

The Consistent Cuckoo Filter

*Lailong Luo *✉Deke Guo †Ori Rottenstreich ‡Richard T.B Ma *Xueshan Luo *Bangbang Ren

*Science and Technology on Information System Engineering Laboratory, National University of Defence Technology

† Computer Science and the Department of Electrical Engineering, Technion

‡School of Computing, National University of Singapore

{luolailong09, dekeguo, xsluo, renbangbang11}@nudt.edu.cn, or@cs.technion.ac.il, tbma@comp.nus.edu.sg

Abstract—The emergence of large-scale dynamic sets in networking applications attaches stringent requirements to approximate set representation. The existing data structures (including Bloom filter, Cuckoo filter, and their variants) preserve a tight dependency between the cells or buckets for an element and the lengths of the filters. This dependency, however, degrades the capacity elasticity, space efficiency and design flexibility of these data structures when representing dynamic sets. In this paper, we first propose the Index-Independent Cuckoo filter (I2CF), a probabilistic data structure that decouples the dependency between the length of the filter and the indices of buckets which store the information of elements. At its core, an I2CF maintains a consistent hash ring to assign buckets to the elements and generalizes the Cuckoo filter by providing optional k candidate buckets to each element. By adding and removing buckets adaptively, I2CF supports the bucket-level capacity alteration for dynamic set representation. Moreover, in case of a sudden increase or decrease of set cardinality, we further organize multiple I2CFs as a Consistent Cuckoo filter (CCF) to provide the filter-level capacity elasticity. By adding untapped I2CFs or merging under-utilized I2CFs, CCF is capable of resizing its capacity instantly. The trace-driven experiments indicate that CCF outperforms its alternatives and realizes our design rationales for dynamic set representation simultaneously, at the cost of a little higher complexity.

Index Terms—Cuckoo filter, consistent hashing, elasticity

I. INTRODUCTION

Set representation while supporting membership queries is a fundamental problem in databases, caches, routers, storage, and networking applications. These systems often represent set elements with a probabilistic data structure and support constant-time approximate membership query with small false positive probability. The most widely-used probabilistic data structures for approximate membership query are Bloom filter [1] [2] [3], Cuckoo filter [4] and their variants [5] [6] [7] [8].

Bloom filter and Cuckoo filter represent sets in diverse ways. Bloom filter is a fixed-length array of bits which are initialized as 0s. To insert an element, k independent hash functions are employed to map the element into the bit vector. Thereafter, the corresponding bits are set to 1s. When testing the membership of any element x , Bloom filter just checks the k corresponding bits are non-zero. If they are all ones, Bloom filter concludes that x is a member of the set (possibly implying a false positive); otherwise, it correctly indicates that x is not a member (no false negatives). Unlike Bloom filter, Cuckoo filter stores the fingerprints of the elements with their candidate buckets directly. Cuckoo filter derives two candidate buckets for each element with the partial cuckoo hashing strategy [9] and tries to store the fingerprint into one of the candidate buckets. An element is identified as a member of the set if its fingerprint can be found in either of its candidate buckets. Bloom filter and Cuckoo filter, however, fail

to represent dynamic set members because of their incapability of resizing their capacities.

To this end, Dynamic Bloom filter (DBF) [5] and Dynamic Cuckoo filter (DCF) [6] have been developed. Both DBF and DCF attempt to add and merge homogenous Bloom filters and Cuckoo filters to extend and downsize their capacities on demand. In both DBF and DCF, the length of each filter is predefined and cannot be altered since the indices of cells or candidate buckets are determined by calculating the modulus based on the length of the filter. As a consequence, they can only resize their capacity by adding or merging homogenous filters. In the worst case where one filter has to be added to store only one additional element, the resultant space utilization can be even less than 50%. Therefore, in space-scarce scenarios, the bucket-level capacity alteration is necessary to save space. Moreover, a major weakness of DBF is that it fails to support reliable element deletion [6] since there may be multiple BFs which satisfy the membership query condition. Although DCF guarantees reliable element deletion, it employs the XOR operation to derive the second candidate bucket during reallocations. Therefore, the length of each Cuckoo filter can only be of the form $m=2^\gamma$ ($\gamma \geq 0$). If not, the XOR results may go out of range.

Consequently, we envision a design of probabilistic data structure which properly concerns the following three design rationales for dynamic set representation.

- **Capacity elasticity (CE).** The data structure's capacity is adaptively adjustable according to the set cardinality. Despite the unpredictability of the number of elements to represent, the offered capacity shows coincident changing trends as the set cardinality adaptively.
- **Space efficiency (SE).** The space utilization remains at a high level irrespective of the variation of set cardinality. This is extremely important for space-scarce scenarios, e.g., wireless sensor networks.
- **Design flexibility (DF).** All the parameters are adjustable so that users can customize their own configurations according to their design goals. For example, the number of hash functions may be increased for higher space utilization or decreased for better query throughput.

These rationales, if realized, will bring unprecedented benefits for set representation and membership query, in terms of space-saving and quality of service. The design flexibility further extends the applicability of the data structure to more general scenarios with diverse requirements.

The existing probabilistic data structures, however, fail to achieve the three rationales properly and simultaneously. As shown in Table 1, Bloom filter and DBF achieve low space

TABLE I
THE FEATURES OF SET REPRESENTATION DATA STRUCTURES. THE
NUMBER OF “+” QUANTIFIES THAT FEATURE.

Name	BF	DBF	CF	DCF	ACF	SCF	I2CF	CCF
CE		++		++			++	+++
SE	+	+	++	++	++	++	+++	+++
DF	++	++	+	+	+	+	+++	+++

utilization. The reason is that they keep half of the bits as 0s, in order to incur the least false positive rate. By contrast, Cuckoo filter and its variants improve their space utilization with the reallocation strategy during each insertion. DBF and DCF offer capacity elasticity to some extent by adding and merging filters dynamically. However, in reality, a more fine-grained capacity scaling is needed to handle small-scale capacity overflows and recycle space timely when a few elements are removed. Furthermore, existing data structures are somehow hindered by their limited design flexibility. In the framework of Bloom filters, the parameters have to be carefully designed to guarantee their target false positive rate. Meanwhile, current proposals of Cuckoo filters must use a fixed number of hash functions and a power of two number of buckets.

A common reason for the existing data structures’ deficiency of achieving the three rationales is that they preserve a tight dependency between the cells or buckets for an element and the lengths of the filters. As a result, their capacities have to be predefined and remain immutable irrespective of the change of dynamic sets. Therefore, in this paper, we first propose I2CF, a probabilistic data structure which decouples the dependency between the length of the filter and the indices of buckets which store the information of elements. At its core, an I2CF maintains a *consistent hash ring* [10] [11] to assign buckets to the elements and generalizes the Cuckoo filter by providing optional k candidate buckets to each element. By adding and removing buckets adaptively, I2CF supports bucket-level capacity alteration for dynamic set representation.

Moreover, in case of a sudden increase or sharp decrease of set cardinality, we further organize multiple I2CFs as a CCF to provide filter-level capacity elasticity. By adding untapped I2CFs or merging under-utilized I2CFs, CCF is capable of resizing its capacity instantly. As shown in Table 1, both I2CF and CCF offer elegant space efficiency and design flexibility. CCF has better capacity elasticity than I2CF, since I2CF only provides bucket-level capacity alteration, while CCF additionally supports filter-level capacity adjustment for dynamic sets. In fact, I2CF is a special case of CCF when only one I2CF is maintained. To summarize, we achieve the following contributions.

- We first design I2CF (Index Independent Cuckoo filter), a probabilistic data structure which decouples the dependency between the length of the filter and the indices of buckets which store the information of elements. It allows flexibility in the memory size without the need for reallocating most elements. Thereafter, we organize multiple I2CFs as a CCF and present the algorithms for dynamic set representation and capacity resizing.
- For any I2CF with given parameters in a CCF, we present a new threshold for the ratio between the number of represented elements and the number of buckets. Additionally, we derive an upper bound for the probability that the

given number of elements can be successfully stored in a given I2CF.

- Trace-driven evaluations are conducted to measure the performance of our proposals. The results show that CCF outperforms DCF and realizes the three design rationales simultaneously at the cost of a little higher time-complexity.

The rest of this paper is organized as follows. Section II introduces the background and related work. Section III presents the I2CF and CCF design and their operations. Section IV presents the performance analysis for CCF theoretically. Section V reports the evaluation results and at last Section VI concludes the whole paper.

II. BACKGROUND AND RELATED WORK

Cuckoo filter (CF) [4] is a light-weight probabilistic data structure to support constant-time membership query. Unlike Bloom filter, CF stores the fingerprint of each element directly. Structurally, a CF consists of m buckets, each is capable of residing b fingerprints. An element x is associated with a f -bit fingerprint η_x which is derived out by a hash function h_0 . CF offers 2 candidate buckets to each element, and the fingerprint can be stored in each of the candidates. If both candidates are occupied, Cuckoo filter randomly kicks out a fingerprint in one of the candidates and reinserts the victim in its other candidate bucket. This reallocation ends successfully when a bucket has available space and fails when the number of such reallocations reaches a given threshold max . During reallocation, the alternative bucket can be derived out by executing an XOR operation towards the current bucket and the fingerprint of the victim. That is, the two bucket locations are derived as $h_1(x)=hash(x)$ and pair-wisely $h_2(x)=h_1(x)\oplus hash(\eta_x)$.

To query whether an element y is a member of set A or not, CF checks the two corresponding buckets of y . If the fingerprint η_y is found in one of these two buckets, CF judges $y\in A$; otherwise CF concludes $y\notin A$. Due to the potential hash collisions of the fingerprints, CF may suffer from false positive errors (referring to elements which do not belong to A as members of A). The false positive rate of CF satisfies $\xi_{CF}\leq 1-(1-\frac{1}{2^f})^{2b}$. Note that there are no false negative errors for the stored elements.

Most recently, several CF variants have been proposed to further improve its performance [6] [7] [8]. The Simplified Cuckoo filter [7] (SCF) calculates the indices of buckets for an element x as $h_1(x)$ and $h_1(x)\oplus\eta_x$. In this way, SCF provides a theoretical guarantee to its performance. The Adaptive cuckoo filter [8] (ACF) tries to remove false positive errors from the CF vector by resetting the collided fingerprints with optional hash functions. Inspired by the Dynamic Bloom filter [5], Dynamic Cuckoo filter [6] (DCF) dynamically maintains multiple homogenous CFs to enable elastic capacity. Initially, only one CF is maintained and marked as active. The subsequent homogenous CFs will be introduced with an either active or passive manner. A recycling mechanism is suggested to merge under-loaded CFs, thereby improving the space utilization. The upper bound of false positive rate in DCF is $1-(1-\xi_{CF})^s$, where ξ_{CF} is the false positive rate of each CF vector (from above) and s is the number CFs in DCF.

The above variants of CF, however, fail to realize our design rationales properly. SCF and ACF leverage the employed hash




Fig. 1. An illustrative example of I2CF with $k=1$, $b=2$ and $v=1$. It represents a set with 4 elements, including x , y , z and w . Initially, the I2CF employs 3 buckets to store these elements. When a new bucket B_3 is added as a successor of B_2 , the fingerprint of z is reallocated from B_0 to B_3 , since the hash value of η_z is mapped between B_2 and B_3 . When B_0 is removed, the fingerprint η_y is reallocated to B_0 's successor B_1 .

functions only, but their capacities cannot be resized after implementation. DCF supports filter level capacity alteration but incurs limited design flexibility and untimely space recycling. Consequently, we present CCF, a novel probabilistic data structure which promises capacity elasticity, high space utilization, and design flexibility simultaneously.

III. CONSISTENT CUCKOO FILTER

We describe the design of the CCF (Consistent Cuckoo filter) in detail here, including its data structure, operations for set representation and resizing strategies. Before that, we introduce the I2CF (Index-Independent Cuckoo filter), which is the basic component of CCF.

A. Design of Index-Independent Cuckoo Filter (I2CF)

To represent dynamic sets, the employed data structure should offer elastic capacity. Although DBF and DCF are capable of filter-level capacity elasticity, they fail to provide the ability of fine-grained capacity alteration. The reason is that their lengths of filters are predefined and immutable throughout their lifetimes. The using of XOR operations to compute hash values in Cuckoo filter further exacerbates the capacity elasticity by restricting the filter length to be a power of two. Therefore, we redesign the framework of Cuckoo filter and propose the Index-Independent Cuckoo filter (I2CF) here.

Basically, I2CF consists of multiple buckets, each of which has b slots. That is, each bucket can accommodate b fingerprints at most. As shown in Fig. 1, the buckets are mapped onto a consistent hash ring [10] [11] ranging from 1 to $M-1$. To ensure better load balance in the consistent hash ring, each bucket has $v \geq 1$ virtual nodes in the consistent hash ring. I2CF also stores the fingerprints of elements instead of the actual contents by offering each fingerprint $k \geq 1$ candidate buckets. An element is successfully represented if its fingerprint is stored in one of its candidate buckets. To determine the candidate buckets of an element x , k independent hash functions are employed to map the fingerprint η_x onto the consistent hash ring. Thereafter, the k nearest buckets (in a clockwise order by default) of the k hash values are regarded as the candidate buckets of η_x . In this way, the candidate buckets are index-independent and decided by the consistent hashing. A fingerprint can be stored in any one of these candidate buckets. If all the candidate buckets are fully occupied, I2CF randomly kicks out an existing fingerprint from one of these buckets to

store the fingerprint to insert. The victim will be reallocated to one of the other candidate buckets. The reallocation ends successfully when a bucket has available space and fails when the number of such reallocations reaches a given threshold max .

Compared with Cuckoo filter, I2CF has two major improvements. First, I2CF organizes the buckets as a consistent hash ring to decouple the dependency between candidate buckets and the length of the filter. As a consequence, I2CF naturally enables the capability of adding and removing buckets on demand. A toy example of adding and removing buckets from an I2CF is given in Fig. 1. Second, I2CF generalizes the number of candidate buckets from the fixed two in Cuckoo filter as a mutable variable k . This generalization further improves its design flexibility. Moreover, as analyzed later, larger k values also guarantee higher space utilization. With these improvements, I2CF achieves the bucket-level capacity elasticity and high space utilization to represent dynamic sets.

B. Overview of Consistent Cuckoo Filter (CCF)

I2CF provides bucket-level capacity elasticity, but when set cardinality increases drastically, a single I2CF may fall short of offering enough space timely. Therefore, we further generalize I2CF as CCF which dynamically maintains multiple I2CFs. Just like existing CF variants, CCF also leverages fingerprints to represent elements in a set. The fingerprint for an element x is generated by mapping x into a given range $[0, 2^f - 1]$ with a hash function h_0 . Basically, a CCF consists of s ($s \geq 1$ and initialized as 1) *heterogeneous* I2CFs. An arbitrary I2CF $_i$ ($i \in [0, s-1]$) has $m_i \geq 1$ buckets with $b_i \geq 1$ slots. The employed number of hash functions k_i and the value of M_i in I2CF $_i$ are also allowed to be different from other I2CFs. With such a framework, CCF enables ultimate design flexibility. Note that to multiplex the calculated hash values for each fingerprint among the I2CFs, we prefer $k_0 = \dots = k_i = \dots = k_{s-1} = k$, and $M_0 = \dots = M_i = \dots = M_{s-1} = M$ by default. More importantly, CCF provides capacity elasticity at both the bucket level and filter level. That is, its capacity can be altered by adding or removing buckets in any I2CF, as well as introducing untapped or compacting under-utilized I2CFs. The details are given in Section III-D. When an I2CF is extended or introduced, it will be marked as active to store new elements.

Theorem 1: For an I2CF $_i$ ($i \in [0, s-1]$) in CCF, let b_i and k_i denote the number of slots in each bucket and the number of

candidate buckets in the filter I2CF_{*i*}, respectively. The false positive rate for a CCF query can be calculated as:

$$\xi_{CCF} = 1 - \prod_{i=0}^{s-1} (1 - \xi_i) = 1 - \prod_{i=0}^{s-1} \left(1 - \frac{1}{2^f}\right)^{k_i \cdot b_i}. \quad (1)$$

When $k_0 = \dots = k_i = \dots = k_{s-1} = k$, $b_0 = \dots = b_i = \dots = b_{s-1} = b$,

$$\xi_{CCF} = 1 - \prod_{i=0}^{s-1} (1 - \xi_i) = 1 - \left(1 - \frac{1}{2^f}\right)^{s \cdot k \cdot b} \approx \frac{s \cdot k \cdot b}{2^f}. \quad (2)$$

The false positive error of CCF stems from the hash collisions of the fingerprints. If two elements $x \in A$ and $y \notin A$ share the same fingerprint, i.e., $\eta_x = \eta_y$, the membership query of y implies a false positive error due to the existence of x . Within the CCF framework, a membership query may check all of the s I2CF vectors. For I2CF_{*i*}, the false positive rate is $\xi_i = 1 - \left(1 - \frac{1}{2^f}\right)^{k_i \cdot b_i}$. The global false positive rate is thus derived out as $\xi_{CCF} = 1 - \prod_{i=0}^{s-1} (1 - \xi_i)$. Note that, both DCF and CCF have multiple filters and share the same false positive rate. Generally, larger f leads to lower false positive rate, while larger k , b and s result in higher false positive rate. However, DCF fails to support runtime false positive rate guarantee since the value of s will be increased continually with the increase of set cardinality. As a consequence, the false positive of DCF keeps increasing when more CFs are launched. CCF, on the contrary, provides runtime false positive rate guarantee by setting a threshold for s . If the value of s reaches the threshold, on one hand, CCF can conduct the compact operation to hopefully merge some I2CF vectors. On the other hand, CCF only employs the bucket-level resizing strategy to accommodate the coming elements, thereby the value of s will not be increased any more. Thus, the false positive rate can be reasonably bounded.

C. Dynamic Set Representation with CCF

In this subsection, we present the basic operations of CCF for dynamic set representation, including insertion, query, and deletion of elements. The associated resizing strategies are detailed in later Section III-D.

Insertion. The CCF tracks the number of elements inserted into each of its I2CF and thereafter marks the I2CF represents the least elements as an active I2CF. To insert an element x , CCF first generates its fingerprint by mapping x into the range $[0, 2^f - 1]$. Then k independent hash functions map η_x onto the consistent hash ring. Based on the generated hash values, the consistent hashing determines the candidate buckets for η_x in the active I2CF. After that, we try to insert η_x into the active I2CF by following the strategy provided by the cuckoo hashing [4]. If the active I2CF can successfully store fingerprint η_x , the insertion algorithm will be terminated. Otherwise, CCF capacity has to be extended at either the bucket level or filter level. Thereafter, η_x will be inserted into the extended or added I2CF. The pseudo-code is shown in Algorithm 1. Note that, when extending, CCF marks the manipulated I2CF as active, so that the coming elements will be stored by this I2CF vector. We suggest choosing the I2CF with least buckets for better balance when bucket level extension is performed. Sometimes multiple buckets have to be added to successfully reside η_x . If there are still many elements to be inserted after x , CCF will

Algorithm 1: CCF insertion(η_x)

Input: The fingerprint to insert η_x

- 1 Calculate the hash values $h_1(\eta_x), \dots, h_k(\eta_x)$;
- 2 Decide the candidate buckets for η_x in the active I2CF based on $h_1(\eta_x), \dots, h_2(\eta_x)$;
- 3 **if** η_x can be successfully inserted into the active I2CF within max reallocations **then**
- 4 **return** True;
- 5 **else**
- 6 Extend CCF with its resizing strategy;
- 7 Insert η_x into the extended or added I2CF;
- 8 **return** True;

Algorithm 2: CCF query(x)

Input: The element to query x

- 1 $\eta_x = h_0(x) \bmod 2^f$;
- 2 Calculate the hash values $h_1(\eta_x), \dots, h_k(\eta_x)$;
- 3 **for** $i=0$ to $s-1$ **do**
- 4 Determine candidate buckets $B_i^1(x), \dots, B_i^k(x)$ in I2CF_{*i*};
- 5 **for** $j=0$ to $k-1$ **do**
- 6 **if** $B_i^j(x)$ has η_x **then**
- 7 **return** True;
- 8 **return** False;

introduce a new I2CF vector such that the coming elements will be stored immediately.

Query. Membership query with CCF may check every I2CF vector. Let s denote the number of I2CF vectors in CCF. We need to check $s \cdot k$ buckets in the worst case. Algorithm 2 presents a membership query in detail. The fingerprint η_x is hashed by k hash functions to determine the locations of η_x in the hash ring for I2CF_{*i*} ($i \in [0, s-1]$). Based on the hash values, the consistent hashing tells CCF the candidate buckets for η_x in I2CF_{*i*}. Then, if any bucket holds η_x , the membership query will be terminated and return true. By contrast, if η_x cannot be found in all I2CFs, CCF judges $x \notin A$ and returns false. There may be a potential false positive error for any queried element, but no false negative errors for the stored elements.

Deletion. The deletion of an element x needs to first perform a membership query for finding its possible locations. If a corresponding fingerprint η_x is found, then the matched fingerprint will be removed from CCF. Algorithm 3 shows the details of the delete operation. If the fingerprint η_x is not found in CCF, the deletion algorithm returns fail. When a sufficient number of elements have been deleted from CCF, the resizing operations will be executed to downsize CCF capacity and maintain high space utilization. CCF prefers filter-level resizing since a smaller s ensures a lower false positive rate.

D. Resizing of I2CF and CCF

An essential challenge for dynamic set representation is the unpredictable set cardinality n . This challenge puts forward new requirement for the employed data structure, i.e., the capability of capacity resizing. Moreover, the set cardinality n may vary irregularly, i.e., n may increase or decrease progressively or dramatically. To handle that, the data structure must be resized in diverse granularity. Therefore, we propose two options to extend the capacity of CCF, i.e., a scale-

Algorithm 3: CCF deletion(x)

Input: The element to delete x

- 1 $\eta_x = h_0(x) \bmod 2^f$;
- 2 Calculate the hash values $h_1(\eta_x), \dots, h_k(\eta_x)$;
- 3 **for** $i=0$ to $s-1$ **do**
- 4 Determine candidate buckets $B_i^1(x), \dots, B_i^k(x)$ in I2CF $_i$;
- 5 **for** $j=0$ to $k-1$ **do**
- 6 **if** $B_i^j(x)$ has η_x **then**
- 7 Remove η_x from $B_i^j(x)$;
- 8 Downsize the CCF when necessary;
- 9 **return** True;
- 10 **return** False;

Algorithm 4: CCF compact()

Input: The current CCF

- 1 success = True;
- 2 **while** success **do**
- 3 Select the least-loaded I2CF vector I2CF $_L$ in CCF;
- 4 Declare a new CCF named CCF $_T$;
- 5 Let CCF $_T = \text{CCF.remove}(\text{I2CF}_L)$;
- 6 **for** all η_x stored in I2CF $_L$ **do**
- 7 **if** !CCF $_T$.insertion(η_x) **then**
- 8 success = False;
- 9 **break**;
- 10 **if** success **then**
- 11 CCF = CCF $_T$;

up method which adds buckets into an I2CF, and a scale-out method which adds an untapped I2CF into CCF. Pairwisely, the CCF capacity can be downsized by either removing buckets from one specific I2CF or compacting sparse I2CFs. *Scale up* and *scale down* support the bucket-level capacity alterations, meanwhile, *scale out* and *compact* achieve the filter-level capacity adjustments. These methods generate ultimate elasticity for CCF when representing dynamic sets.

Scale up. When a new bucket is added into an I2CF, only the fingerprints stored in the bucket's successor may be affected. We consider that a new bucket B_{new} is mapped between two existing buckets B_i and B_j ($i, j \in [0, m-1]$), and B_j is the successor of B_{new} . In this case, only the fingerprints in B_j might have to be reallocated to bucket B_{new} . Specifically, if a fingerprint in B_j is mapped between B_i and B_{new} , it should be moved to B_{new} ; otherwise, it should still remain in B_j . Especially, if B_j is empty, B_{new} will also be empty. A toy example for adding bucket can be found in Fig. 1(b).

Scale down. Correspondingly, CCF can remove buckets from an I2CF for higher space utilization. When an existing bucket is removed from an I2CF, only the fingerprints in this bucket should be reinserted into the CCF. We consider two buckets B_i and B_j in the hash ring such that B_j is the successor of B_i . CCF tries to store the fingerprints in B_i by pushing them into bucket B_j preferentially and then reallocating the rest fingerprints to other buckets. If all the fingerprints are successfully stored, B_i will be removed; otherwise, B_i cannot be removed. An illustrative example for removing bucket is shown in Fig. 1(c). When scaling down, CCF prefers removing empty or under-utilized buckets for time-saving.

Scale out. Another method to increase the capacity of CCF is to add untapped I2CFs. Initially, CCF maintains a single I2CF and scale up or scale down this filter according to the real demand. When the number of elements to represent increases dramatically, the capacity of CCF can be extended immediately by adding one or multiple untapped I2CFs into the system. Note that, the added I2CFs are allowed to be heterogeneous since they are totally independent. The number of buckets and the number of slots are all mutable.

Compact. When an I2CF becomes sparse due to the removal of elements from the set, CCF tries to remove this I2CF through the compact operation. As shown in Algorithm 4, CCF first selects a least-loaded I2CF vector I2CF $_L$ and removes it. The updated CCF is denoted as CCF $_T$. Thereafter, we try to reinsert the fingerprints in I2CF $_L$ into CCF $_T$. If all the fingerprints in I2CF $_L$ can be successfully inserted into CCF $_T$, the selected I2CF $_L$ is allowed to be removed; otherwise, the CCF is already condensed enough and cannot be further compressed. The compact algorithm keeps removing I2CF vectors until an undeletable I2CF is reached.

In practice, the set cardinality n varies due to the join or removal of elements. When the value of n increases (decreases) gradually, CCF executes the scale up (scale down) algorithm to adaptively adjust its capacity. In the case of dramatic growth (reduction) of n , the scale out (compact) operation will be employed to extend (downsize) the CCF instantly. With these strategies, CCF ensures capacity elasticity and high space utilization simultaneously.

IV. PERFORMANCE ANALYSIS OF CCF

In this section, we first theoretically analyze the time-complexity of CCF. Then, we present a new method to calculate the threshold of the ratio between the number of elements to represent n_i and the number of buckets m_i for a given I2CF $_i$. Lastly, we offer an upper-bound probability of successfully inserting a given number of elements with a given I2CF $_i$.

A. Time-complexities of CCF

Theorem 2: Consider a CCF with s I2CFs and $k_0 = \dots = k_i = \dots = k_{s-1} = k$, $b_0 = \dots = b_i = \dots = b_{s-1} = b$. Let max and m denote the allowed reallocation times and the lengths of I2CFs. The I2CFs may have unequal lengths, for simplicity, we treat them as a uniform m . Then the time-complexities for CCF insertion, query and deletion are $O(max \cdot \log m)$, $O(s \cdot k \cdot b \cdot \log m)$ and $O(s \cdot k \cdot b \cdot \log m)$, respectively.

CCF introduces the consistent hashing to achieve capacity elasticity. Therefore the time-complexity of query and deletion is not constant anymore. Basically, whenever we need to know the indices of candidate buckets for an element, CCF has to refer to the underlying consistent hash ring. In a real implementation, the hash values of these buckets are organized as a binary search tree. Consequently, given a hash value of an element, the corresponding candidate bucket will be searched out in $O(\log m)$ time. To insert an element into the active I2CF, at most max reallocations is allowed, thus the time-complexity is $O(max \cdot \log m)$. As for a query and deletion, CCF has to go through all the I2CFs in the worst cases, therefore the resultant time-complexity is $O(s \cdot k \cdot b \cdot \log m)$.

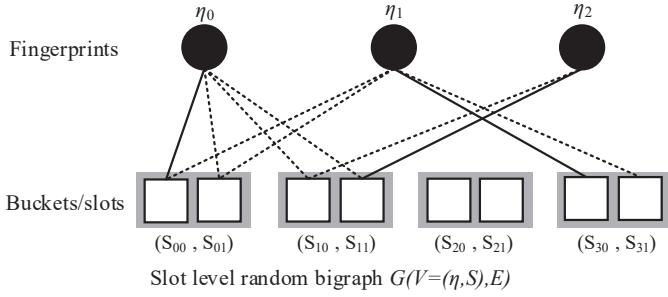


Fig. 2. An instance of slot level random bigraph for an I2CF_i with $m_i=4$, $k_i=2$, $b_i=2$ and $n_i=3$. A “mapping conflict” happens to η_2 since only bucket 1 is assigned to it. The three solid edges forms a complete matching, meaning that the three fingerprints can stored successfully.

Compared with DCF, the time-complexities of CCF are a little higher with an additional multiplicand of $\log m$. Logarithm level complexity, in fact, is acceptable in practice since the logarithmic function increases slowly when the value of m grows drastically. Distributed systems that employ the consistent hashing techniques all incur log level complexity. They are proved to function well. Examples include Cassandra [12], Akamai [13], Swift [14], Dynamo [15], etc.

B. Threshold for CCF Insertion

Each I2CF in CCF can be extended or downsized by adding or removing buckets dynamically. But for a static I2CF with given parameters, we need to explore how many fingerprints can be successfully inserted. That is, given the number of the elements to be represented n_i , a derivative problem is to seek a threshold T_i for the ratio between n_i and m_i . When $\frac{n_i}{m_i} \leq T_i$, I2CF_i can successfully store the n_i elements with probability $1-o(1)$; otherwise, I2CF_i may fail to record all the n_i elements with probability $1-o(1)$.

The mapping between elements and buckets in I2CF_i can be abstracted as a k_i -uniform hypergraph with m_i nodes and n_i hyperedges each of which is of fixed size k_i chosen independently from the m_i nodes. Based on the *core theory* of hypergraph, T_i can be derived out as a function of k_i and b_i . The details are given in literatures [16] and [17]. In fact, the resulted hypergraph may not be k_i -uniform since the k_i independent hash functions may select the same buckets from I2CF_i for an element x . We call this phenomenon “mapping conflict”. These mapping conflicts violate the k_i -uniform assumption towards the hypergraph. However, [16] and [17] did not consider the impact of the potential mapping conflicts. Therefore, in this paper, we present a novel new abstraction to I2CF and other Cuckoo filter-like data structures.

We notice that an I2CF_i can be naturally represented as a slot level random bipartite graph (or bigraph) $G(V=(\eta, S), E)$, where η and S denote the fingerprints to be stored and the slots in I2CF_i, respectively. As shown in Fig. 2, each slot has two subscripts which demonstrate its host bucket and its location in that bucket. For example, S_{01} means the second slot of the first bucket. In the bigraph, the edges demonstrate the assignment between the fingerprints and slots. If a bucket is a candidate bucket of a fingerprint, all slots of the bucket will have an edge to that fingerprint, to explicitly indicate that these slots can be employed to store that fingerprint. In the generated bigraph, a matching indicates a possible way to store these fingerprints.

Besides, this abstraction naturally provides us the important theorem in bigraphs, i.e., Hall’s Theorem [18].

Theorem 3: (Hall’s Theorem) Let $G(V=(X, Y), E)$ be a bigraph with bipartite sets X and Y . For a set of nodes $W \subseteq X$, let $N_G(W)$ denote the set of neighbors of W in G , i.e., the set of all nodes in Y which are adjacent to some element of W . There is a matching that entirely covers X if and only if for every subset W of X :

$$|W| \leq |N_G(W)|. \quad (3)$$

Additionally, given the parameters of I2CF_i as m_i , k_i , b_i and the number of fingerprints to be inserted n_i , we have the following observations.

Observation 1: For an insertion of an arbitrary element x , let Θ ($\Theta \in [0, k_i]$) be a variable and denote the total times that element x are mapped into a bucket. The value of Θ follows a typical binomial distribution since the employed hash functions are independent. Specifically, the probability that $\Theta = \theta$ can be calculated as:

$$p\{\Theta = \theta\} = \binom{k_i}{\theta} \left(\frac{1}{m_i}\right)^\theta \left(1 - \frac{1}{m_i}\right)^{(k_i - \theta)}. \quad (4)$$

Let p_0 denote the probability that an element x is mapped into the bucket. Since $\Theta \geq 1$ means x is mapped into the bucket, the value of p_0 can be derived out as:

$$p_0 = 1 - p\{\Theta = 0\} = 1 - \left(1 - \frac{1}{m_i}\right)^{k_i}. \quad (5)$$

Observation 2: Let $\Phi \in [0, n_i]$ be a variable and denote the total number of elements that mapped into a bucket. Then the value of Φ also follows a typical binomial distribution since the insertions of elements are independent. To be specific, the probability that $\Phi = \phi$ is:

$$p\{\Phi = \phi\} = \binom{n_i}{\phi} p_0^\phi (1 - p_0)^{(n_i - \phi)}. \quad (6)$$

By jointly considering the observations and Hall’s Theorem, we provide a new threshold \hat{T}_i for an I2CF_i. When $\frac{n_i}{m_i}$ is less than \hat{T}_i , the fingerprints can be stored successfully with high probability; by contrast, when $\frac{n_i}{m_i}$ is larger than \hat{T}_i , I2CF_i may fail to store some fingerprints with high probability.

Theorem 4: If $\Phi < b_i$, it is impossible for I2CF_i to store any fingerprint with the remained $b_i - \Phi$ slots. We consider the situations where $\Phi < b_i$ and derive out the fraction of space that may be utilized for a bucket as:

$$\hat{u} = 1 - \sum_{\phi=0}^{b_i-1} \left(1 - \frac{\phi}{b_i}\right) p\{\Phi = \phi\}, \quad (7)$$

where $1 - \frac{\phi}{b_i}$ is the fraction of bucket space which will never be utilized when $\Phi < b_i$. By contrast, if all the fingerprints are successfully stored, then the space utilization of I2CF_i is:

$$\bar{u} = \frac{n_i}{m_i \cdot b_i}. \quad (8)$$

Then, the threshold \hat{T}_i can be derived as the unique value of $\frac{n_i}{m_i}$ such that:

$$\hat{u} = \bar{u}. \quad (9)$$

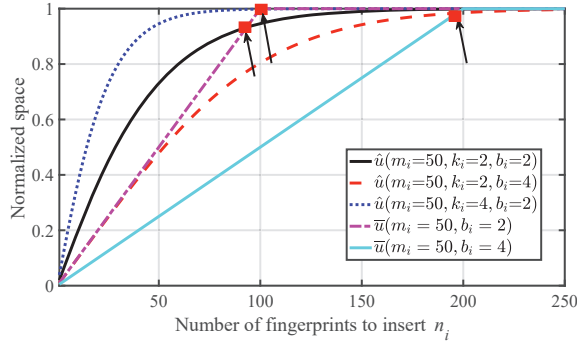


Fig. 3. The \hat{T}_i with diverse parameter settings. The value of \hat{T}_i can be derived out by dividing the x-axis value of the crossover point with b_i .

Theorem 4 can be proved by jointly considering observation 1, observation 2 and the Hall's Theorem. Intuitively, when $\frac{n_i}{m_i}$ is less than \hat{T}_i , \hat{u} is larger than \bar{u} , meaning there is sufficient space for the n_i fingerprints. As a result, when $\hat{u} \geq \bar{u}$, I2CF_i satisfies the requirement of Hall's Theorem with high probability. In contrast, if $\frac{n_i}{m_i} \geq \hat{T}_i$, $\hat{u} \leq \bar{u}$ and the space that can be utilized is not sufficient enough to accommodate the n_i fingerprints. In this situation, I2CF_i will fail to satisfy Hall's Theorem with high probability. As shown in Fig. 3, given $m_i=50$ and $b_i=2$, the value of \hat{T}_i grows significantly when k_i increases. Besides, with given m_i and k_i , the growth of b_i results in increasing of \hat{T}_i . Table II further presents the derived \hat{T}_i when m_i is taken as 2^{30} while k_i and b_i varies. This threshold provides a guide for the users of CCF and I2CF for their parameter settings in practice. Intuitively, larger \hat{T}_i guarantees higher space utilization. Therefore, given the same value of b_i , I2CF has the potentiality of realizing better space utilization than DCF by increasing the value of k_i .

C. Probability of Successful Representation

Theorem 4 and the literature [16] present the threshold of $\frac{n_i}{m_i}$ for a given I2CF. When $\frac{n_i}{m_i}$ is less than the threshold, the n_i fingerprints can be successfully stored with high probability. However, they fail to settle the derived question: what exactly is the probability of successfully inserting n_i fingerprints with a given I2CF, or with concession, what is the upper bound of that probability? We try to answer this question with the following observation.

Observation 3: For n_i given fingerprints, the number of edges in the maximum matching of the resultant $G(V=(\eta, S), E)$ implies the maximum number of fingerprints which can be successfully inserted into I2CF_i. If the maximum matching is a complete matching, then all the given fingerprints can be successfully stored by I2CF_i.

Note that the maximum matching in a specific bigraph can be solved by existing algorithms such as the Hungarian algorithm [19], Ford-Fulkerson algorithm [20], Hopcroft-Karp algorithm [21], etc. Let Ψ be a variable describing the number of successful inserted fingerprints in an I2CF_i with the parameters of m_i , n_i , k_i and b_i . A brute force method to calculate the probability distribution of Ψ is possible by exploring the probability space of the $G(V=(\eta, S), E)$ and then count those bigraphs in which the maximum matching contains a number of $\Psi=\psi$ edges. This method, however, suffers from exponentially growing time complexity since it has to test all

TABLE II
THE THRESHOLD \hat{T}_i FOR I2CF_i (WHEN $m_i = 2^{30}$).

$k_i \setminus b_i$	1	2	3	4
2	0.796812130	1.861790807	2.905683863	3.934728166
3	0.940479791	1.979049536	2.992264312	3.997079786
4	0.980172599	1.996604114	2.999390447	3.999888473
5	0.993022846	1.999453835	2.999955482	3.999996295
6	0.997483538	1.999913939	2.999996938	3.999999888
7	0.999082240	1.999986694	2.999999798	3.999999996

the $m_i^{n_i \cdot k_i}$ possible bigraphs. Therefore, alternatively, we derive out an upper bound of $p\{\Psi=n_i\}$ ($n_i \in [1, m_i \cdot b_i]$) based on the Hall's Theorem [18] and observation 3.

Theorem 5: For a given I2CF_i with m_i , k_i , b_i and n_i ($n_i \in [1, m_i \cdot b_i]$) fingerprints to be inserted, the probability that all the n_i fingerprints can be successfully accommodated has the following upper bound:

$$p\{\Psi=n_i\} \leq \sum_{j=\lfloor n_i/b_i \rfloor}^{\max\{n_i \cdot k_i, m_i\}} p\{\Omega=j\}, \quad (10)$$

where $p\{\Omega=j\}$ denotes the probability that the n_i fingerprints are mapped into exactly j buckets of I2CF_i. And $p\{\Omega=j\}$ can be calculated as:

$$p\{\Omega=j\} = \frac{\binom{m_i}{j} \sum_{l=0}^{F(j, n_i, k_i)} \left[D_l \prod_{r=0}^{j-1} \binom{n_i - \sum_{q=0}^{r-1} Q[l][q]}{Q[l][r]} \right]}{m_i^{n_i \cdot k_i}}, \quad (11)$$

where Q is an array of vectors each of which has j positive integers and the sum of these integers is exactly $n_i \cdot k_i$. The number of vectors in Q is denoted as $F(j, n_i, k_i)$ and can be calculated with the input value of j , n_i and k_i . D_l is the number of possible combinations of the j integers in $Q[l]$. The factor $\prod_{r=0}^{j-1} \binom{n_i - \sum_{q=0}^{r-1} Q[l][q]}{Q[l][r]}$ counts all possible cases when the $n_i \cdot k_i$ mappings are distributed into the selected j buckets according to the distribution given by $Q[l]$.

Theorem 5 can be proved by considering both observation 3 and the Hall's Theorem. Basically, $p\{\Omega=j\}$ only counts the probability that the n_i fingerprints are mapped into exactly j buckets, but fails to consider the situation where a subset of the n_i fingerprints may not satisfy Hall's Theorem. Therefore, Equ (10) offers the upper bound of $p\{\Psi=n_i\}$. Equ (11) can be derived out by regarding each mapping as a ball then formulating it as a typical balls and bins problem.

We give a walk-through example by calculating the upper bound of $p\{\Psi=3\}$ with $m_i=5$, $b_i=2$, $k_i=2$ and $n_i=3$, respectively. According to Equ (10), we have $p\{\Psi=3\} \leq p\{\Omega=2\} + p\{\Omega=3\} + p\{\Omega=4\} + p\{\Omega=5\}$. Then, according to Equ (11), $p\{\Omega=2\}=0.03968$, $p\{\Omega=3\}=0.3456$, $p\{\Omega=4\}=0.4992$, and $p\{\Omega=5\}=0.1152$. Therefore, the upper bound for $p\{\Psi=3\}$ is 0.99968. When calculating $p\{\Omega=3\}$, we have $n_i k_i = 6 = 1+1+4 = 1+2+3 = 2+2+2$. Thus $F(j, n_i, k_i) = 3$, $Q = \{[1, 1, 4], [1, 2, 3], [2, 2, 2]\}$, $D_0 = 3$ since $[1, 1, 4]$ has three permutations, i.e., $\{1, 1, 4\}$, $\{1, 4, 1\}$, and $\{4, 1, 1\}$, $D_1 = 6$, and $D_2 = 1$. Consequently, $p\{\Omega=3\} = \left[\binom{5}{3} \cdot (3 \cdot \binom{6}{1} \binom{5}{1} + 6 \cdot \binom{6}{1} \binom{5}{2} + \binom{6}{2} \binom{4}{2}) \right] / 5^6 = 5400 / 15625 = 0.3456$.

With the above analysis, we provide a better understanding of the proposed data structures, as well as a guide to the potential users about the parameter settings whenever CCF or I2CF are within their considerations.

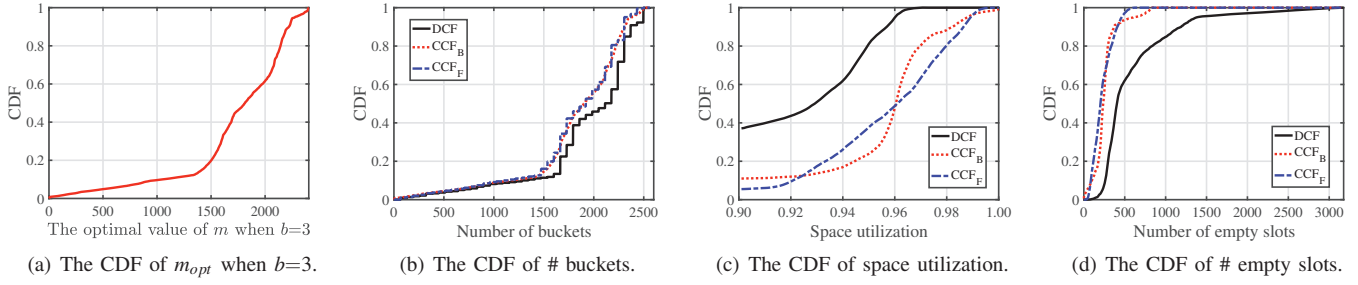


Fig. 4. The comparison between CCF and DCF with Yahoo! trace.

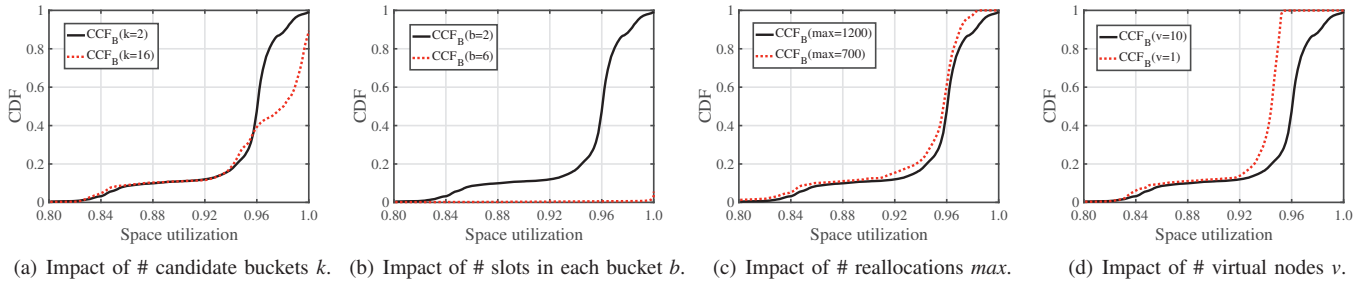


Fig. 5. The impact of different parameters in CCF.

V. EVALUATION

In this section, we compare the performance of CCF with DCF for dynamic set representation and then quantify the impact of the parameters. All of the experiments are conducted on a host with 8 GB RAM and 3.4 GHz CPU. Especially, we conduct the evaluations based on the real network flow trace from Yahoo! [22]. The Yahoo! trace records the basic information for each flow in its 6 distributed data centers, including the IP addresses of both source and destination servers, the arriving and terminating time of each flow, etc. In this paper, we regard the combination of the source and destination IP addresses as the content of an element, coupled with the start and end time. Since there are millions of flows in the trace, we only fetch the flows within 20 minutes. There are in total 58,941 flows and Fig. 4(a) shows the CDF of the optimal number of buckets m_{opt} . We have $m_{opt} = \lceil n/b \rceil$, since each bucket can store b fingerprints at most.

A. Comparison with DCF

We implement two versions of CCF, i.e., CCF with only bucket-level alteration CCF_B and CCF with only filter-level alteration CCF_F. For fairness, the parameters for CCF_F and DCF are set as the same, i.e., $m_i=64$, $b_i=3$, $k_i=2$, $f=30$. For both CCF_B and CCF_F, the value of M and the number of virtual nodes in the consistent hash ring v are given as 5×10^{10} and 10 respectively. Fig. 4(b), (c) and (d) depict the CDF of the resultant number of buckets, space utilization and the number of empty slots, respectively.

By jointly considering Fig. 4(a) and (b), we characterize the capacity elasticity of DCF and CCF. Obviously, CCF_B achieves the best elasticity and keeps capacity up and down whenever the number of elements increases or decreases. The curve of CCF_B in Fig. 4(b) matches the variation of m_{opt} in Fig. 4(a) perfectly. The CCF_F also responds to the changes of m_{opt} rapidly by executing its compact and scale out algorithms dynamically. The DCF, however, fails to

compact under-utilized CFs instantly when the value of m_{opt} decreases. The reason is that DCF moves any fingerprint in the under-utilized CF into its corresponding buckets of other CFs. Therefore, successful compaction is hard to achieve. Our CCF_F, in contrast, always tries to insert the fingerprints in an under-utilized I2CF into other I2CFs, thereby releasing the fingerprints from their locations in the under-utilized I2CF. Thus, both CCF_B and CCF_F have better elasticity than DCF.

Moreover, the CDF of space utilization is depicted in Fig. 4(c). For DCF, about 37% percent of resultant space utilization is less than 0.90. However, CCF_B and CCF_F have less than 10% percent of results which are below 0.90. Moreover, the maximum space utilization for DCF is 0.970, which is much lower than that of CCF_B (1.0) and CCF_F (0.999). To be accurate, on average, the space utilization for DCF, CCF_B and CCF_F are 0.8809, 0.9481 and 0.9425, respectively. Correspondingly, the CDF of the number of empty slots is shown in Fig. 4(d). For CCF_B and CCF_F, 93% and 97% results have less than 500 empty slots, while that value for DCF is 62% only. In the worst case, DCF remains 3,176 slots empty. And more than 16% DCF results incur more than 1,000 empty slots. The reason is that DCF can only compact an under-utilized filter if all the stored fingerprints find their corresponding unfilled buckets in other CFs. As a result, when the value of n decreases but DCF may fail to recycle under-utilized CFs timely. Notice that CCF_B has more empty slots than CCF_F. The reason is that we only try to merge the buckets which store less than 2 fingerprints in our experiments. At the end of our tests, due to the removals of flows, the proportion of buckets which accommodate 2 fingerprints gets higher, while CCF_B doesn't recycle the empty slots immediately.

From the above experiments, we conclude that CCF achieves better capacity elasticity and higher space utilization than DCF. The feature of design flexibility, on the other hand, may not be quantified directly. Intuitively, DCF only adds or merges homogenous CFs, while the I2CFs in CCF is allowed

to have diverse parameter settings. This flexibility makes CCF more suitable for dynamic set representation than DCF.

B. Impact of Parameter Settings

In this subsection, we quantify the impact of parameters for CCF. Especially, we consider four main parameters, i.e., the number of candidate buckets k , the number of slots in a bucket b , the maximum reallocations max , and the number of virtual nodes in the consistent hash ring v . Note that, we evaluate the space utilization of CCF_B when the above parameters vary. The reference is set as a CCF_B with $k=2$, $b=3$, $max=1200$, and $v=10$. We then vary the four parameters separately and show the results in Fig. 5.

As shown in Fig. 5(a), when k is increased from 2 to 16, CCF achieves higher space utilization (rising from 0.9481 to 0.9599 on average). When $k=16$, nearly half of the results have more than 0.98 space utilization. However, less than 12% percent of the results achieve more than 0.98 space utilization when $k=2$. An element has more candidate buckets implies that a bucket may be assigned to more elements. Consequently, the probability that a bucket is assigned to less than b elements gets lower, which leads to higher space utilization. When b is increased from 3 to 6, we can see from Fig. 5(b) that the space utilization increases dramatically. To be specific, the space utilization is 0.9481 for $b=3$, but 0.9986 for $b=6$ on average. This phenomenon is reasonable since, with a larger b , there are fewer buckets in the CCF_B . In the Yahoo! trace, the maximum number of flows to store is about 7,290. So $max=1,200$ means the reallocations when inserting an element may cover the whole filter to explore potential empty slots. Also, with less number of buckets in the filter, the probability of buckets which is assigned to less than b elements gets lower. Accordingly, the resultant space utilization increases.

When the value of max is decreased from 1200 to 700, the CDF of the resultant space utilization is recorded in Fig. 5(c). Obviously, with more allowed reallocations, CCF achieves higher space utilization. The reason is that, with larger max , the insertion will search more buckets, and hopefully CCF may find an empty slot to store the fingerprint. Moreover, as depicted in Fig. 5(d), when the number of virtual nodes in the consistent hash ring decreases from 10 to 1, the space utilization experiences a significant drop (decreasing from 0.9481 to 0.9298 on average). When $v=1$, only about 16% percent of the results realize more than 0.95 space utilization. By contrast, when $v=10$, about 76% percent of the results realize more than 0.95 space utilization. Basically, with more virtual nodes, consistent hashing generates better load balance among the buckets. Therefore, the probability that a bucket is assigned with less than b elements gets lower and thereby resulting in higher space utilization.

From the above results, we conclude that the parameters of CCF have diverse impacts on its performance. The users can customize their own configurations to achieve their goals by leveraging these parameters.

VI. CONCLUSION

In this paper, we present the CCF design for dynamic set representation and membership query, with the targets of capacity elasticity, space efficiency, and design flexibility. CCF is composed of an adjustable number of I2CFs. At its core,

each I2CF enables bucket-level capacity alteration with the using of consistent hashing. At the filter level, CCF resizes its capacity by adding untapped I2CFs or merging under-utilized I2CFs adaptively. Without any inner dependency and constraints, all the parameters of CCF are mutable and can be customized by its users. Theoretical analysis and trace-driven experiments show that CCF outperforms DCF and achieves the design rationales simultaneously at the cost of a little higher time-complexity.

REFERENCES

- [1] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Communications of the ACM*, vol. 13, no. 7, pp. 422-426, 1970.
- [2] S. Tarkoma, C.E. Rothenberg, and E. Lagerspetz, "Theory and practice of Bloom filters for distributed systems," *IEEE Communications Surveys and Tutorials*, vol.14, no. 1, pp.131-155, 2012.
- [3] L. Luo, D. Guo, R. Ma, O. Rottenstreich, and X. Luo, "Optimizing Bloom filter: Challenges, solutions, and comparisons," *IEEE Communications Surveys and Tutorials*, accepted to appear, 2018.
- [4] B. Fan, D. Andersen, M. Kaminsky, and M. Mitzenmacher, "Cuckoo filter: Practically better than bloom," in *ACM CoNEXT*, Sydney, Australia, 2-5 December, 2014.
- [5] D. Guo, J. Wu, H. Chen, Y. Yuan, and X. Luo, "The dynamic Bloom filters," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 1, pp. 120-133, 2010.
- [6] H. Chen, L. Liao, H. Jin, and J. Wu, "The dynamic cuckoo filter," in *IEEE ICNP*, Toronto, Canada, 10-13 October, 2017.
- [7] D. Eppstein, "Cuckoo filter: Simplification and analysis," *arXiv preprint*, arXiv:1604.06067, 2016.
- [8] M. Mitzenmacher, S. Pontarelli, and P. Reviriego, "Adaptive Cuckoo filters," *arXiv preprint*, arXiv:1704.06818, 2017.
- [9] B. Fan, D. G. Andersen, and M. Kaminsky, "MemC3: Compact and concurrent memcache with dumber caching and smarter hashing," in *USENIX NSDI*, Lombard, IL, USA, 2-5 April, 2013.
- [10] D.Karger, E. Lehman, F. Leighton, R. Panigrahy, M. Levine, and D. Lewin, "Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the World Wide Web," in *ACM STOC*, Texas, USA, 4-6 May, 1997.
- [11] I. Stoica, R. Morris, D. Liben-Nowell, D. Karger, M. Kaashoek, F. Dabek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup protocol for internet applications," *IEEE/ACM Transactions on Networking*, vol. 11, no. 1, pp. 17-32, 2003.
- [12] A. Lakshman, and P. Malik, "Cassandra: A decentralized structured storage system," *ACM SIGOPS Operating Systems Review*, vol. 44, no.2, pp. 35-40, 2010.
- [13] L. Wang, V. Pai, and L. Peterson, "The effectiveness of request redirection on CDN robustness," *ACM SIGOPS Operating Systems Review*, vol. 36, no. SI, pp. 345-360, 2002.
- [14] T. Chekam, E. Zhai, Z. Li, Y. Cui, and Kui Ren, "On the synchronization bottleneck of OpenStack Swift-like cloud storage systems," in *IEEE INFOCOM*, San Francisco, CA, USA, 11-15 April, 2016.
- [15] G. DeCandia, D. Hastorun, M. Jampani, G.Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels, "Dynamo: Amazon's highly available key-value store," in *ACM SOSP*, Stevenson, Washington, USA, 14-17 October, 2007.
- [16] M. Dietzfelbinger, A. Goerd, M. Mitzenmacher, A. Montanari, R. Pagh, and M. Rink, "Tight thresholds for cuckoo hashing via XORSAT," in *ICALP*, Bordeaux, France, 5-10 July, 2010.
- [17] N. Fountoulakis, M. Khosla, and K. Panagiotou, "The multiplicity-orientability thresholds for random hypergraphs," in *ACM-SIAM SODA*, San Francisco, CA, USA, 23-25 January, 2011.
- [18] P. Hall, "On representatives of subsets," *Journal of the London Mathematical Society*, vol. 1, no. 1, pp. 26-30, 1935.
- [19] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 1, no. 1-2, pp. 83-97, 1955.
- [20] L. Ford and D. Fulkerson, "Flows in networks," *Princeton University Press*, 1962.
- [21] J. E. Hopcroft and R. M. Karp, "An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs," *SIAM Journal on Computing*, vol. 2, no.4, pp. 225-231, 1973.
- [22] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, "DCTCP: Efficient packet transport for the commoditized data center," in *ACM SIGCOMM*, New Delhi, India, 30 August-3 September, 2010.