoccupied).

## 3 Bounding the False Positive Rate by the Number of Distinct Queries

In this section we show that the Cuckooing ACF is support optimal: it achieves strong performance against skewed datasets, where the queries are taken from a relatively small set of elements.

Our analysis focuses on a Cuckooing ACF with $k=2$ hash tables, $b=1$ slots per bin, and $N=n$ slots per hash tabl ${ }^{6}$ (corresponding to the classic Cuckoo Hashing analysis). The experiments in Section 4 indicate that our analysis likely extends to broader parameter ranges. However, formally completing the analysis for all parameters would require significant new structural insights in our proofs (e.g. Lemma 2); we leave this to future work.

Theorem 1. Consider a sequence of at most n queries $q_{1}, \ldots q_{n}$ to a Cuckooing ACF $\mathcal{F}$ with $k=2$ hash tables, $N=n$ slots per table, and fingerprints of length $f=\log 1 / \epsilon$ bits. Let $Q=\bigcup_{i=1}^{n}\left\{q_{i}\right\}$. Then the expected number of false positives incurred by $\mathcal{F}$ while querying $q_{1}, \ldots, q_{n}$ is $\epsilon|Q|+O\left(\epsilon^{2}|Q|+\log ^{4} n\right)$.

Thus, for any sequence of $n$ queries with a support of size $|Q|=\omega\left(\log ^{4} n / \epsilon\right)$, the Cuckooing ACF is support optimal.

In contrast, for a worst-case input sequence, the Cyclic ACF and the Swapping ACF do not perform much better than a Cuckoo Filter. Taking the Cyclic ACF as an example, consider a sequence of $n$ queries, each chosen uniformly at random from a randomly-selected set of size $|Q|=1 / \epsilon^{2^{s}}$. Each of these queries collides with some $x \in S$ under every choice of hash selector bits with probability $\Omega\left(\epsilon^{2^{s}}\right)$. Thus, over $n$ queries, the

[^3]Cyclic ACF incurs $\Omega(n)$ false positives for constant $\epsilon$ and $s$, compared to $O\left(\log ^{4} n\right)$ false positives for the Cuckooing ACF via Theorem 1. See the proof of Theorem 9 for a more detailed explanation.

### 3.1 Proof of Theorem 1

Without loss of generality we assume that each false positive query only collides in one of the hash tables. Since $k=2$, fixing a query that collides in both hash tables can be simulated by executing the fixing function for each hash table separately.

To simplify notation, we define $B(i, C)=h_{C[i]}^{\ell}\left(x_{i}\right)$ to be the bin storing $x_{i}$ under configuration $C$, and $B^{\prime}(i, C)=h_{1-C[i]}^{\ell}\left(x_{i}\right)$ to be the alternate bin that can store $x_{i}$. Because $b=1$ for this analysis, we can refer to "slot" $B(i, C)$ and "bin" $B(i, C)$ equivalently.

Let $C_{0}$ be the configuration of $\mathcal{F}$ before the first query $q_{1}$, and for $1 \leq i \leq n$ let $C_{i}$ be the configuration after query $q_{i}$. For each $1 \leq i \leq n$, if $q_{i}$ is a false positive under $C_{i-1}$, let $k_{i}$ be the number of elements moved when fixing query $q_{i}$; otherwise let $k_{i}=0$. We denote the sequence of elements moved when fixing $q_{i}$ as $x_{i_{1}}, x_{i_{2}}, \ldots, x_{i_{k_{i}}}$. Thus, $q_{i}$ collides with $x_{i_{1}}$ under $C_{i-1}$. We call the sequence of slots affected by these movements $B\left(i_{1}, C_{i-1}\right), B\left(i_{2}, C_{i-1}\right), \ldots B\left(i_{k_{i}}, C_{i-1}\right), B^{\prime}\left(i_{k_{i}}, C_{i-1}\right)$ the path on $C_{i-1}$ of $q_{i} \|^{7}$

We say that $q_{i}$ loops if one of the moved elements repeats; i.e. there exist $1 \leq \ell_{1}<\ell_{2} \leq k_{i}$ such that $i_{\ell_{1}}=i_{\ell_{2}}$. Interestingly, classic Cuckoo Hashing analysis generally only needs to bound the number of queries that loop twice, as only twice-looping queries force a rebuild. However, even a query that loops once cannot be fixed in a Cuckooing ACF, so we must bound how frequently this happens in our analysis.

Let the initial false positives be the queries in $Q$ that are false positives for $\mathcal{F}$ in configuration $C_{0}$. For a fixed $C, \mathcal{F}$, and $Q$, we denote the set of initial false positives as

$$
F_{0}(C, \mathcal{F}, Q)=\left\{q \in Q \mid q \text { is a false positive for } \mathcal{F} \text { under } C_{0}\right\}
$$

Fixing $C, \mathcal{F}$, and $Q$, we refer to $F_{0}:=F_{0}(C, \mathcal{F}, Q)$.
We start with a structural lemma: the elements moved when fixing any query consist of a (possibly empty) sequence of elements stored in the slot they occupied in $C_{0}$, followed by a (possibly empty) sequence of elements not stored in the slot they occupied in $C_{0}$.

Lemma 2. If a query $q_{i}$ on a configuration $C_{i-1}$ moves an element $x_{i_{\ell}}$ satisfying $C_{i-1}\left[i_{\ell}\right] \neq C_{0}\left[i_{\ell}\right]$, and $q_{i}$ does not loop, then all $j$ with $\ell \leq j \leq k_{i}$ satisfy $C_{i-1}\left[i_{j}\right] \neq C_{0}\left[i_{j}\right]$.

Proof. This proof is by induction on $j$; the base case $j=\ell$ is satisfied by assumption.
Assume by induction that $C_{i-1}\left[i_{j-1}\right] \neq C_{0}\left[i_{j-1}\right]$ for some $j>\ell$. Since $q_{i}$ does not loop, when $x_{i_{j-1}}$ is moved, it cannot have been moved previously while fixing $q_{i}$, and thus must be stored in slot $B\left(i_{j-1}, C_{i-1}\right)$. Then after $x_{i_{j-1}}$ is moved it must be stored in slot $B\left(i_{j-1}, C_{0}\right)$; this must be equal to the slot storing $x_{i_{j}}$. Because $q_{i}$ does not loop, $x_{i_{j}}$ must be stored where it was when the fixing began; i.e. in $B\left(i_{j}, C_{i-1}\right)$. Thus $B\left(i_{j}, C_{i-1}\right)=B\left(i_{j-1}, C_{0}\right)$, so $C_{0}\left[i_{j}\right] \neq C_{i-1}\left[i_{j}\right]$, as otherwise $x_{i_{j}}$ and $x_{i_{j-1}}$ would be stored in the same slot in $C_{0}$.

Lemma 2 immediately gives structure to the problem in two key ways. First, it limits how queries can break one another: if $q$ is a false positive, but is not an initial false positive, then there must be some initial false positive $q_{i}$ that caused $q$ to become a false positive. We do not need to worry about non-initial false positives causing other, new false positives. Second, it ties the behavior of all elements to how they behave on the initial configuration $C_{0}$. This means that we can make statements about how queries interact using $C_{0}$; we do not need to reset our analysis every time the filter configuration changes.

Our analysis depends heavily on how queries behave on configuration $C_{0}$, so we introduce notation to help discuss queries on $C_{0}$. Let $k_{i}^{0}$ be the number of elements moved when querying $q_{i}$ on configuration $C_{0}$. Let $x_{i_{1}^{\prime}}, x_{i_{2}^{\prime}}, \ldots x_{i_{k_{i}^{\prime}}^{\prime}}$ be the set of elements moved when fixing $q_{i}$, if $q_{i}$ were queried on configuration $C_{0}$.

[^4]We now define the notion of costly queries. In short, costly queries are either difficult to fix, or break other queries. However, costly queries will prove to be rare, and will not substantially increase the total number of false positives.

A query $q_{i}$ is a costly query if $q_{i}$ is a false positive on $C_{i-1}$ and $q_{i}$ meets one of the following criteria. We refer to these as Criteria 1-4 in our analysis.

1. If $q_{i}$ is an initial false positive, then there exists a $q_{j}$ such that the path of $q_{j}$ when queried on $C_{0}$ collides with the path of $q_{i}$ when queried on $C_{0}$; if $q_{i}$ is not an initial false positive, then there exists a $q_{j}$ such that the path of $q_{j}$ when queried on $C_{0}$ is hashed to by $q_{i}$. In other words, if $q_{i} \in F_{0}$ there exists a $q_{j} \in F_{0} \backslash\left\{q_{i}\right\}$, and a pair of indices $\ell_{j}, \ell_{i}$ with $1 \leq \ell_{j} \leq k_{j}^{0}$ and $1 \leq \ell_{i} \leq k_{i}^{0}$ and a hash index $\beta \in\{0,1\}$ where $h_{\beta}^{\ell}\left(x_{j_{\ell_{j}}^{\prime}}\right)=h_{\beta}^{\ell}\left(x_{i_{\ell_{i}}^{\prime}}\right)$; if $q_{i} \notin F_{0}$ then there exists a $q_{j} \in F_{0} \backslash\left\{q_{i}\right\}$ and an index $\ell_{j}$ with $h_{C_{0}\left[\ell_{j}\right]}^{\ell}\left(q_{i}\right)=h_{C_{0}\left[\ell_{j}\right]}^{\ell}\left(x_{j_{\ell_{j}}^{\prime}}\right)$.
2. $q_{i}$ loops when queried on $C_{i-1}$.
3. There exists a false positive $q_{j} \in Q$ that loops on $C_{j-1}$ and whose path contains a slot that $q_{i}$ hashes to. In other words, there exists a $q_{j} \in Q$ that loops on $C_{j-1}$, an integer $1 \leq \ell \leq k_{j}$, and a hash index $\beta \in\{0,1\}$ such that $B\left(j_{\ell}, C_{j-1}\right)=h_{\beta}^{\ell}\left(q_{i}\right)$ or $B^{\prime}\left(j_{\ell}, C_{j-1}\right)=h_{\beta}^{\ell}\left(q_{i}\right)$.
4. $q_{i}$ hashes to a slot along its own path (excluding the slot storing the first moved element): thus, there exists a $1 \leq \ell \leq k_{i}$ such that $B^{\prime}\left(i_{\ell}, C_{i-1}\right)=h_{1-C_{i-1}\left[i_{\ell}\right]}^{\ell}\left(q_{i}\right)$.
To begin, we show that all all false positives queries are either an initial false positive, or a costly query.
Lemma 3. Let $q_{i}$ be a false positive on $C_{i-1}$ for filter $\mathcal{F}$. Then $q_{i} \in F_{0}$ or $q_{i}$ is costly.
Proof. Assume that $q_{i}$ is a false positive with $q_{i} \notin F_{0}$; we show that $q_{i}$ is costly. Let $1 \leq j<i$ be the smallest $j$ such that $q_{i}$ is a false positive under $C_{j}\left(j\right.$ must exist because $q_{i}$ is a false positive under $C_{i-1}$, but $q_{i}$ is not an initial false positive). If $q_{j}$ loops, then $q_{i}$ is costly by Criterion 3 ,

Now, assume $q_{j}$ does not loop. There must be some $x_{j_{\ell}}$ moved by $q_{j}$ (when queried on $C_{j-1}$ ) such that $q_{i}$ collides with $x_{j_{\ell}}$ in $C_{j}$. Since $q_{i}$ is not an initial false positive and $q_{j}$ does not loop, we must have that $q_{i}$ collides with $x_{j_{\ell}}$ when $x_{j_{\ell}}$ is stored in slot $B^{\prime}\left(j_{\ell}, C_{j-1}\right)=B^{\prime}\left(j_{\ell}, C_{0}\right)$. However, this means that $C_{j-1}\left[j_{\ell}\right]=C_{0}\left[j_{\ell}\right]$, and therefore $C_{j-1}\left[j_{1}\right]=C_{0}\left[j_{1}\right]$ as otherwise we reach a contradiction with Lemma 2 . This means that $q_{j}$ is an initial false positive, and thus $q_{i}$ is costly by Criterion 1 .

We bound the cost of all queries using a potential function argument. The potential of a configuration is the number of elements $x \in S$ stored in their original position that collide with an initial false positive. Specifically, for all $0 \leq t \leq n$, define

$$
\Phi(t)=\mid\left\{x_{i} \in S, q \in F_{0} \mid x_{i} \text { collides with } q \text { under } C_{t} \text { and } C_{0}[i]=C_{t}[i]\right\} \mid
$$

Let the amortized cost of query $q_{i}$ be $1+\Phi(i)-\Phi(i-1)$ if $q_{i}$ is a false positive, and $\Phi(i)-\Phi(i-1)$ if $q_{i}$ is not a false positive. Summing (and taking $\Phi(n) \geq 0$ ), the expected number of false positives incurred during the queries on $q_{1}, \ldots, q_{n}$ is at most the expected number of amortized false positives plus $\mathbb{E}[\Phi(0)]$. The proof proceeds in three parts: first, bounding $\mathbb{E}[\Phi(0)]$, then bounding the cost of queries that are not costly, and finally bounding the cost of costly queries.

The expected potential before any queries follows immediately from linearity of expectation and $f=$ $\log (1 / \epsilon)$.

Lemma 4. $\mathbb{E}[\Phi(0)] \leq \epsilon|Q|$
Proof. For a given $x_{i} \in S$ and $q \in Q$, the probability that $h_{C_{0}[i]}\left(x_{i}\right)=h_{C_{0}[i]}(q)$ is at most $\epsilon / n$. Summing over the $|Q| n$ pairs obtains $\mathbb{E}[\Phi(0)] \leq \epsilon|Q|$.

Now we show that non-costly queries have no amortized cost. This is immediate for any $q_{i}$ that is not a false positive. For a non-costly false positive $q_{i}$, we show that this cost is offset by a decrease in potential from moving $x_{i_{1}}$.
Lemma 5. If a query $q_{i}$ is not costly, then the amortized cost of $q_{i}$ is at most 0 .
Proof. If $q_{i}$ is not a false positive, then no elements are moved, and $\Phi(i)=\Phi(i-1)$. Thus the amortized cost of a query that is not a false positive is 0 .

If $q_{i}$ is a false positive but is not costly, then $q_{i}$ is an initial false positive by Lemma furthermore $q_{i}$ does not loop by Criterion 2. Since $q_{i}$ is an initial false positive, moving $x_{i_{1}}$ decreases $\Phi$ by 1 ( $q_{i}$ does not collide with $x_{i_{1}}$ in its new position by Criterion 4 and $x_{i_{1}}$ is not moved again while fixing $q_{i}$ because $q_{i}$ does not loop by Criterion 2).

We now show that this is the only change to $\Phi$. Assume the contrary: there is a $q_{j} \in F_{0}$ and an $\ell \leq k_{i}$ such that $x_{i_{\ell}}$ collides with $q_{j}$ under $C_{i}$, and $C_{i}\left[i_{\ell}\right]=C_{0}\left[i_{\ell}\right]$. By Lemma 2 since $q_{i}$ does not loop we can partition the elements moved by $q_{i}$ into at most two sequences of elements: the possibly-empty sequence $x_{i_{1}}, \ldots x_{i_{\lambda}}$, and the sequence $x_{i_{\lambda+1}}, \ldots x_{i_{k_{i}}}$ where for all $\kappa \in\{1, \ldots, \lambda\}, C_{i}\left[i_{\kappa}\right] \neq C_{0}\left[i_{\kappa}\right]$, and for all $\kappa \in\left\{\lambda+1, \ldots, k_{i}\right\}$, $C_{i}\left[i_{\kappa}\right]=C_{0}\left[i_{\kappa}\right]$. Thus, $\ell>\lambda$. Recall that $j_{1}^{\prime}, \ldots j_{k_{j}^{0}}^{\prime}$ is the sequence of elements moved when querying $q_{j}$ on $C_{0}$. The sequence $i_{\ell}, i_{\ell-1}, \ldots, i_{\lambda+1}$ (notice that $\ell, \ell-1 \ldots, \lambda+1$ is in decreasing order) must be a subsequence of $j_{1}^{\prime}, \ldots j_{k_{j}^{0}}^{\prime}$. Similarly, if $i_{1}, \ldots, i_{\lambda}$ exists it must be a subsequence of $i_{1}^{\prime}, \ldots, i_{\lambda}^{\prime}$. If $q_{i}=q_{j}$, then $q_{i}$ must be costly by Criterion 4. If $q_{i} \neq q_{j}$, then $q_{i}$ must be costly by Criterion 11, since if $x_{i_{\lambda}}$ exists then $h_{C_{0}\left[i_{\lambda}\right]}^{\ell}\left(x_{i_{\lambda}}\right)=h_{1-C_{0}\left[i_{\lambda+1}\right]}^{\ell}\left(x_{i_{\lambda+1}}\right)$, and if $x_{i_{\lambda}}$ does not exist then $h_{1-C_{0}\left[i_{\lambda+1}\right]}^{\ell}\left(q_{i}\right)=h_{1-C_{0}\left[i_{\lambda+1}\right]}^{\ell}\left(x_{i_{\lambda+1}}\right)$. Thus $\Phi(i)-\Phi(i-1) \leq-1$, and the amortized cost of $q_{i}$ is at most 0 .

We begin analyzing costly queries by bounding their cost in terms of their path length. Intuitively, this bound reflects that each of the $k_{i}$ elements moved when fixing $q_{i}$ collides with an element of $Q$ with probability $\epsilon / n$.
Lemma 6. If $q_{i}$ is costly, then its expected amortized cost is at most $4+2 \epsilon|Q| k_{i} / n$.
Proof. We upper bound the difference in $\Phi$ using a very pessimistic case: if some moved element $x_{i_{\ell}}$ collides with some $q^{\prime} \in F_{0}$ under some hash $h_{\beta}$, then we assume that $x_{i_{\ell}}$ collides with $q^{\prime}$ under $C_{i}$ but not $C_{i-1}$. This leads to the bound

$$
\mathbb{E}[\Phi(i)-\Phi(i-1)] \leq \mathbb{E}\left[\sum_{\ell=1}^{k_{i}} \mid\left\{q \in F_{0} \mid h_{\beta}(q)=h_{\beta}\left(x_{i_{\ell}}\right) \text { for some } \beta \in\{0,1\}\right\}| | q_{i} \text { is costly }\right]
$$

If the path of $q_{i}$ collides with the path of $q_{j}$ (under Criterion 3 ), then $x_{j_{1}}$ may collide with $q_{j}$ under $C_{i}$ and $C_{0}$; the same holds if the path of $q_{i}$ collides with the path of $q_{j}$ when queried on $C_{0}$ (under Criterion 1). If $q_{i}$ hashes to a slot along its path, it may collide with the moved element under both $C_{0}$ and $C_{i}$. If $q_{i}$ loops (under Criterion 2), then $q_{i}$ collides with $x_{i_{1}}$ under $C_{i-1}$ and may collide with $x_{i_{1}}$ under $C_{i}$ and $C_{0}$; however this does not increase $\Phi$. We sum to obtain a (loose) upper bound of 3 collisions from the cases listed in this paragraph.

For any other pair $\left(q^{\prime}, \lambda\right)$ where $q^{\prime} \in F_{0}$ and $\lambda \in\left\{1, \ldots k_{i}\right\}, q^{\prime}$ and $x_{i_{\lambda}}$ collide under $\beta$ with probability $\epsilon / n$ for each hash index $\beta \in\{0,1\}$. Summing via union bound over all $q^{\prime}, \lambda, \beta$ obtains

$$
\mathbb{E}\left[\sum_{\ell=1}^{k_{i}} \mid\left\{q \in F_{0} \mid h_{\beta}(q)=h_{\beta}\left(x_{i_{\ell}}\right) \text { for some } \beta \in\{0,1\}\right\}| | q_{i} \text { is costly }\right] \leq 3+\left(2 k_{i}|Q|\right)(\epsilon / n)
$$

Adding in the cost of the false positive $q_{i}$, the expected amortized cost of $q_{i}$ is at most $4+2 \epsilon|Q| k_{i} / n$.
The following lemma is dedicating to bounding the increase in potential from costly queries that satisfy Criterion 1. These are by far the most common and expensive costly queries, so our analysis of this case needs to be tight up to (essentially) constants to obtain our desired bounds. That said, we did not focus on minimizing the constants themselves.

Lemma 7. Assume that all initial false positives $q_{i} \in F_{0}$ satisfy $k_{i}^{0}=O(\log n)$. Then the expected number of queries satisfying Criterion 1 is at most $8\left|F_{0}\right|^{2} / n+2 \epsilon\left|Q \| F_{0}\right| / n$. Furthermore,

$$
\mathbb{E}\left[k_{i} \mid q_{i} \text { satisfies Criterion } 1 \leq 96\right.
$$

Proof. We begin by bounding the number of queries $q_{i} \in F_{0}$ that satisfy Criterion 1 . Taking a union bound, for any $q_{j} \in F_{0}$, the probability that an element moved when querying $q_{i}$ shares a slot with an element moved when querying $q_{j}$ is $k_{i}^{0} k_{j}^{0} / n$. So long as $q_{i}$ and $q_{j}$ are moving distinct elements when being fixed under $C_{0}$, the hashes of each element are independent, and therefore the number of elements each moves in independent $\square^{8}$ If $q_{i}$ and $q_{j}$ ever move the same element, they satisfy Criterion 1 and we are done. Thus, we can upper bound the expected number of pairs $q_{i}, q_{j} \in F_{0}$ satisfying Criterion 1 by $\sum_{q_{i}, q_{j} \in F_{0}} \mathbb{E}\left[k_{i}^{0} k_{j}^{0}\right] / n$.

From [21, Section 3.1.2] we have (in the notation of this paper, and recalling that the filter may access fully random (i.e. $(1, n)$-universal) hash functions, and recalling that we assume that $q_{i}$ and $q_{j}$ collide with at least one element), $\operatorname{Pr}\left[k_{i}^{0}=\hat{k}\right] \leq 2(1 / \gamma)^{-\hat{k}-1}=2^{2-\hat{k}}$ for $\hat{k} \geq 1$. As above, we can assume that $k_{i}^{0}$ and $k_{j}^{0}$ are independent for the purpose of our upper bound. Summing,

$$
\begin{aligned}
\mathbb{E}\left[k_{i}^{0} k_{j}^{0}\right] & \leq \sum_{k_{i}^{0}=1}^{\infty} \sum_{k_{j}^{0}=1}^{\infty}\left(k_{i}^{0}\right)\left(k_{j}^{0}\right) 2^{2-k_{i}^{0}} 2^{2-k_{j}^{0}} \\
& =16 \sum_{k_{i}^{0}=1}^{\infty}\left(k_{i}^{0}\right) 2^{-k_{i}^{0}}\left(\sum_{k_{j}^{0}=1}^{\infty}\left(k_{j}^{0}\right) 2^{-k_{j}^{0}}\right) \\
& \leq 32 \sum_{k_{i}^{0}=1}^{\infty}\left(k_{i}^{0}\right) 2^{-k_{i}^{0}} \leq 64
\end{aligned}
$$

Thus, the expected number of pairs that satisfy Criterion 1 is at most $64\left|F_{0}\right|^{2} / n$ for $n \geq 64$.
We now consider the case when $q_{i} \notin F_{0}$. This means that $q_{i}$ must hash to the same slot as an element moved by some $q_{j}$ when $q_{j}$ is being fixed on $C_{0}$; for a given element this occurs with probability $\epsilon / n$. This means that the probability that $q_{i}$ satisfies Criterion 1 is at most $\sum_{q_{j} \in F_{0}} \epsilon k_{j}^{0} / n$. Summing obtains $2 \epsilon\left|Q \| F_{0}\right| / n$ expected false positives over all $q_{i} \in Q$.

Finally, we bound $\mathbb{E}\left[k_{i} \mid q_{i}\right.$ satisfies Criterion 11. By definition,

$$
\begin{aligned}
\mathbb{E}\left[k_{i} \mid q_{i} \text { satisfies Criterion } 1\right] & =\sum_{k^{\prime}=1}^{\infty} k^{\prime} \operatorname{Pr}\left[k_{i}=k^{\prime} \mid q_{i} \text { satisfies Criterion } 1\right] \\
& =\frac{\operatorname{Pr}\left[k_{i}=k^{\prime} \text { and } q_{i} \text { satisfies Criterion } 1\right.}{\operatorname{Pr}\left[q_{i} \text { satisfies Criterion } 1\right.}
\end{aligned}
$$

As above,

$$
\begin{aligned}
\operatorname{Pr}\left[k_{i}=k^{\prime} \text { and } q_{i} \text { satisfies Criterion } 1\right] & \leq k^{\prime} 2^{2-k^{\prime}} \sum_{q_{j} \in F_{0}} \sum_{k_{j}^{0}=0}^{\infty} k_{j}^{0} 2^{2-k_{j}^{0}} / n \\
& \leq 4 k^{\prime} 2^{-k^{\prime}}\left|F_{0}\right| \sum_{k_{j}^{0}=0}^{\infty} k_{j}^{0} 2^{-k_{j}^{0}} / n \\
& \leq 8\left|F_{0}\right| k^{\prime} 2^{-k^{\prime}} / n
\end{aligned}
$$

[^5]A lower bound for $\operatorname{Pr}\left[q_{i}\right.$ satisfies Criterion [1] is the event that $h_{1-C\left[j_{1}\right]}^{\ell}\left(q_{i}\right)=h_{1-C\left[j_{1}\right]}^{\ell}\left(x_{j_{1}}\right)$. This occurs with probability $1-(1-1 / n)^{\left|F_{0}\right|} \geq 1-1 / 2^{\left|F_{0}\right| / n}$. Thus,

$$
\begin{aligned}
\mathbb{E}\left[k_{i} \mid q_{i} \text { satisfies Criterion } 1\right] & =\sum_{k^{\prime}=1}^{\infty} k^{\prime} \operatorname{Pr}\left[k_{i}=k^{\prime} \mid q_{i} \text { satisfies Criterion } 1\right] \\
& \leq \sum_{k^{\prime}=1}^{\infty} k^{\prime} \frac{8\left|F_{0}\right| k^{\prime} 2^{-k^{\prime}} \mid n}{1-1 / 2^{\left|F_{0}\right| / n}} \\
& \leq 8 \frac{\left|F_{0}\right| / n}{1-1 / 2^{\left|F_{0}\right| / n}} \sum_{k^{\prime}=1}^{\infty} k^{\prime 2} 2^{-k^{\prime}} \\
& \leq 16 \sum_{k^{\prime}=1}^{\infty} k^{\prime 2} 2^{-k^{\prime}} \leq 96
\end{aligned}
$$

Now we bound the total cost of all costly queries. The additive polylog term is due to applying tail bounds to rare but problematic events like looping queries - it is possible that a sufficiently involved analysis of these events would reduce or remove this term.

Lemma 8. For any sequence of queries $q_{1}, \ldots, q_{n}$

$$
\mathbb{E}\left[\sum_{\substack{q_{i} i s \text { costly } \\ 1 \leq i \leq n}}\left(4+2 \epsilon|Q| k_{i} / n\right)\right] \leq O\left(\epsilon^{2}|Q|^{2} / n+\log ^{4} n\right) .
$$

Proof. From classic Cuckoo Hash analysis, $k_{i}=O(\log n)$ for all $i$ with probability at least $1-1 / n$ (see i.e. 21]). We upper bound the number of costly $q_{i}$ by summing the expected number of $q_{i}$ meeting each criterion. Assume that $k_{i}=O(\log n)$ and $k_{i}^{0}=O(\log n)$ for all $i$ (if not, then at most all $n$ queries are false positives; since this edge case happens with probability $\leq 2 / n$ this adds 2 to the expected cost).

Criterion 1; From Lemma 7, the expected number of queries satisfying Criterion 1 is at most $8\left|F_{0}\right|^{2} / n+$ $2 \epsilon\left|Q \| F_{0}\right| / n$. Substituting $\mathbb{E}\left[k_{i} \mid q_{i}\right.$ satisfies Criterion $1 \leq 96$ in Lemma 6 , each of these queries costs at most $(3+192 \epsilon|Q| / n)=O(1)$; thus the total cost is $O\left(\left|F_{0}\right|^{2} / n+\epsilon|Q|\left|F_{0}\right| / n\right)$.

To bound $E\left[\left|F_{0}\right|^{2}\right]$, notice that this expression represents the square of the sum of $|Q|$ independent Bernoulli trials, each of which succeeds with probability $\epsilon$. Thus, the variance of $\left|F_{0}\right|$ is $\epsilon|Q|(1-\epsilon)$; substituting $\mathbb{E}\left[\left|F_{0}\right|\right]=\epsilon|Q|$ into $\operatorname{Var}(X)=\mathbb{E}\left[X^{2}\right]-(\mathbb{E}[X])^{2}$ obtains $\mathbb{E}\left[\left|F_{0}\right|^{2}\right] \leq \epsilon^{2}|Q|^{2}+\epsilon(1-\epsilon)|Q|$. By linearity of expectation, $\mathbb{E}\left[|Q|\left|F_{0}\right| / n\right] \leq \epsilon|Q|^{2} / n$. Summing obtains a total cost of $O\left(\epsilon^{2}|Q|^{2} / n\right)$.

Criterion 2. Query $q_{i}$ loops if and only if there exists an element $x_{i_{\ell}}$ (with $1 \leq \ell \leq k_{i}$ ) such that the second hash location of $x_{i \ell}$ collides with another one of the $k_{i}$ slots containing elements moved by $q_{i}$. For a given slot $\sigma$, the probability that $h_{C_{i}\left[i_{\ell}\right]}^{\ell}\left(x_{i_{\ell}}\right)=\sigma$ is $1 / n$. Taking a union bound, $q_{i}$ loops with probability at most $k_{i}^{2} / n$. Summing over all $2 n$ queries, the expected number of queries that loop is $O\left(\log ^{2} n\right)$. Since $k_{i}=O(\log n)$, Lemma 6 gives that each costs $O(\log n)$, giving a total cost of $O\left(\log ^{3} n\right)$.

Criterion 3: If there are $L$ queries that loop, since each looping query moves $O(\log n)$ elements, the probability that $q_{i}$ collides with any element moved by a looping query is $O(L \log n / n)$. By linearity of expectation, since $\mathbb{E}[L]=O\left((\log n)^{2}\right)$ (shown in the proof for Criterion 22 , summing over all $q_{i}$, at most $O\left(\log ^{3} n\right)$ queries meet Criterion 3 in expectation. Using Lemma 6 obtains a total cost of $O\left(\log ^{4} n\right)$.

Criterion 4 Each element $x_{i_{\ell}}$ moved when querying $q_{i}$, from $1 \leq \ell \leq k_{i}$, hashes to the same slot as $q_{i}$ with probability $1 / n$; thus $q_{i}$ collides with an element after it is moved with probability $O(\log n) / n$. Summing over all queries, there are $O(\log n)$ queries that are costly in expectation by Criterion 4 using Lemma 6 obtains a total cost of $O\left(\log ^{2} n\right)$.

Summing, all costly queries have a total cost of $O\left(\epsilon|Q|^{2} / n+\log ^{4} n\right)$.
Summing the starting potential cost (Lemma 4), the total amortized cost of all queries that are not costly (Lemma 5), and the total amortized cost of all costly queries (Lemma 8), we obtain Theorem 1.

## 4 Experiments

In this section, we examine how the Cuckooing ACF performs on network trace datasets. There are two main takeaways from this section. First, the design of the Cuckooing ACF results in better practical performance than previous adaptive filters on network trace datasets. Second, the analysis of Section 3 extends to practice: an adaptive cuckoo filter with practical parameter settings (including very high load factor) still achieves strong performance.


Figure 1: We examine the false positive rate of each adaptive filter, varying the ratio of the number of queries to the number of stored elements. The bottom right figure normalizes the false positive rate by the number of false positives incurred by the Cuckooing ACF. It summarizes the results for the Swapping ACF and the Cyclic ACF with $s=1$, for all three datasets, for $f=8,12,16$.

Our experiments use three network traces from the CAIDA 2014 dataset, specifically:

- equinix-chicago.dirA. 20140619 (which we call "Chicago A")
- equinix-chicago.dirB.20140619-432600 ("Chicago B"), and
- equinix-sanjose.dirA.20140320-130400 ("San Jose")
following the experiments of Mitzenmacher et al. in 17 . In each test, the elements stored in the filter and the query elements both come from the network trace. Let $A$ be the set of query elements. We perform tests for different $|A| /|S|$ ratios; specifically $|A| /|S|=\{1,3,5,10,20,30,40,50,60,70,80,90,100\}$.

We begin by setting $n=|S|$ using the prescribed $|A| /|S|$ ratio and the total number of unique flows in the trace - in particular, $n=(\#$ unique flows $) /(1+|A| /|S|)$. The first $n$ unique flows seen in the trace are inserted into each filter. The remaining flows in the trace (those not in $S$ ) are used as queries.

We consider six data structures in our experiments:

- a classic Cuckoo Filter, with $k=4$ hash tables and $b=1$ slots per bucket;
- the Cuckooing ACF, with $k=4$ hash tables and $b=1$ slots per bucket;
- three implementations of the Cyclic ACF described in 17, with $s=1, s=2$, and $s=3$ hash selector bits. To ensure a fair comparison in space usage, each hash selector bit used is accounted for with a corresponding decrease in the number of fingerprint bits (for example, a Cuckoo Filter with fingerprints of length 8 is compared to a Cyclic ACF with fingerprints of length 7 and $s=1$ hash selector bit); and
- a Swapping ACF as described in 17, with $b=4$ slots per bin and $k=2$ hash tables

All filters are at $95 \%$ occupancy in all of our experiments (i.e. $\gamma=1 / .95$ ). We give results for fingerprints of length $f=8$ bits on Chicago A, and summarize key results for Chicago B and San Jose, as well as results with $f=12$ and $f=16$ bits on all datasets. We repeat each experiment 10 times; all results given are the average performance over these 10 trials.

### 4.1 Experimental Results

The three plots in Figure 1 excluding the bottom-right plot show that the Cuckooing ACF has the strongest performance of all adaptive filters on network trace datasets with fingerprints of size $f=8$. The left hand plot in Figure 1 show that the Cuckooing ACF has the strongest performance of all adaptive filters on the Chicago A dataset with $f=8$. Its performance is particularly strong for low values of $|A| /|S|$ - that is to say, its performance is strong when the number of queries is small relative to $n$ (as one may expect given Theorem 1).

We ran further experiments, using fingerprints of size $f=8, f=12$, and $f=16$ on Chicago A, Chicago B, and San Jose datasets, achieving similar (in fact slightly better) results. These experiments are summarized in the right hand plot of Figure 1; we also give full results in Appendix B. The y-axis in this figure indicates the false positive rate of the given filter divided by the false positive rate of the Cuckooing ACF. This plot only includes the two best filters: the Swapping ACF, and the Cyclic ACF with $s=1$. We run the experiments for three fingerprint sizes $\{8,12,16\}$ on all three datasets, giving 18 total lines in the plot. Note that there is some overlap with the left hand plot-one of the bottommost two lines in the plot correspond to the Swapping ACF with $f=8$ on Chicago A. Specifically, the Cuckooing ACF does even better with larger fingerprints like $f=12$ and $f=16$ compared to $f=8$.

Overall, the Cuckooing ACF always performs at least as well as every other cuckoo filter on these datasets, frequently outperforming them by nearly a factor of 2 .

We believe that the simplicity of the Cuckooing ACF is, in fact, the source of this performance improvement. The Swapping ACF uses two bins of size 4, so each query is compared to twice as many fingerprints as in the Cuckooing ACF or Cuckoo Filter ${ }^{9}$ Meanwhile, the adaptivity bit of the Cyclic ACF causes each fingerprint to be one bit shorter, again doubling the false positive rate. Thus, the simple swapping strategy means we can avoid metadata bits with $b=1$, giving the Cuckooing ACF has an immediate factor- 2 advantage over its competitors.

## 5 Adersarial Adaptivity

Previous work leaves a dichotomy: the Adaptive Cuckoo Filters of Mitzenmacher et al. 17 work well in practice, whereas the "Broom Filter" of Bender et al. [3] is effective even against an adversary that tries to "learn" a filter's internal state. In Sections 3 and 4 we showed that the Cuckooing ACF is practical while retaining theoretical bounds. But our theoretical bounds are not adversarial; they are based on the number of unique queries made to the filter. Can our results be taken further - is there an ACF that adapts effectively even against an adversary?

In this section we give a general lower bound showing that an adversary can obtain a false positive rate of $\Omega(1)$ against any space-efficient ACF. This result is closely tied to a key structural distinction: the Broom Filter is difficult to implement because the length of the stored fingerprint may be different for each element. Our lower bound shows that this flexibility is, in fact, necessary in order to achieve adaptivity.

[^6]
### 5.1 Definition

Bender et al. [3] defined a notion of adaptivity that captures a worst-case adversary attempting to maximize the filter's false positive rate. We summarize this model in this subsection, and refer readers to [3] for a more thorough discussion.

In the adaptivity game, an adversary generates a sequence of queries. After each query $q$, the adversary and filter both learn if $q$ was a false positive. The filter may change its internal representation in response. The adversary will use whether or not $q$ was a false positive to determine the further queries.

At any time, the adversary may name a special element $\hat{q}$ - the adversary is asserting that this query is likely to be a false positive. The adversary "wins" if $\hat{q}$ is a false positive, and the filter "wins" if $\hat{q}$ is not a false positive.

The sustained false positive rate of a filter is the maximum probability $\epsilon$ with which the adversary can win the adaptivity game. We call a filter $\mathcal{F}$ adaptive if $\mathcal{F}$ can achieve a sustained false positive rate of $\epsilon$ for any constant $\epsilon$.

### 5.2 Lower Bounds

To begin, we note that the Cyclic ACF is not adaptive. A nearly-identical proof shows that the Swapping ACF is not adaptive.

Theorem 9. Let $\mathcal{F}$ be a Cyclic ACF with $k=O(1)$ hash tables, each with $N=\Theta(n)$ slots. Then there exists an adversarial strategy, making $\Theta\left(2^{s} / \epsilon^{2^{s}}\right)$ queries, which wins the adaptivity game against $\mathcal{F}$ with probability $\Omega(1)$. Thus the sustained false positive rate of $\mathcal{F}$ is $\Omega(1)$.

Proof. The adversary begins by selecting a query element $q_{1}$ at random. The adversary queries $q_{1} 2^{s}$ times. If $q_{1}$ is a false positive every time it is queried, the adversary sets $\hat{q} \leftarrow q_{1}$; otherwise the adversary chooses a new query element $q_{2}$ and repeats. This process is repeated until $O\left(1 / \epsilon^{2^{s}}\right)$ query elements have been chosen, requiring $O\left(2^{s} / \epsilon^{2^{s}}\right)$ queries overall.

We show that the adversary finds a $\hat{q}$ with probability $\Omega(1)$, and that $\hat{q}$ will be a false positive with probability 1.

Each time $q$ collides with an element $x_{i} \in S$, the hash selector bits associated with $x_{i}$ are incremented; thus, if $q$ does not collide with $x_{i}$ on the $j$ th query, it will not collide on the $j^{\prime}$ th query for $j^{\prime}>j$. Then if $q$ is a false positive on all $2^{s}$ collisions, there is an $x_{j} \in S$ such that $h^{f}(q, \alpha)=h^{f}\left(x_{j}, \alpha\right)$ for all $\alpha \in\left\{0, \ldots, 2^{s}-1\right\}$. We immediately obtain that any $\hat{q}$ found by the adversary is a false positive with probability 1.

We now bound the probability that $\hat{q}$ is found. For a given query element $q_{i}$ and a given element $x_{j} \in S$, the probability that $h_{\beta_{j}}^{\ell}\left(q_{i}\right)=h_{\beta_{j}}^{\ell}\left(x_{j}\right)$ is $1 / n$. The probability that, for all $\alpha, h^{f}\left(q_{i}, \alpha\right)=h^{f}\left(x_{j}, \alpha\right)$ is $1 / \epsilon^{2^{s}}$.

Thus, after the algorithm has considered $1 / \epsilon^{2^{s}}$ query elements, the probability that there exists a query element $q^{\prime}$ and an $x^{*} \in S$ such $x^{*}$ causes $q^{\prime}$ to be a false positive $2^{s}$ times is $1-\left(1-\epsilon^{2^{s}} / n\right)^{n / \epsilon^{2^{s}}}=\Omega(1)$. Thus, the adversary finds a $\hat{q}$ with constant probability.

We can also show that the Swapping ACF is not adaptive; this proof is essentially identical to that of Theorem 9 .

Theorem 10. Let $\mathcal{F}$ be a Swapping ACF with $N=\Theta(n)$ slots. Then there exists an adversarial strategy, making $\Theta\left(b / \epsilon^{b}\right)$ queries, which wins the adaptivity game against $\mathcal{F}$ with probability $\Omega(1)$.

One might think that hashing elements to another bucket (as is done in the Cuckooing ACF) is sufficient to make a filter adaptive. The reason Theorem 9 gives such a strong lower bound for the Cyclic ACF is that when we move an element to the next fingerprint, it is still a false positive with probability $\epsilon$. A constant number of these movements still leaves a significant probability that a false positive is not yet fixed. In contrast, when a colliding element is moved in the Cuckooing ACF, it still collides with the query with probability only $\epsilon / n$ - this seems low enough that almost all queries are successfully fixed after only a single movement.

Nonetheless, the adversary can use a birthday attack to find a small set of elements that cannot all be simultaneously fixed.

We obtain lower bounds for a fairly broad class of filters, where the total total information stored (i.e. hash index plus location plus fingerprint) for each element is at most $\log (n / \epsilon)+O(1)$ bits. This stands in contrast to the Broom Filter of Bender et al. [3], which is adaptive and which stores an average of $\log (n / \epsilon)+O(1)$ bits-in short, this proof shows that the nonuniformity of hash lengths in 3 is crucial to achieving adaptivity.

Definition 11. A deterministic $k$-adaptive filter $\mathcal{F}$ on $n$ elements with false positive rate $\epsilon$ is a filter satisfying the following:

- $\mathcal{F}$ has access to $k$ uniform random hash functions $h_{0}, \ldots, h_{k-1}$. Each hash has length at most $\log (N / \epsilon)$, for some $N=O(n / k)$ with $N \geq n / k$.
- For every configuration $C$ of $\mathcal{F}$, each $x_{i} \in S$ is stored using at least one hash $h_{C[i]}(x), 0 \leq C[i] \leq k-1$.
- The filter answers present to a query $q$ on configuration $C$ if there exists an $x_{i}$ such that $h_{C[i]}(q)=$ $h_{C[i]}\left(x_{i}\right)$. Otherwise, it answers absent.
- On a false positive $q, \mathcal{F}$ updates $C$ to a new configuration $C^{\prime}$ in round-robin order. In particular, if a query $q$ collides with an element $x_{i} \in S$ stored using $h_{\beta}$, then $x_{i}$ is stored in $C^{\prime}$ using $h_{\beta^{\prime}}$ satisfying $\beta^{\prime}=\beta+1(\bmod k)$.

By setting each hash $h_{i}$ in Definition 11 so that for any $i \in\{0, \ldots, k-1\}$ and $x \in U, h_{i}(x)$ is the concatenation of $h_{i}^{\ell}(x)$ and $h^{f}(x, i)$, the Cuckooing ACF is a deterministic $k$-adaptive filter. By setting $h_{(i, \alpha)}(x)$ to be the concatenation of $h_{i}^{\ell}(x)$ and $h^{f}(x, \alpha)$, the Cyclic ACF is a deterministic $k 2^{s}$-adaptive filter 10

The round-robin ordering requirement stands out as being a bit artificial, but our proof can fairly easily be generalized to handle other deterministic methods to update the configuration.

Theorem 12. There exists an adversarial strategy making $O(n)$ queries such that, for any deterministic $k$-adaptive filter $\mathcal{F}$ with $k<\log n / 6 \log \log n$ and $\epsilon>1 /\left(n^{1 / k}\right)$, the adversary wins the adaptivity game with probability $\Omega(1)$.

The rest of this section proves Theorem 12 .

Querying to Find a Mutually Unfixable Set. The proof of Theorem 12 begins with the adversary searching for a structure that "blocks" the filter, preventing it from fixing a false positive.

Consider a stored element $x_{i} \in S$, and fix a filter $\mathcal{F}$ with $k$ hash functions $h_{0}, \ldots h_{k-1}$. A set of queries $K$ is called mutually unfixable for $x_{i}$ if

- for all $\beta \in\{0, \ldots, k-1\}$, there exists a $q^{\prime} \in K$ such that $h_{\beta}\left(x_{i}\right)=h_{\beta}\left(q^{\prime}\right)$, and
- for all $q^{\prime} \in K$ there exists a $\beta, \in\{0, \ldots, k-1\}$ such that $h_{\beta}\left(x_{i}\right)=h_{\beta}\left(q^{\prime}\right)$.

The goal of our adversary is to find a mutually unfixable subset of the queries, as in any filter configuration, at least one element in a mutually unfixable set is a false positive.

The adversary begins by choosing a set $Q$ of $(1+1 / k) N /\left(\epsilon n^{1 / k}\right)$ queries selected uniformly at random from $U$. We show that if $Q$ is this size, then with constant probability $Q$ will contain $\Theta(1)$ mutually unfixable sets, each of size $O(k)$.

The adversary then queries members of $Q$ for $2 k$ rounds; any query that is a false positive during the second set of $k$ rounds is stored in a set $Q_{d}$. We show that $Q_{d}$ will be the union of some mutually unfixable subsets of $Q$.

[^7]Finally, the adversary repeatedly selects $k$ elements from $Q_{d}$ and randomly selects a permutation $P$ on these elements. The adversary queries these elements in order, twice. We show that, over $O(n)$ total queries, the adversary will (with constant probability) pick $k$ elements corresponding to a mutually unfixable set, and query them in an order such that each is a false positive every time it is queried. With this strategy, the adversary can find a false positive $\hat{q}$ with constant probability.

For each $\beta \in\{0, \ldots, k-1\}$, let $R_{Q}(\beta)$ be the set of values hashed to by $Q$ via $h_{\beta}$. In other words, $y \in R_{Q}(\beta)$ when there exists $q \in Q$ such that $h_{\beta}(q)=y$. We begin by showing that for all $\beta,\left|R_{Q}(\beta)\right|$ is very close to $|Q|$ with constant probability.

Lemma 13. Let $Q$ be any set of elements of size $(1+1 / k) N /\left(\epsilon n^{1 / k}\right)$, and assume $k \leq \log n / 6 \log \log n$. Then with probability at least $1 / 2, \min _{1 \leq \beta<k}\left|R_{Q}(\beta)\right|>|Q| /(1+1 / k)$.

Proof. We say that a $q \in Q$ is redundant for a hash $h_{\beta}$ if there is a $q^{\prime} \in Q \backslash\{q\}$ such that $h_{\beta}(q)=h_{\beta}\left(q^{\prime}\right)$. Then $|Q|-\left|R_{Q}(\beta)\right|$ is upper bounded by the number of redundant $q$ for $h_{\beta}$.

Via a union bound, the probability of a given query $q \in Q$ hashing under a given $h_{\beta}$ to one of the $\leq k|Q|$ values hashed to by another $q^{\prime} \in Q \backslash\{q\}$ is at most $\epsilon|Q| / N$. Thus, we can upper bound the number of redundant $q$ for $h_{\beta}$ by the sum of $|Q|$ Bernoulli trials, each of which succeeds with probability $\epsilon|Q| / N$. Then the total number of redundant $\left(q, h_{\beta}\right)$ in expectation is $\epsilon|Q|^{2} / N$. Using Chernoff bounds 18, Exercise 4.7], since $\epsilon|Q|^{2} / N \geq N / n^{1 / k}=\Omega(\log n)$, the probability that there are more than $2 \epsilon|Q|^{2} / N$ redundant pairs is at most $1 / n$. Taking the union bound over all $\beta \in\{0, \ldots, k-1\}$, all $R(\beta)$ have at most $2 \epsilon|Q|^{2} / N$ redundant pairs with probability $k / n$.

Plugging into our bound for $|Q|-\left|R_{Q}(\beta)\right|$, we obtain that $\left|R_{Q}(\beta)\right|>|Q|(1-2 \epsilon|Q| / N)$ for all $\beta$ with probability $\geq 1-k / n \geq 1 / 2$. (We weaken this bound to simplify.) Since $k \leq \log n / 6 \log \log n$ implies $n^{1 / k}>2(k+1) /(1+1 / k)$, we have $|Q|<N /(2 \epsilon(k+1))$; and therefore $1-2 \epsilon|Q| / N>1 /(1+1 / k)$.

We observe that there exists a mutually unfixable set $K \subseteq Q$ for some $x_{i} \in S$ if and only if, for all $\beta \in\{0, \ldots, k-1\}$, element $x_{i}$ hashes to an element of $R_{Q}(\beta)$ under $h_{\beta}$. Let $X_{Q}$ consist of all $x_{j} \in S$ such that there is a $K^{j} \subseteq Q$ that is mutually unfixable for $x_{j}$. That is, $X_{Q}=\left\{x_{j} \in S \mid h_{\beta}\left(x_{j}\right) \in R_{Q}(\beta)\right.$ for all $\beta \in$ $\{0, \ldots, k-1\}\}$. We begin by showing that $X_{Q}$ is nonempty with constant probability.

Lemma 14. For any deterministic $k$-adaptive filter $\mathcal{F}$ storing a set of elements $S$ of size $n$, and for any set of $(1+1 / k) N /\left(\epsilon n^{1 / k}\right)$ query elements $Q, Q$ contains a $K \subseteq Q$ such that $K$ is mutually unfixable for some $x_{j} \in S$. with probability $>(1-1 / e) / 2$.

Proof. By independence of the hashes, the probability that a given $x$ has a mutually unfixable set is $\prod_{\beta=1}^{k}\left(\epsilon\left|R_{Q}(\beta)\right| / N\right)$. By Lemma 13 . with probability $1-1 / n$, this is at least $\left(\frac{\epsilon|Q|}{N(1+1 / k)}\right)^{k}$. Substituting for $|Q|$, this means that each $x_{i} \in S$ has a mutually unfixable set with probability at least $1 / 2$.

Thus, if Lemma 13 is satisfied, the probability that no $x_{i}$ has a mutually unfixable set is $(1-1 / n)^{n} \leq 1 / e$. Therefore, $Q$ has a mutually unfixable set for some $x$ with probability at least $(1 / 2)(1-1 / e)$.

The next lemma upper bounds $\left|X_{Q}\right|$; with the above we have $\left|X_{Q}\right|=\Theta(1)$ with constant probability.
Lemma 15. For any deterministic $k$-adaptive filter $\mathcal{F}$ storing a set of elements $S$ of size $n$, and any set of $(1+1 / k) N /\left(\epsilon n^{1 / k}\right)$ query elements $Q$, let $X_{Q}$ be the elements of $S$ with a mutually unfixable subset of $Q$. Then with constant probability, $\left|X_{Q}\right|=O(1)$.

Proof. We have $\left|R_{Q}(\beta)\right| \leq|Q|$ for all $0 \leq \beta \leq k-1$ by definition; thus, the probability that a given $x$ has a mutually unfixable subset is $\left(\epsilon\left|R_{Q}\right| / N\right)^{k} \leq e / n$. Thus, $e$ elements $x_{i} \in S$ have mutually unfixable subsets of $Q$ in expectation; by Markov's inequality there are $O(1)$ elements in $S$ with mutually unfixable subsets with constant probability. Each maximal mutually unfixable set must be unfixable for some $x_{i} \in S$, so this proves $u=O(1)$.

Finally, we bound the size of each mutually unfixable set.

Lemma 16. If $k<\log n / 6 \log \log n$, then with constant probability, for all mutually unfixable $K \subseteq Q$, $k<|K|=O(k)$.
Proof. Consider an element $x_{i}$ such that some $K^{i} \subseteq Q$ of size $\left|K^{i}\right|=k$ is mutually unfixable for $x_{i}$. A further element $q \in Q \backslash K^{i}$ collides with $x_{i}$ (on some hash) with probability at most $\epsilon / N$; summing over $|Q|<n$ elements we obtain $\epsilon(n / N)<O(k)$ additional elements in expectation. By Markov's inequality, there are $O(k)$ such elements with constant probability. From Lemma 15 there are $O(1)$ elements $x \in S$ with mutually unfixable sets with constant probability. Therefore, this bound applies to all $O(1)$ mutually unfixable sets with constant probability.

To lower bound the size of these sets, we have that if the size of a set $K^{j}$ that is mutually unfixable for some $x_{j} \in S$ is less than $k$, there exists a $q_{1} \in K^{j}$ and two hash indices $\beta_{1}, \beta_{2} \in\{0, \ldots, k-1\}$ such that $q_{1}$ collides with $x_{j}$ under both $h_{\beta_{1}}$ and $h_{\beta_{2}}$. The probability that there exists such a pair $\left(q_{1}, x_{j}\right)$ is at most $k^{2}|Q| / n$, which is much less than a constant since $k<\log n / 6 \log \log n$. Thus, $|K| \geq k$ for all mutually unfixable $K$ with constant probability.

### 5.3 Lower bound for deterministic k-adaptive filters

The adversary begins by selecting a set of $(1+1 / k) N /\left(\epsilon n^{1 / k}\right)$ random queries $Q$. The adversary performs $2 k$ rounds, where in each round, the adversary queries every element of $Q$ in random order. We call the first $k$ such rounds the preliminary rounds, and we call the second $k$ rounds the testing rounds.

Let $Q_{d}$ be the set of elements that are false positives during any of the testing rounds. If $\left|Q_{d}\right| \neq O(k)$, the adversary fails ${ }^{11}$

The adversary performs a further $\binom{\left|Q_{d}\right|}{k} k$ ! rounds, the permutation rounds. In each permutation round, the adversary picks a random subset of $Q_{d}$ of size $k$, which we call $Q_{r}$. The adversary then picks a random permutation $P$ of $Q_{r}$. The adversary queries all elements in $Q_{r}$ in the order given by $P$, then repeats: again querying all elements in $Q_{r}$ in the order given by $P$. If every query in these repeated queries is a false positive, the adversary stops, and sets $\hat{q}$ to be the first element in $P$.

We now analyze this adversary. The idea is that after the $k$ preliminary rounds, the false positives during the testing rounds will consist of a union of mutually unfixable subsets of $Q$. We prove this in the following lemma.

Lemma 17. For any $Q_{d}$ obtained by the adversary, there exists a sequence of $u$ sets $K_{1}, \ldots, K_{u} \subseteq Q$ such that for all $\lambda, K_{\lambda}$ is a mutually unfixable set for some element $x_{i} \in S$, and $Q_{d}=\bigcup_{\lambda=1}^{u} K_{\lambda}$.

Proof. First, we show that every $q \in Q_{d}$ is a member of a mutually unfixable subset $K$. Consider an element $x_{i} \in S$ that does not have a mutually unfixable $K^{i} \subseteq Q$. In each round where $x_{i}$ collides with a false positive of $Q, x_{i}$ must be moved to the next hash. After less than $k$ movements, $x_{i}$ is moved to some hash $h_{\beta}$ where $x_{i}$ does not collide with any element of $Q$ under $\beta$; that is, $h_{\beta}\left(x_{i}\right) \notin R_{Q}(\beta)$. After this, $x_{i}$ will not be moved during subsequent queries from $Q$, and will not cause any elements of $Q$ to be a false positive.

Thus, any $x_{i}$ without a mutually unfixable $K^{i} \subseteq Q$ will not collide with any query in the testing rounds. This means that all $q^{\prime} \in Q_{d}$ collide with an $x_{j}$ that has a mutually unfixable subset $K^{i} \subseteq Q$. By definition, $q^{\prime} \in K^{i}$.

Now, consider a maximal mutually unfixable set $K^{\ell} \subseteq Q$, which is mutually unfixable for some $x_{\ell} \in S$. We show that $Q_{d}$ contains a mutually unfixable $K^{\prime} \subseteq K^{\bar{\ell}}$. During each round, at least one member of $K^{\ell}$ collides with $x_{\ell}$ by definition; thus $x_{\ell}$ is moved at least $k$ times during the testing rounds. This means that for each $\beta \in\{0, \ldots, k-1\}, x_{\ell}$ collides with a $q \in Q$ while being stored using hash $h_{\beta}$. Thus, by definition, the set of elements that collide with $x_{\ell}$ during the testing rounds form a mutually unfixable $K^{\prime} \subseteq K^{\ell}$.

With the above, we show that the adversary wins the adaptivity game with constant probability.
Lemma 18. With probability $\Omega(1)$, the adversary finds a $\hat{q}$ which is a false positive during the permutation rounds.

[^8]Proof. We assume that Lemma 15 is satisfied; this happens with probability $\Omega(1)$. Combining Lemma 15 with Lemma 17 . $Q_{d}$ consists of $u=O(1)$ mutually unfixable sets $K_{1}^{\prime}, \ldots K_{u}^{\prime}$, each of size $O(k)$.

Let $K^{i}$ be a mutually unfixable set of an element $x_{i} \in S$ such that $K^{i} \subseteq Q_{d}$ and $\left|K^{i}\right|=k$. We analyze the probability that, in a given permutation round, the adversary chooses $Q_{r}=K^{i}$, and chooses a permutation $P$ that causes each query to be a false positive; this means that the adversary selects a $\hat{q}$ with probability $\Omega(1)$. Then we show that if the adversary selects a $\hat{q}$, it is a false positive with probability $\Omega(1)$.

Fix a permutation round, and let $\beta_{i}$ be the hash index of $x_{i}$ at the beginning of this permutation round. Let $P_{K^{i}}=\left(q_{1}, q_{2}, \ldots, q_{k}\right)$ be a permutation on $K^{i}$ such that for each $q_{j}, h_{\beta_{i}+j(\bmod k)}\left(q_{j}\right)=h_{\beta_{i}+j(\bmod k)}\left(x_{i}\right)$.

If $P_{K^{i}}$ is selected, then all queries to $P_{K^{i}}$ will be false positives, and the adversary wins the adaptivity game with probability 1 . We select $Q_{r}$ to consist of the members of $K^{i}$ with probability $1 /\binom{\left|Q_{d}\right|}{k}$, and we query them in order of $P_{K^{x}}$ with probability $1 /(k!)$. Repeating this process for $\binom{\left|Q_{d}\right|}{k} k$ ! rounds gives constant probability of success.

Now, we show that, with constant probability, the adversary only sets $\hat{q}$ for a permutation $P^{\prime}=$ $\left(q_{1}^{\prime}, q_{2}^{\prime}, \ldots, q_{k}^{\prime}\right)$ if $P^{\prime}$ satisfies, for some $x_{\ell}$ with hash index $\beta_{\ell}, h_{\beta_{\ell}+j(\bmod k)}\left(q_{j}\right)=h_{\beta_{\ell+j}(\bmod k)}\left(x_{\ell}\right)$ for all $1 \leq j \leq k^{\prime}$. Recall that the adversary sets $\hat{q}$ when the $j$ th query to $P$ is a false positive for all $j$; thus for each $q_{j}$ there exists an $x_{j} \in S$ and a $0 \leq \beta_{j}<k$ such that $h_{\beta_{j}}\left(q_{j}\right)=h_{\beta_{j}}\left(x_{j}\right)$. If $x_{j}=x_{\ell}$ for all $j$ then we are done.

Otherwise, let $\hat{X} \subseteq S \backslash\left\{x_{\ell}\right\}$ be the elements of $S$ (other than $x_{i}$ ) that collide with an element of $Q_{r}$. If all $x_{j} \in \hat{X}$ only collide with a member of $Q_{r}$ under a single hash function $h_{\beta}$ (i.e. there do not exist $q_{1}, q_{2} \in Q_{r}$ and distinct $\beta_{1}, \beta_{2} \in\{0, \ldots, k-1\}$ with $h_{\beta_{1}}\left(x_{j}\right)=h_{\beta_{1}}\left(q_{1}\right)$ and $\left.h_{\beta_{2}}\left(x_{j}\right)=h_{\beta_{2}}\left(q_{2}\right)\right)$, then $x_{j}$ will only be a false positive once. Thus, all collisions in the repeated queries will collide with $x_{i}$, and again we are done.

The probability that there exists an $x_{\ell} \in S$ and distinct $\beta_{1}, \beta_{2} \in\{0, \ldots, k-1\}$ with $h_{\beta_{1}}\left(x_{\ell}\right)=h_{\beta_{1}}\left(q_{1}\right)$ and $h_{\beta_{2}}\left(x_{\ell}\right)=h_{\beta_{2}}\left(q_{2}\right)$ for some $q_{1}, q_{2} \in Q_{d}$ is at most $k^{2}\left|Q_{d}\right|^{2} \frac{\epsilon^{2}}{n^{2}}=o(1)$.

The number of queries required for this adversary is, for some constant $c$ depending on the maximum size of $Q_{d}, O\left(n^{1-1 / k} / \epsilon+c^{k} k!\right)$; this is $O(n)$ since $k<\log n / 6 \log \log n$ and $\epsilon>1 /\left(n^{1 / k}\right)$. This completes the proof of Theorem 12

## References

[1] Yuriy Arbitman, Moni Naor, and Gil Segev. De-amortized cuckoo hashing: Provable worst-case performance and experimental results. In International Colloquium on Automata, Languages, and Programming, pages 107-118. Springer, 2009.
[2] Michael A Bender, Rathish Das, Martín Farach-Colton, Tianchi Mo, David Tench, and Yung Ping Wang. Mitigating false positives in filters: to adapt or to cache? In Symposium on Algorithmic Principles of Computer Systems (APOCS), pages 16-24. ACM-SIAM, 2021.
[3] Michael A Bender, Martin Farach-Colton, Mayank Goswami, Rob Johnson, Samuel McCauley, and Shikha Singh. Bloom filters, adaptivity, and the dictionary problem. In Foundations of Computer Science (FOCS), pages 182-193. IEEE, 2018.
[4] 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. Don't thrash: how to cache your hash on flash. Proc. VLDB Endowment, 5(11):1627-1637, 2012.
[5] Burton H Bloom. Space/time trade-offs in hash coding with allowable errors. Communications of the ACM, 13(7):422-426, 1970.
[6] Andrei Broder and Michael Mitzenmacher. Network applications of bloom filters: A survey. Internet mathematics, 1(4):485-509, 2004.
[7] J Bruck, Jie Gao, and Anxiao Jiang. Weighted bloom filter. In IEEE International Symposium on Information Theory, 2006.
[8] Larry Carter, Robert Floyd, John Gill, George Markowsky, and Mark Wegman. Exact and approximate membership testers. In Symposium on Theory of Computing (STOC), pages 59-65, 1978.
[9] Bernard Chazelle, Joe Kilian, Ronitt Rubinfeld, and Ayellet Tal. The bloomier filter: an efficient data structure for static support lookup tables. In Symposium on Discrete Algorithms, pages 30-39, 2004.
[10] David Eppstein. Cuckoo filter: Simplification and analysis. In Scandinavian Symposium and Workshops on Algorithm Theory (SWAT), volume 53, pages 8:1-8:12, 2016.
[11] Bin Fan, Dave G Andersen, Michael Kaminsky, and Michael D Mitzenmacher. Cuckoo filter: Practically better than Bloom. In International Conference on Emerging Networking Experiments and Technologies (CoNEXT), pages 75-88. ACM, 2014.
[12] Dimitris Fotakis, Rasmus Pagh, Peter Sanders, and Paul Spirakis. Space efficient hash tables with worst case constant access time. Theory of Computing Systems, 38(2):229-248, 2005.
[13] Alan Frieze, Páll Melsted, and Michael Mitzenmacher. An analysis of random-walk cuckoo hashing. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, pages 490-503. Springer, 2009.
[14] Shahabeddin Geravand and Mahmood Ahmadi. Bloom filter applications in network security: A state-of-the-art survey. Computer Networks, 57(18):4047-4064, 2013.
[15] Adam Kirsch, Michael Mitzenmacher, and Udi Wieder. More robust hashing: Cuckoo hashing with a stash. SIAM Journal on Computing, 39(4):1543-1561, 2009.
[16] Shachar Lovett and Ely Porat. A lower bound for dynamic approximate membership data structures. In Foundations of Computer Science (FOCS), pages 797-804. IEEE, 2010.
[17] Michael Mitzenmacher, Salvatore Pontarelli, and Pedro Reviriego. Adaptive cuckoo filters. In Workshop on Algorithm Engineering and Experiments (ALENEX), pages 36-47, 2018.
[18] Michael Mitzenmacher and Eli Upfal. Probability and computing: randomization and probabilistic techniques in algorithms and data analysis. Cambridge university press, 2017.
[19] Moni Naor, Gil Segev, and Udi Wieder. History-independent cuckoo hashing. In International Colloquium on Automata, Languages, and Programming, pages 631-642. Springer, 2008.
[20] Anna Pagh, Rasmus Pagh, and S Srinivasa Rao. An optimal bloom filter replacement. In Symposium on Discrete Algorithms (SODA), pages 823-829. ACM-SIAM, 2005.
[21] Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. Journal of Algorithms, 51(2):122-144, 2004.
[22] Prashant Pandey, Michael A. Bender, Rob Johnson, and Rob Patro. A general-purpose counting filter: Making every bit count. In International Conference on Management of Data (SIGMOD), pages 775-787. ACM, 2017.
[23] Ely Porat. An optimal bloom filter replacement based on matrix solving. In International Computer Science Symposium in Russia (CSR), pages 263-273, 2009.
[24] Minmei Wang, Mingxun Zhou, Shouqian Shi, and Chen Qian. Vacuum filters: more space-efficient and faster replacement for bloom and cuckoo filters. Proc. VLDB Endowment, 13(2):197-210, 2019.

## A Symbol Table

The following symbols are used throughout the paper:

| Symbol | Usage |
| :---: | :---: |
| $S$ | The set of elements stored in the filter |
| $x_{i}$ | The $i$ th element stored in the filter |
| $\epsilon$ | The static false positive rate; used as a parameter in adaptive filters |
| $U$ | Universe of possible elements; all queries and stored elements are from $U$ |
| $Q$ | The set of elements queried in a query sequence |
| $h_{i}^{\ell}$ | The $i$ th location hash function |
| $h_{i}^{f}$ | The $i$ th fingerprint hash function |
| $f$ | The size of the fingerprint |
| $k$ | The number of hash tables |
| $b$ | The number of slots per bin |
| $N$ | The number of bins in a hash table; $N=\gamma n /(b k)$ |
| $\beta_{i}$ | The hash index of $x_{i}\left(\right.$ i.e. the table storing $\left.x_{i}\right)$ |
| $B\left(x_{i}\right)$ | The bin storing the fingerprint of $x_{i} ; B\left(x_{i}\right)=h_{\beta_{i}}^{\ell}\left(x_{i}\right)$ |
| $C$ | A configuration of a filter, defined using the hash index of each stored element $x_{i}$ |
| $C[i]$ | The hash index of $x_{i}$ in configuration $C$ |
| $s$ | The number of hash selector bits in a Cyclic ACF |
| $\alpha$ | The value stored using the hash selector bits for a given slot |
| $A$ | The set of elements queried during a given experiment |

The following symbols are used in the proof that the Cuckooing ACF is adaptive in Section 3

| Symbol | Usage |
| :---: | :---: |
| $C_{0}$ | The configuration before any query is performed |
| $C_{i}$ | The configuration after query $q_{i}$ |
| $k_{i}$ | The number of elements moved when fixing $q_{i}$ on configuration $C_{i-1}$ |
| $x_{i_{1}}, x_{i_{2}}, \ldots x_{i_{k_{i}}}$ | The sequence of elements moved when fixing $q_{i}$ on configuration $C_{i-1}$ |
| $B(i, C)$ | The bin storing $x_{i}$ in configuration $C ; B(i, C)=h_{C[i]}^{\ell}\left(x_{i}\right)$ |
| $B^{\prime}(i, C)$ | The alternate bin that can store $x_{i}$ in configuration $C ; B^{\prime}(i, C)=h_{1-C[i]}^{\ell}\left(x_{i}\right)$ |
| $F_{0}$ | The set of initial false positives |
| $k_{i}^{0}$ | The number of elements moved when fixing $q_{i}$ on configuration $C_{0}$ |
| $x_{i_{1}^{\prime}}, x_{i_{2}^{\prime}}, \ldots x_{i_{k_{i}^{\prime}}^{\prime}}$ | The sequence of elements moved when fixing $q_{i}$ on configuration $C_{0}$ |
| $\Phi(t)$ | The potential after query $t$ |

The following symbols are used in the lower bound (the proof of Theorem 12 in Section 5 .

| Symbol | Usage |
| :---: | :---: |
| $\hat{q}$ | A special query that the adversary names as a likely false positive |
| $R_{Q}(\beta)$ | The set of values that can be obtained by applying hash $h_{\beta}$ to any element in $Q$ |
| $X_{Q}$ | The set of all $x_{j}$ such that there exists a $K^{j} \subseteq Q$ that is mutually unfixable for $x_{j}$ |
| $Q_{d}$ | The set of queries that are false positives during any of the testing rounds |
| $Q_{r}$ | A random subset of $Q_{d}$ of sized $k ;$ this subset is queried during a permutation round |
| $P$ | A permutation applied to some $Q_{r}$ during a permutation round |
| $P_{K^{i}}$ | A permutation on a mutually unfixable set of $x_{i}$ with certain hash collision properties |
| $\hat{X}$ | The elements of $S$ other than a fixed $x_{i}$ that collide with any element of $Q_{r}$ |

## B Experiments On Network Traces




Figure 2: The false positive rate incurred for each filter on the Chicago A dataset, with 8, 12, and 16 fingerprint bits.


Figure 3: The false positive rate incurred for each filter on the Chicago B dataset, with 8, 12, and 16 fingerprint bits.



Figure 4: The false positive rate incurred for each filter on the San Jose dataset, with 8, 12, and 16 fingerprint bits.


[^0]:    *This work was supported in part by ISF grants no. 1278/16 and 1926/19, by a BSF grant no. 2018364, and by an ERC grant MPM under the EU's Horizon 2020 Research and Innovation Programme (grant no. 683064). A preliminary version of this paper appeared at the 17th Algorithms and Data Structures Symposium.
    †Bar-Ilan University, Ramat Gan, Israel. Email: kopelot@gmail.com
    $\ddagger$ Williams College, Williamstown, MA 01267 USA. Email: srm2@williams.edu
    §Bar-Ilan University, Ramat Gan, Israel. Email: porately@cs.biu.ac.il

[^1]:    ${ }^{1}$ Note that the filter does not have access to this sequence ahead of time; it must process the queries online. We fix a predetermined sequence of queries to clarify that, unlike in 3 , we do not allow each new query to be determined adversarially based on the result of previous queries.
    ${ }^{2}$ In fact, if $S$ is stored to allow for efficient, in-place reverse hash lookups-so that the normal hash lookup and the adaptive rehashing can be performed simultaneously - then some false positives may be fixed without increasing the number of remote accesses at all.

[^2]:    ${ }^{3}$ We assume $\gamma n$ is an integer multiple of $b k$ for simplicity.
    ${ }^{4}$ When treating the hash value as a bit string we assume that $N$ is a power of two for simplicity; this assumption is not necessary for the implementation.
    ${ }^{5}$ While such strong hashes are not useable in practice, this analysis is generally predictive of experimental results (see i.e. $11,17 \mid 22$ ).

[^3]:    ${ }^{6}$ That is to say, $\gamma=2$.

[^4]:    ${ }^{7}$ The final term $B^{\prime}\left(i_{i_{k_{i}}}, C_{i-1}\right)$ denotes that the last element is swapped to a new position, and does not "cuckoo" any further elements.

[^5]:    ${ }^{8}$ Strictly speaking, these events are not independent. After all, if $k_{i}^{0}=n$, then $k_{j}^{0}$ is very likely to be $\Omega(n)$. However, since we assume $k_{i}^{0}=O(\log n)$ and $k_{j}^{0}=O(\log n)$, only $O(\log n)$ hash table slots are touched by either path; this affects the expectation of each by at most $O((\log n) / n)$.

[^6]:    ${ }^{9}$ If the bins are decreased in size to, say, 2, the Swapping ACF requires extremely frequent rebuilds with the desired load factor-the factor-2 decrease in efficiency is intrinsic to the Swapping ACF.

[^7]:    ${ }^{10}$ Strictly speaking the Cyclic ACF does not satisfy Definition 11 because the hashes are not all independent. However, this only increases the ability of an adversary to find false positives.

[^8]:    ${ }^{11}$ Specifically, for a sufficiently large constant $c$, the adversary fails if $\left|Q_{d}\right|>c k$. The constant $c$ gives a tradeoff between the lower-order terms in the number of queries performed, and the constant probability that the adversary wins.

