

# Nearly Optimal Static Las Vegas Succinct Dictionary

Huacheng Yu\*

## Abstract

Given a set  $S$  of  $n$  (distinct) keys from key space  $[U]$ , each associated with a value from  $\Sigma$ , the *static dictionary* problem asks to preprocess these (key, value) pairs into a data structure, supporting value-retrieval queries: for any given  $x \in [U]$ ,  $\text{valRet}(x)$  must return the value associated with  $x$  if  $x \in S$ , or return  $\perp$  if  $x \notin S$ . The special case where  $|\Sigma| = 1$  is called the *membership* problem. The “textbook” solution is to use a hash table, which occupies linear space and answers each query in constant time. On the other hand, the minimum possible space to encode all (key, value) pairs is only  $\text{OPT} := \lceil \lg_2 \binom{U}{n} + n \lg_2 |\Sigma| \rceil$  bits, which could be much less.

In this paper, we design a randomized dictionary data structure using

$$\text{OPT} + \text{poly } \lg n + O(\lg \lg \lg \lg U)$$

bits of space, and it has *expected constant* query time, assuming the query algorithm can access an external lookup table of size  $n^{0.001}$ . The lookup table depends only on  $U$ ,  $n$  and  $|\Sigma|$ , and not the input. Previously, even for membership queries and  $U \leq n^{O(1)}$ , the best known data structure with constant query time requires  $\text{OPT} + n/\text{poly } \lg n$  bits of space (Pagh [Pag01a] and Pătraşcu [Păt08]); the best known using  $\text{OPT} + n^{0.999}$  space has query time  $O(\lg n)$ ; the only known non-trivial data structure with  $\text{OPT} + n^{0.001}$  space has  $O(\lg n)$  query time and requires a lookup table of size  $\geq n^{2.99}$  (!). Our new data structure answers open questions by Pătraşcu and Thorup [Păt08, Tho13].

We also present a scheme that compresses a sequence  $X \in \Sigma^n$  to its zeroth order (empirical) entropy up to  $|\Sigma| \cdot \text{poly } \lg n$  extra bits, supporting decoding each  $X_i$  in  $O(\lg |\Sigma|)$  expected time.

---

\*Department of Computer Science, Princeton University. yuhch123@gmail.com

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Our contributions . . . . .	2
1.2	Related work . . . . .	3
1.3	Technical contributions . . . . .	4
1.3.1	Strings with fractional lengths . . . . .	4
1.3.2	Data interpretation . . . . .	6
1.4	Overview of Pătraşcu’s data structure . . . . .	6
<b>2</b>	<b>Overview</b>	<b>7</b>
2.1	Random inputs . . . . .	8
2.2	Using data interpretation . . . . .	9
2.3	Constructing data interpretation . . . . .	10
2.4	Worst-case input . . . . .	11
2.5	Organization . . . . .	12
<b>3</b>	<b>Preliminaries and Notations</b>	<b>12</b>
3.1	Random access machine . . . . .	12
3.2	Notations . . . . .	12
<b>4</b>	<b>Fractional-length Strings</b>	<b>13</b>
<b>5</b>	<b>Reductions to Perfect Hashing</b>	<b>16</b>
<b>6</b>	<b>Perfect Hashing for Medium-Sized Sets</b>	<b>17</b>
6.1	No bad block pair . . . . .	19
6.2	At least one bad block pair . . . . .	23
6.3	Final data structure for medium size sets . . . . .	28
<b>7</b>	<b>Data Structure Pair for Block Pair</b>	<b>31</b>
7.1	Data interpretation . . . . .	39
7.2	Small sets . . . . .	43
<b>8</b>	<b>Perfect Hashing for Sets of Any Size</b>	<b>55</b>
<b>9</b>	<b>Discussions and Open Problems</b>	<b>59</b>
<b>A</b>	<b>Proofs for Fractional-length Strings</b>	<b>62</b>
<b>B</b>	<b>Approximating Binomial Coefficients</b>	<b>72</b>
<b>C</b>	<b>Dictionary with Linear Redundancy</b>	<b>73</b>

# 1 Introduction

Given  $n$  (key, value) pairs  $\{(k_i, v_i)\}_{i=1, \dots, n}$  with distinct keys  $k_i \in \{0, \dots, U-1\}$  and (possibly duplicated) values  $v_i \in \{0, \dots, \sigma-1\}$ , the *static dictionary* problem asks to preprocess them into a data structure, supporting value-retrieval queries

- $\text{valRet}(x)$ : return  $v_i$  if  $x = k_i$ , and return  $\perp$  if  $x \neq k_1, \dots, k_n$ .

When  $\sigma = 1$ , it is called the *membership* problem, i.e., preprocessing a set  $S$  of  $n$  keys into a data structure, supporting queries of form “is  $x \in S$ ?”

Dictionaries are very fundamental data structures, which have been extensively studied in theory [CW79, TY79, Yao81, FKS84, FNSS92, FN93, BM99, Pag01a, Pag01b, FM95, Mil96, MNSW98, BMRV02]. They are also one of the most basic data structures in practice, included in standard libraries for most of the popular programming languages (e.g., `std::unordered_map` for C++, `HashMap` for Java, and language-level built-in support for JavaScript, Python, Ruby, etc).

The “textbook” implementation of a dictionary is to store a hash table: use a hash function to map all keys to  $O(n)$  buckets, and store each (key, value) pair in the corresponding bucket. Simple hash functions (e.g.  $(kx \bmod p) \bmod n$ ) have low collision probabilities, and resolving collisions by chaining leads to a dictionary data structure with *expected* constant query time. Using perfect hashing (e.g. [FKS84]), one can further improve the query time to *worst-case* constant. However, such data structures use at least  $n \lg U + n \lg \sigma$  bits of space, even just to write down all (key, value) pairs in the buckets, while the information theoretical space lower bound for this problem is only<sup>1</sup>

$$\text{OPT} := \lceil \lg \binom{U}{n} + n \lg \sigma \rceil$$

bits, which is much less than  $n \lg U + n \lg \sigma$  (note that  $\lg \binom{U}{n} = n \lg(U/n) + O(n)$ ).

It turns out that it is possible to not *explicitly* store all pairs, and beat  $n \lg U + n \lg \sigma$  bits. For *membership* queries ( $\sigma = 1$ ), the previously best known data structure by Pagh [Pag01a] (and later improved by Pătraşcu [Păt08]) uses  $\text{OPT} + O(n/\text{poly} \lg n + \lg \lg U)$  bits of space, and answers queries in constant time. This data structure also gives a smooth tradeoff between time and space: for query time  $O(t)$ , it uses space

$$\text{OPT} + n/r + O(\lg \lg U),$$

where  $r = (\frac{\lg n}{t})^{\Omega(t)}$ . To achieve this query time, it is assumed that the query algorithm has access to an external lookup table of size  $\min\{n^3, r^6\}$ , which depends only on  $U$  and  $n$ , and not the input. In particular, when  $U = \text{poly } n$ , if the number of extra bits is  $n^{0.99}$ , the query time becomes  $O(\lg n)$ . If we want the space to be very close to  $\text{OPT}$ , the query time is  $O(\lg n)$ , but the lookup table size becomes about  $n^3$  (it may even be larger than the data structure itself). For  $\sigma > 1$ , only  $(\text{OPT} + O(n + \lg \lg U))$ -bit data structures were known [Pag01a]. While these data structures have deterministic query algorithms (and worst-case query-time guarantee), no better zero-error randomized data structure was known. To the best of our knowledge, data structures with Las Vegas query algorithms have never been the state-of-the-art for this problem since perfect hashing [FKS84].<sup>2</sup> Therefore, it was unclear if randomization is even useful for static dictionaries.

<sup>1</sup>Throughout the paper,  $\lg$  is the binary logarithm.

<sup>2</sup>Monte Carlo algorithms, where the query is allowed to err with a small probability, would have a different space benchmark. Thus, they are not the focus of this paper.

## 1.1 Our contributions

In this paper, we show that if we allow randomization, near-optimal space and optimal time can be achieved simultaneously. We design a dictionary data structure with  $\text{poly } \lg n + O(\lg \lg U)$  extra bits and expected constant query time, making a step towards the optimal static dictionary. The query algorithm only needs to access a small lookup table.

**Theorem 1.** *There is a randomized algorithm that preprocesses  $n$  (key, value) pairs into a data structure with*

$$\text{OPT} + \text{poly } \lg n + O(\lg \lg U)$$

*bits, such that for any given query  $x$ , the query algorithm answers  $\text{valRet}(x)$  in expected constant time on a random access machine with word-size  $w \geq \Omega(\lg U + \lg \sigma)$ , assuming it can access an external lookup table of size  $n^\epsilon$ , for any constant  $\epsilon > 0$ .*

Same as the previous data structures, the lookup table depends only on  $U$ ,  $n$  and  $\sigma$ , and not the input. In fact, the  $\lg \lg U$  term can be improved to  $\lg \lg \cdots \lg U$  for logarithm iterated any constant number of times. Hence, when  $U$  is at most  $2^{2^{\cdots 2^n}}$  with  $O(1)$  levels, this term can be removed. In this case, among the  $\text{OPT} + \text{poly } \lg n$  bits of the data structure, the first  $\text{poly } \lg n$  are the (plain) random bits used by the preprocessing algorithm, and the “actual” data structure only occupies the next (and last)  $\text{OPT} + 1$  bits. The expectation of the query time is taken over these random bits, which we assume the worst-case input data and query do not see. Moreover, the query time is  $O(1)$  with probability  $1 - o(1)$ , and is  $\text{poly } \lg n$  in worst-case.

By storing the lookup table as part of the data structure, Theorem 1 implies a data structure with  $\text{OPT} + n^\epsilon + O(\lg \lg U)$  bits and expected constant query time, which is still an improvement over the previous best known. In the *cell-probe* model, where we only count how many times the query algorithm accesses the memory and the computation is free, the lookup table is not necessary, because it does not depend on the input and can be computed without accessing the data structure.

In the theorem, we assume that each (key, value) pair fits in  $O(1)$  words, which is necessary to obtain constant query time on random access machines. We will discuss larger keys or values in Section 9.

**Perfect hashing.** In general, a perfect hashing maps  $n$  input keys to distinct buckets, and it is called *minimal* if it maps them to exactly  $n$  distinct buckets, labeled from 0 to  $n - 1$ . En route to the new dictionary data structure, the key component is a succinct membership data structure, equipped with a *two-sided* minimal perfect hashing. Specifically, given a set  $S \subseteq [U]$  of size  $n$ , we would like to construct a data structure  $\mathcal{D}$ , which not only supports membership queries, but also defines

- a bijection  $h$  between  $S$  and  $[n]$ , and
- a bijection  $\bar{h}$  between  $[U] \setminus S$  and  $[U - n]$ .

That is, we want to perfectly hash all keys, as well as all non-keys. The data structure must support  $\text{hash}(x)$  queries, which returns a pair  $(b, v)$  such that

- if  $x \in S$ ,  $b = 1$  and  $v = h(x)$ ;
- if  $x \notin S$ ,  $b = 0$  and  $v = \bar{h}(x)$ .

**Theorem 2 (informal).** *There is a randomized algorithm that preprocesses a set  $S \subseteq [U]$  of size  $n$  into a data structure with*

$$\lg \binom{U}{n} + \text{poly } \lg n + O(\lg \lg U)$$

bits, such that it defines a bijection  $h$  between  $S$  and  $[n]$  and a bijection between  $[U] \setminus S$  and  $[U - n]$ . For any  $x$ , the query algorithm outputs  $\text{hash}(x)$  in expected constant time on a random access machine with word-size  $w \geq \Omega(\lg U)$ , assuming it has access to an external lookup table of size  $n^\epsilon$ , for any constant  $\epsilon > 0$ .

See Section 8 for the formal statement.

**Locally decodable arithmetic codes.** We also show that the above perfect hashing data structure can be applied to obtain a version of locally decodable arithmetic codes with a better space [Păt08]. This problem asks to compress a sequence  $X = (x_1, \dots, x_n) \in \Sigma^n$  for some (small) alphabet set  $\Sigma$ , such that each symbol  $x_i$  can be recovered efficiently from the compression. We should think of a sequence  $X$  that is sampled from some low entropy source, and the encoding should take much less than  $n \lg |\Sigma|$  bits. The arithmetic codes match the zeroth order entropy of  $X$ , i.e., if each symbol in the sequence has entropy  $H$  (marginally), then the encoding has length  $\sim n \cdot H$ . Pătraşcu [Păt08] gave a data structure whose size is

$$\sum_{\sigma \in \Sigma} f_\sigma \lg \frac{n}{f_\sigma} + O(|\Sigma| \lg n) + n / \left( \frac{\lg n}{t} \right)^t + \tilde{O}(n^{3/4}),$$

where  $f_\sigma$  is the number of occurrences of  $\sigma$ . It supports single-element access in  $O(t)$  time on word RAM. Note that when each symbol  $x_i$  is sample independent from a source of entropy  $H$ , then the empirical entropy  $\sum_{\sigma \in \Sigma} f_\sigma \lg \frac{n}{f_\sigma} \approx n \cdot H$  with high probability.

**Theorem 3.** *There is a randomized algorithm that preprocesses a sequence  $(x_1, \dots, x_n) \in \Sigma^n$  into a data structure with*

$$\lg \left( \frac{n!}{f_{\sigma_1}! f_{\sigma_2}! \dots} \right) + |\Sigma| \cdot \text{poly } \lg n$$

*bits, where  $f_\sigma$  is the number of occurrences of  $\sigma$ . For any index  $i$ , the query algorithm recovers  $x_i$  in  $O(\lg |\Sigma|)$  time in expectation on a word RAM with word-size  $w \geq \Omega(\lg n)$ , assuming it has access to an external lookup table of size  $n^\epsilon$ , for any constant  $\epsilon > 0$ .*

Note that the first term in the space is the minimum possible space to store a sequence with frequencies  $(f_\sigma)_{\sigma \in \Sigma}$ , which is at most  $\sum_{\sigma} f_\sigma \lg \frac{n}{f_\sigma}$ .

## 1.2 Related work

The perfect hashing scheme by Fredman, Komlós and Szemerédi [FKS84] maps  $[U]$  to  $[n]$  such that for any given set  $S$  of size  $n$ , there is a hash function  $h$  that maps all elements in  $S$  to different buckets (i.e., no hash collision) such that  $h(x)$  can be evaluated in constant time. This hash function takes  $O(n\sqrt{\lg n} + \lg \lg U)$  bits to store, and it is later improved to  $O(n + \lg \lg U)$  bits by Schmidt and Siegel [SS90]. By storing the (key, value) pair in the corresponding bucket, the perfect hashing scheme solves the dictionary problem with  $O(n)$  words of space and constant query time. Fiat, Naor, Schmidt and Siegel [FNSS92] showed that only  $O(\lg n + \lg \lg U)$  extra bits are needed to store *both* the hashing function and the table, obtaining space of  $n \lceil \lg U \rceil + n \lceil \lg \sigma \rceil + O(\lg n + \lg \lg U)$ . Fiat and Naor [FN93] further removed the  $O(\lg n)$  term, as well as the  $O(\lg \lg U)$  term when  $U$  is not too large.

The first dictionary data structure that achieves nearly optimal space is due to Brodnik and Munro [BM99]. Their data structure uses  $\text{OPT} + O(\text{OPT} / \lg \lg \lg U)$  bits, and it has constant query time. Pagh [Pag01a] reduced the dictionary problem to the *rank* problem (see below, also Section 1.4 for definition of the rank

problem), and designed a data structure for membership queries using  $\text{OPT} + O(n \lg^2 \lg n / \lg n + \lg \lg U)$  bits for  $n < U / \lg \lg U$ , and  $\text{OPT} + O(U \lg \lg U / \lg U)$  for  $n \geq U / \lg \lg U$ . Pagh’s dictionary uses rank data structures as subroutines. By improving the rank data structures, Pătraşcu [Păt08] improved the bound to  $\text{OPT} + n / \text{poly} \lg n + O(\lg \lg U)$ , as we mentioned earlier. Such data structures using  $\text{OPT} + o(\text{OPT})$  bits are called *succinct* data structures [Jac89], where the number of extra bits  $o(\text{OPT})$  is called the *redundancy*.

It is worth mentioning that when  $n = U$ , i.e., when the input is a sequence of values  $v_1, \dots, v_U \in [\sigma]$ , Dodis, Pătraşcu and Thorup [DPT10] designed a data structure using optimal space. Their data structure uses a lookup table of  $\text{poly} \lg n$  size. We get rid of this lookup table (see Lemma 16 in Section 7) as an application of our new technique.

No non-trivial lower bounds are known without restrictions on the data structure or model. Fich and Miltersen [FM95] and Miltersen [Mil96] proved  $\Omega(\lg n)$  and  $\Omega(\lg \lg n)$  lower bounds in the RAM model with restricted operations. Buhrman, Miltersen, Radhakrishnan and Venkatesh [BMRV02] proved that in the *bit-probe* model (where the word-size  $w = 1$ ), any data structure using  $O(\text{OPT})$  space must have query time at least  $O(\lg \frac{U}{n})$ . Viola [Vio12] proved a higher lower bound for the case where  $U = 3n$ , that any bit-probe data structure with query time  $q$  must use space  $\text{OPT} + n/2^{O(q)} - \log n$ .

Raman, Raman and Rao [RRR07] considered the *indexable dictionary* problem, which generalizes membership. Given a set  $S$  of  $n$  keys, it supports rank and select queries:  $\text{rank}_S(x)$  returns  $\perp$  if  $x \notin S$ , and returns  $i$  if  $x$  is the  $i$ -th smallest in  $S$ ;  $\text{select}_S(i)$  returns the  $i$ -th smallest element in  $S$ . They obtained a data structure using  $\text{OPT} + o(n) + O(\lg \lg U)$  bits and constant query time. Grossi, Orlandi, Raman and Rao [GORR09] studied the *fully indexable dictionary* problem. It generalizes the indexable dictionary problem to let  $\text{rank}_S(x)$  return the number of elements in  $S$  that are at most  $x$  (also for  $x \notin S$ ). They obtained a data structure using space  $\text{OPT} + n^{1+\epsilon} + U^\delta \cdot n^{-1/\epsilon}$ . In fact, this problem is much harder. It was observed in [PV10] that rank queries can be reduced from *colored predecessor search*, which has a query time lower bound of  $\Omega(\lg \lg n)$  even when the space is  $O(n \lg U)$  [PT06, PT07] (not to say the succinct regime). When  $U > n^2$ , the problem requires  $n^{1+\epsilon}$  space to get constant query time (when the word-size is  $\lg n$ ), even only supporting rank queries.

The locally decodable source coding [MHMP15] studies (almost) the same problem as Theorem 3, in a slightly different setting. They consider  $X$  that consists of i.i.d samples from a source of entropy  $H$ . However, the main focus is non-adaptive bit-probe query algorithms. That is, the query algorithm has to decide which  $t$  bits of the encoding to access based only on the queried index  $i$ . They studied the lossy case, where the encoding is equipped with the error correcting ability.

### 1.3 Technical contributions

We make two technical contributions to succinct data structures: We summarize the “spillover representation”, introduced by Pătraşcu [Păt08], to define binary strings with *fractional lengths* and build a “toolkit” of black-box operations; we study the “opposite” of data structures, called the *data interpretation*. We believe they will have more applications to other problems in succinct data structures.

#### 1.3.1 Strings with fractional lengths

A data structure is simply a bit string, and its length (or size) is the number of bits. Under standard notions, an  $s$ -bit string is only well-defined for integer  $s$ . Here, we show how to define such strings when  $s$  is fractional. We will see why this notion is useful later (or see [Păt08]).

Let  $(M, K)$  be a pair such that  $M \in \{0, 1\}^m$  is a bit string, and  $K \in [R]$  is an integer. Such a pair is viewed as a “binary string” of length  $m + \lg_2 R$ . When  $R$  is a power of two, this matches the standard notion of length, as we could simply write  $K$  in its binary representation using  $\lg R$  bits and append it to  $M$ . As we increase  $R$ , such a pair could potentially represent more information. Only when  $R$  is increased by a factor of two, does the pair correspond to a string with one more bit. That is, by restricting  $R \in [2^\kappa, 2^{\kappa+1})$  for some fixed parameter  $\kappa$ , we essentially “insert”  $2^\kappa - 1$  valid lengths between adjacent integers. It makes the measure of space more fine-grained. Also note that a uniformly random pair of this size has binary entropy exactly equal to  $m + \lg_2 R$ . In this paper,  $R$  is always set to  $2^{O(w)}$  (i.e.  $\kappa = O(w)$ ), where  $w$  is the word-size. Thus,  $K$  is an  $O(w)$ -bit integer, and the algorithms are able to do arithmetic operations on  $K$  in constant time.

We summarize a few black-box operations on fractional-length strings. The two major ones are *concatenation* and *fusion*.

**Concatenation.** Given  $B$  (fractional-length) strings  $\mathcal{S}_1, \dots, \mathcal{S}_B$  of lengths  $s_1, \dots, s_B$ , we show that they can be “concatenated” into one string of length  $s \approx s_1 + \dots + s_B$  (note that we do not get exactly  $s_1 + \dots + s_B$ , because the set of valid lengths is not closed under addition). This is trivial for integral-length strings, as we could simply connect all strings. Moreover for integral-lengths, suppose for a given  $i$ ,  $s_1 + \dots + s_{i-1}$  can be computed efficiently, then we will be able to find where  $\mathcal{S}_i$  starts, and access it within the long string. Likewise, we prove the same is true for fractional-length strings. That is, we show that if  $s_1 + \dots + s_{i-1}$  can be *well-approximated*, then  $\mathcal{S}_i$  can be *decoded*, i.e., it may be accessed within the long string, as if it was stored independently. We emphasize that decoding an input string  $\mathcal{S}_i = (M_i, K_i)$  *does not* mean reconstructing the entire string. Instead, the decoding algorithm only recovers  $K_i$ , and *finds where*  $M_i$  is located within the long string (where  $M_i$  is guaranteed to be a consecutive substring). Thus, the decoding algorithm can be *very efficient*. Nevertheless, after decoding,  $\mathcal{S}_i$  can still be accessed as if it was stored independently. In particular, by storing the prefix sums  $s_1 + \dots + s_i$  in a lookup table, we will be able to decode any  $\mathcal{S}_i$  in *constant time*.

Concatenation is useful when the data structure needs multiple subordinates. We could simply construct each subordinate separately and then concatenate them. It also demonstrates, to some extent, why fractional-lengths are useful and necessary. If we only use integral-length strings, then each  $\mathcal{S}_i$  will have length (at least)  $\lceil s_i \rceil$ . The length of the concatenated string becomes  $\lceil s_1 \rceil + \dots + \lceil s_B \rceil$ , which could be  $B - 1$  bits longer than  $\lceil s_1 + \dots + s_B \rceil$ .

**Fusion.** The other major operation is to *fuse* an integer into a string. Roughly speaking, it is to jointly store a pair  $(i, \mathcal{S}_i)$ , where  $i \in [C]$  is an integer, and  $\mathcal{S}_i$  is a string of length  $s_i$ . More specifically, let us first fix lengths  $s_1, \dots, s_C$ . We are then given a pair  $(i, \mathcal{S}_i)$  such that  $\mathcal{S}_i$  is guaranteed to have length  $s_i$ . We show that such a pair can be encoded by a string  $\mathcal{S}$  of (fixed) length  $s \approx \lg(2^{s_1} + \dots + 2^{s_C})$ . This length is the best possible, because there are  $2^{s_i}$  different possible strings of length  $\mathcal{S}_i$ . Therefore, there are  $2^{s_1} + \dots + 2^{s_C}$  different pairs  $(i, \mathcal{S}_i)$  in total. To encode such a pair,  $\lg(2^{s_1} + \dots + 2^{s_C})$  bits are necessary. Furthermore, suppose for *every*  $i$ ,  $\lg(2^{s_1} + \dots + 2^{s_i})$  can be *well-approximated*, then we will be able to recover the value of  $i$  and decode  $\mathcal{S}_i$ .

The fusion operation is useful when we study different cases of the input, and construct a data structure for each case separately. For example, suppose we wish to construct a data structure on a subset  $S \subseteq [n]$  (of arbitrary size), using close to  $n$  bits (and supporting some queries). Now suppose when  $|S| = i$ , we already know how to construct a data structure using  $\approx \lg \binom{n}{i}$  bits, such that the queries can be answered efficiently when the query algorithm is *given the value of*  $i$  for free. If this is the case for every  $i$ , then by

applying the fusion operation, we automatically obtain a data structure for subsets of *arbitrary* size using  $\approx n$  bits, *without* giving  $|S|$  to the query algorithm for free. To see this, let case  $i$  be all  $S$  with exactly  $i$  elements. Then we are able to construct a data structure  $\mathcal{S}_i$  just for all inputs in case  $i$ . The final data structure is the pair  $(i, \mathcal{S}_i)$  fused into one single data structure. The space bound guarantees that if each  $\mathcal{S}_i$  has nearly optimal size  $\approx \lg \binom{n}{i}$ , then the final data structure also has nearly optimal size of  $\approx n$  bits, since  $\sum_i \binom{n}{i} = 2^n$ . Given a query, we first retrieve the value of  $i$  and decode  $\mathcal{S}_i$ , then run the query algorithm for inputs in case  $i$ , given the value of  $i$ . We may also view  $\mathcal{S}_i$  as the data structure “conditioned on”  $i$ . Suppose all “conditional” data structures almost match their “conditional” optimal sizes, then they can be combined into one single data structure matching the overall optimal space. See Section 1.4 for a more concrete example.

By including a few other operations, we build a “toolkit” for operating on fractional-length strings. The view of fractional-length strings makes the “spillover representation” of Pătraşcu [Păt08] more explicit. The original paper needs huge lookup tables to store truth tables for  $O(w)$ -bit bizarre word operations. The new view assigns semantic meanings to those operations, so that a major part can be efficiently computed *without* lookup tables. This is the main reason why we can reduce the lookup table size.

### 1.3.2 Data interpretation

For a data structure problem, we preprocess a combinatorial object into a binary string. Then this string is stored in memory, which is divided into  $w$ -bit words. In each time step, a query algorithm may access a memory word (i.e. a  $w$ -bit substring), or do local computation. Finally, it computes some function of the input object. The concept of *data interpretation* is to perform the above procedure in the opposite direction. Given a binary string, we preprocess it into a combinatorial object. In each time step, a query algorithm may query an oracle for some function of the object, or do local computation. Finally, it reconstructs a  $w$ -bit substring of the input string.

Since this paper concerns data structures with space almost matching the information theoretical lower bound, we will also make data interpretations space-efficient. We design a data interpretation algorithm which preprocesses an input string of (fractional-)length  $\approx \lg \binom{V}{m}$  into a set  $S \subseteq [V]$  of size  $m$ , such that assuming there is a rank oracle for  $S$  (recall that  $\text{rank}_S(x)$  returns the number of elements in  $S$  that are at most  $x$ ), any designated  $w$  consecutive bits of the input string can be reconstructed in  $\text{poly } \lg V$  time (see Section 7.1). It might not be obvious at this moment why it is beneficial to convert a string (data structure) back to a set, but it turns out to be a key subroutine in our data structure. See Section 2 for more details.

## 1.4 Overview of Pătraşcu’s data structure

In this subsection, we summarize how Pătraşcu’s rank data structure [Păt08] works, which has important ideas to be used in our data structure. We will “rephrase” this data structure using fractional-length strings, which is a non-trivial simplification.

Given a set  $S \subseteq [U]$  of size  $n$ , the rank data structure preprocesses it into  $\approx \lg \binom{U}{n}$  bits, such that for any given query  $x$ , the number of elements in  $S$  that are at most  $x$  can be computed in  $O(\lg U)$  time. The idea is to recursively construct data structures for smaller universes, and then merge the subordinate using concatenation and fusion. Suppose  $S$  has  $i$  elements in  $\{0, \dots, U/2 - 1\}$ , the first half of the universe, and it has  $n - i$  elements in the second half. We first recursively construct (fractional-length) data structures for both halves, using space  $\approx \lg \binom{U/2}{i}$  and  $\approx \lg \binom{U/2}{n-i}$  respectively. Next, we concatenate two data structures, and obtain one single data structure  $\mathcal{S}_i$ , which has length  $\approx \lg \left( \binom{U/2}{i} \binom{U/2}{n-i} \right)$ . Note that the data structure  $\mathcal{S}_i$

encodes an input set  $S$  with exactly  $i$  elements in the first half (and  $n - i$  in the second half), and it *does not* encode the value of  $i$  (likewise in the desired final data structure, the value of  $n$  is assumed to be known, and is not encoded). Finally, we encode the value of  $i$  by fusing it into  $\mathcal{S}_i$ , i.e., we jointly store the pair  $(i, \mathcal{S}_i)$ . The fusion operation guarantees that the pair can be stored using approximately

$$\lg \left( \sum_{i=0}^n 2^{|\mathcal{S}_i|} \right) \approx \lg \left( \sum_{i=0}^n \binom{U/2}{i} \binom{U/2}{n-i} \right) = \lg \binom{U}{n}$$

bits.

This recursion terminates at sets of size  $n = 0$  or  $n = U$ , in which case there is nothing to store (again we assume  $n$  does not need encoding, so it is clear which case we are in). We guarantee that both concatenation and fusion are implemented such that each operation only causes an overhead of no more than  $O(1/U^2)$  bits. Therefore the overall space is no more than  $\lg \binom{U}{n} + O(1/U)$ . For the final (fractional-length) data structure  $(M, K)$ , we simply write  $K$  in its binary representation and append it to  $M$ . This gives us an integral-length data structure using at most  $\lceil \lg \binom{U}{n} \rceil + 1$  bits.

It is then straightforward to answer a rank query on this data structure. Given a query  $x$ , we first recover the value of  $i$ , and decode  $\mathcal{S}_i$  (again, decoding  $\mathcal{S}_i$  does not mean reconstructing it). Then we further decode  $\mathcal{S}_i$  into the two data structures for the two halves. It can be done in constant time using a lookup table. Next, if  $x < U/2$ , we recurse into the first half. If  $x \geq U/2$ , we recurse into the second half (and add  $i$  to the final answer). Since each time  $U$  decreases by a factor of two, the query time is  $O(\lg U)$ .

In [Pät08], it is also shown that when  $U$  is small, we can do a  $B$ -way divide-and-conquer, as long as  $B \lg U \leq O(w)$  (recall that  $w$  is the word-size). Therefore when  $U \leq w^{O(1)}$ , we can afford to set  $B = w^{1/2}$  and have only constant depth of recursion (rather than  $O(\lg U)$ ). This gives us a rank data structure with constant query time for small  $U$ . In this paper, we show that it is possible to further improve it, and we design a constant-query-time data structure when only  $n$  is bounded by  $w^{O(1)}$  (and  $U$  could be still as large as  $2^{\Theta(w)}$ ). This will be the starting point of our new data structure.

## 2 Overview

In this section, we overview our new static dictionary. For simplicity, we will first focus on the membership queries (i.e.,  $\sigma = 1$ ), and assume  $U = \text{poly } n$ . In this case, all previous solutions use hash functions in their main construction, to map the keys into buckets. Our data structure is conceptually different: Instead of random hash functions, we consider *random inputs*. While our data structure works for worst-case inputs, let us first think of the input set being  $n$  uniformly random (distinct) keys. Then with high probability, the input already has the properties we wanted from a random hash function, e.g., by dividing the key space into buckets in some fixed way, we have the sizes of buckets roughly balanced, etc. We first construct a data structure just for those “random-looking” inputs. On the other hand, with low probability, the input may look “non-typical,” e.g., some bucket may have size much larger than average. However, “with low probability” means that only a small fraction of all possible inputs have these non-typical features. Suppose the total number of such inputs is, say  $\frac{1}{n^2} \cdot 2^{\text{OPT}}$ , then only  $\text{OPT} - 2 \lg n$  bits are needed for the optimal encoding. This suggests that we can afford to spend more extra bits on them. Suppose we use  $\text{OPT} - \lg n$  bits ( $\lg n$  extra bits) to encoding these non-typical inputs, it is still negligible overall — among all  $O(2^{\text{OPT}})$  possible data structures (memory states), such an encoding only wastes  $2^{\text{OPT}} \cdot (\frac{1}{n} - \frac{1}{n^2})$ . Another useful way to view it is that if we use  $x$  extra bits for such rare cases, then those  $x$  bits “start” at the  $(\text{OPT} - 2 \lg n)$ -th bit, rather than the  $\text{OPT}$ -th bit. The more non-typical the input is, the more extra bits we can afford to spend. Finally, we will

use the fusion operation to fuse all cases together. Similar strategies for constructing succinct data structures, where we consider random inputs and/or non-typical inputs, have been used in [BL13, Yu19, VWY19].

In the following, we show how to handle the “random-looking” case and the “non-typical” cases for membership.

## 2.1 Random inputs

We partition the universe into  $n/\lg^4 n$  buckets of size  $V$ . Then for a “random-looking” input set  $S$ , there are  $\lg^4 n \pm \lg^3 n$  keys in *every* block. As we mentioned in Section 1.4, for  $\text{poly } \lg n = \text{poly } w$  keys, we can construct a rank data structure with only  $O(1/U)$  extra bits, such that given the number of keys, a query algorithm answers rank queries in constant time. In particular, it supports membership queries (e.g., by asking  $\text{rank}_S(x)$  and  $\text{rank}_S(x - 1)$ ). The high-level idea is to construct a rank data structure for each block, then *concatenate* them. In order to answer a query in block  $i$ , we need to

- recover the number of keys in block  $i$  (as the rank data structure assumes this number is known), and
- approximate the total length of data structures for first  $i - 1$  blocks (to decode the  $i$ -th data structure).

That is, besides the  $n/\lg^4 n$  rank data structures, we need to store their lengths such that any prefix sum can be approximated. Unfortunately, any data structure supporting prefix sum queries cannot simultaneously have “low” query time and “small” space, due to a lower bound of Pătraşcu and Viola [PV10]. The underlying issue in this approach is that the data structure for each block has a variable length (the length depends on the number of keys in the block, which varies based on the input). In order to locate the  $i$ -th data structure from the concatenated string, computing a prefix sum on a sequence of variables seems inevitable. The Pătraşcu-Viola lower bound even prevents us from supporting prefix sums *implicitly*. That is, not only separately storing a prefix sum data structure for the lengths requires “high” query time or “large” space, there is also no “clever” way to jointly store the lengths together with the data structures for the blocks. Hence, this “variable-length encoding” issue is the primary problem we need to tackle for “random-looking” inputs.

To this end, observe that although the number of keys in each block is not fixed, its deviation is actually small compared to the number, i.e., the number of keys cannot be too different for different inputs. Then the main idea is to construct *two* data structures for each block, consisting of

- *a main data structure*, which stores “most of the information” about the block, and importantly, has a fixed length (independent of the number of keys), and
- *an auxiliary data structure*, which stores all “remaining information” about the block (and unavoidably has variable length).

Furthermore, we wish that with high probability, a given query can be answered by only accessing the main data structure (in constant time) *without knowing the number of keys*. If it is possible, then to construct the final data structure, we

- concatenate all main data structures,
- concatenate the auxiliary data structures, and store them together with a prefix sum structure,
- finally concatenate the two.

Now, since all main data structures have fixed lengths, each one can be decoded in constant time without a prefix sum structure (the total length of the first  $i - 1$  data structures is simply  $i - 1$  times the length of a single one). Then to answer a query in block  $i$ , we first decode the  $i$ -th main data structure, and query it in constant time. With high probability, the answer to the query is already found, and we are done. Otherwise, we decode the  $i$ -th auxiliary data structure by querying the prefix sum structure, and query the data structures to find the answer. It may take a longer time, but if the probability that we have to decode the auxiliary data structure is sufficiently low, then the expected query time is still constant.

Next, we describe an approach to construct such two data structures for a block, which uses more space than what we aim for, but exhibits the main idea. For each block of size  $V$ , we pick  $\lg^4 n - \lg^3 n$  *random* keys in the block to store in the main data structure. We also pick  $V - (\lg^4 n + \lg^3 n)$  *random non-keys* (i.e. the elements in the key space but not in the input set), and store them in the main data structure. This is always possible because there are at least  $\lg^4 n - \lg^3 n$  and at most  $\lg^4 n + \lg^3 n$  keys in each block for “random-looking” inputs. Hence, only  $2\lg^3 n$  elements are “unknown” from the main data structure. Then we show that such a separation of the block into  $\lg^4 n - \lg^3 n$  keys,  $V - (\lg^4 n + \lg^3 n)$  non-keys and  $2\lg^3 n$  unknowns can be jointly stored using the near-optimal  $\approx \lg \binom{V}{\lg^4 n - \lg^3 n, 2\lg^3 n}$  bits (this is an easy application of the rank data structures).<sup>3</sup> Its size is independent of the actual input. Then in the auxiliary data structure, we store the remaining information about the block, i.e., among the unknowns, which ones are the keys. For a block with  $m$  keys, it takes  $\approx \lg \binom{2\lg^3 n}{m - (\lg^4 n - \lg^3 n)}$  bits. Then for each query, the answer can be found in the main data structure with probability at least  $1 - O(1/\lg n)$ . Only when the main data structure returns “unknown”, does the query algorithm need to access the auxiliary data structure.

The above construction has all the desired properties, except that it uses too much space. The inherent reason is that it implicitly stores the *randomness* used in deciding which keys and non-keys to store in the main data structure. If we sum up the sizes of the main and auxiliary data structures,

$$\begin{aligned} & \lg \binom{V}{\lg^4 n - \lg^3 n, 2\lg^3 n} + \lg \binom{2\lg^3 n}{m - (\lg^4 n - \lg^3 n)} \\ &= \lg \binom{V}{m} + \lg \binom{m}{\lg^4 n - \lg^3 n} + \lg \binom{V - m}{V - (\lg^4 n + \lg^3 n)}. \end{aligned}$$

Unsurprisingly, the number of extra bits  $\lg \binom{m}{\lg^4 n - \lg^3 n} + \lg \binom{V - m}{V - (\lg^4 n + \lg^3 n)}$  is exactly how much is needed to decide which keys and non-keys to store in the main data structure. These “random bits” are not part of the input, and implicitly storing them causes a large amount of redundancy.

However, when the inputs are uniform, we do not really need any external randomness to decide the two subsets, since the entire data structure is close to a random string. This suggests that for each block, we should treat the data structure from other part of the inputs as the “randomness”. That is, we use the opportunity of implicitly storing the random bits, to store other information that needs to be stored. This is where we use *data interpretation*. We convert existing data structures back to subsets of certain sizes, which correspond to the keys and non-keys in the main data structure. The details are presented in the next subsection.

## 2.2 Using data interpretation

To implement this idea, we will have to slightly modify the construction. Now, the universe is partitioned into *pairs of blocks*. Each pair consists of a *primary* block and a *secondary* block, such that for a “random-looking” input, the primary block contains  $\lg^{2c} n \pm \lg^{c+1} n$  keys, and the secondary block contains  $\Theta(\lg^{c+1} n)$  keys (which plays the role of the “randomness”), for some constant  $c$ . Fix a block pair, let  $V$  be the size of the primary block,  $m$  be the number of keys in the primary block,  $V_{sc}$  be the size of the secondary block, and  $m_{sc}$  be the number of keys in the secondary block. The goal is to construct two data structures using  $\approx \lg \binom{V}{m} + \lg \binom{V_{sc}}{m_{sc}}$  bits in total.

We first construct a rank data structure for the secondary block using  $\approx \lg \binom{V_{sc}}{m_{sc}}$  bits. We then *divide* this data structure into three substrings of lengths approximately  $\lg \binom{m}{\lg^{2c} n - \lg^{c+1} n}$ ,  $\lg \binom{V - m}{V - (\lg^{2c} n + \lg^{c+1} n)}$

---

<sup>3</sup>  $\binom{n}{k_1, k_2} = n! / (k_1! k_2! (n - k_1 - k_2)!)$ .

and  $\lg \binom{V_{\text{sc}}}{m_{\text{sc}}} - \lg \binom{m}{\lg^{2c} n - \lg^{c+1} n} - \lg \binom{V-m}{V - (\lg^{2c} n + \lg^{c+1} n)}$  (we show divisions can also be done for fractional-length strings). Then  $m_{\text{sc}} = \Theta(\lg^{c+1} n)$  guarantees that there are enough bits and such division is possible. Next, we apply a data interpretation algorithm to interpret the first string of length  $\lg \binom{m}{\lg^{2c} n - \lg^{c+1} n}$  as a set of size  $\lg^{2c} n - \lg^{c+1} n$  over a universe of size  $m$ , indicating which of the  $m$  keys in the primary block should be stored in the main data structure. We also interpret the second string as a subset indicating which of the  $V - (\lg^{2c} n + \lg^{c+1} n)$  non-keys should be stored in the main data structure. Moreover, we show that the data interpretation algorithm guarantees that any consecutive  $w$  bits of the original string can be recovered in  $\lg^{O(1)} n$  time, assuming there is a rank oracle of the set generated from the interpretation. Therefore, there is no need to store the first two strings, as they can be implicitly accessed efficiently.

The main data structure is the same as what we stated in the previous subsection: storing  $\lg^{2c} n - \lg^{c+1} n$  keys,  $V - (\lg^{2c} n + \lg^{c+1} n)$  non-keys and  $2 \lg^{c+1} n$  “unknowns”, supporting rank queries in constant time. The auxiliary data structure now consists of two parts:

- among the  $2 \lg^{c+1} n$  unknowns, which  $m - (\lg^{2c} n - \lg^{c+1} n)$  are keys, and
- the third substring from above.

One may verify that the sizes of the two data structures is what we claimed. This leads to our main technical lemma.

**Lemma 4** (main technical lemma, informal). *For  $V \leq \text{poly } n$ , given  $S \subseteq [V]$  of size  $m$  and  $S_{\text{sc}} \subseteq [V_{\text{sc}}]$  of size  $m_{\text{sc}}$ , we can construct a main data structure  $\mathcal{D}_{\text{main}}$  of size*

$$\approx \lg \binom{V}{\lg^{2c} n - \lg^{c+1} n, 2 \lg^{c+1} n}$$

*and an auxiliary data structure  $\mathcal{D}_{\text{aux}}$  of size*

$$\approx \lg \binom{V}{m} + \lg \binom{V_{\text{sc}}}{m_{\text{sc}}} - \lg \binom{V}{\lg^{2c} n - \lg^{c+1} n, 2 \lg^{c+1} n},$$

*such that*

- any given query “ $x \stackrel{?}{\in} S$ ” can be answered in constant time by accessing only  $\mathcal{D}_{\text{main}}$  with probability  $1 - O(\lg^{-c+1} n)$ ;
- any given query “ $x \stackrel{?}{\in} S$ ” or “ $x \stackrel{?}{\in} S_{\text{sc}}$ ” can be answered in  $\text{poly } \lg n$  time by accessing both  $\mathcal{D}_{\text{main}}$  and  $\mathcal{D}_{\text{aux}}$  in worst-case.

See Section 7 for the formal statement. Then, the final data structure will be the concatenation of all main and auxiliary data structures, in a similar way to what we stated in the previous subsection. The auxiliary data structure needs to be decoded only when the main data structure returns “unknown” or the query lands in a secondary block. By randomly shifting the universe, we bound the probability of needing the auxiliary data structure by  $O(\lg^{-c+1} n)$ . By setting  $c$  to be a sufficiently large constant, the expected query time is constant.

### 2.3 Constructing data interpretation

Next, we briefly describe how to design such a data interpretation algorithm, i.e., to convert a (fractional-length) string to a set. The high-level idea is similar to the rank data structure described in Section 1.4, with all steps done in the opposite direction. The data structure uses concatenation and fusion of fractional-length strings. We first show how to do the opposite of the two operations. More concretely, we show that

- given a string  $\mathcal{D}$ , it can be *divided* into two substrings  $\mathcal{D}_1, \mathcal{D}_2$  of given lengths;
- given a string  $\mathcal{D}$ , it can be viewed as a pair  $(i, \mathcal{D}_i)$ , where  $\mathcal{D}_i$  has a given length  $s_i$ , i.e., we *extract* an integer  $i$  from  $\mathcal{D}$  and let  $\mathcal{D}_i$  be the rest.

Both division and extraction are done with negligible space overhead.

Then, given a string  $\mathcal{D}$  of length  $\approx \lg \binom{V}{m}$ , to interpret it as a set of size  $m$ , we first *extract* an integer  $i$  from  $\mathcal{D}$  such that  $i \in \{0, \dots, m\}$  and  $\mathcal{D}_i$  has length  $s_i \approx \lg \binom{V/2}{i} \binom{V/2}{m-i}$ . Then we divide  $\mathcal{D}_i$  of length  $s_i$  into two substrings  $\mathcal{D}_a$  and  $\mathcal{D}_b$  of lengths  $\approx \lg \binom{V/2}{i}$  and  $\approx \lg \binom{V/2}{m-i}$  respectively. The integer  $i$  will represent the number of keys in the first half of the universe, and  $m - i$  is the number of keys in the second half. We recursively construct sets  $S_a, S_b \subseteq [V/2]$  from  $\mathcal{D}_a$  and  $\mathcal{D}_b$  of sizes  $i$  and  $m - i$  respectively. Then the final set is  $S_a \cup (S_b + V/2)$ .

To access  $w$  consecutive bits of  $\mathcal{D}$  given a  $\text{rank}_S(\cdot)$  oracle, we first ask the oracle  $\text{rank}_S(V/2)$ , i.e., the number of keys in the first half. This determines the value of  $i$ , and hence the lengths of  $\mathcal{D}_a$  and  $\mathcal{D}_b$ , which in turn determines whether the  $w$  consecutive bits are entirely in  $\mathcal{D}_a$ , or entirely in  $\mathcal{D}_b$ , or split across the two substrings. If it is entirely contained in one substring, we simply recurse into the corresponding half of the universe. On the other hand, it is possible to show that splitting across the two substrings does not happen too many times, and when it happens, we recurse into both halves. The recursion has depth  $O(\lg V)$ . More details can be found in Section 7.1.

## 2.4 Worst-case input

Applying the above data structure to worst-case input has the following two issues:

1. for each primary block, the sets stored in the main data structure are no longer random, hence, the expected query time may not be constant;
2. some primary block may not have  $\lg^{2c} n \pm \lg^{c+1} n$  keys, and some secondary block may not have  $\Theta(\lg^{c+1} n)$  keys.

The first issue is easy to resolve. We simply sample a uniformly random string  $\mathcal{R}$ , and “XOR” it to the substrings before we do data interpretation. Thus, every string to be interpreted as a set will be a uniformly random string. In particular, for each primary block, the subsets being stored in the main data structure are uniformly random subsets marginally. This is sufficient to guarantee the constant query time. We use the same  $\mathcal{R}$  for all blocks. Storing  $\mathcal{R}$  in the data structure introduces  $\text{poly } \lg n$  extra bits of space.

For the second issue, we use the earlier argument. If the expected number of keys in a primary block is  $\lg^{2c} n$ , then the probability that a random set has some primary block with more than  $\lg^{2c} n + \lg^{c+1} n$  is at most  $\exp(-\Theta(\lg^2 n))$ . Then we can afford to use more extra bits to encode such inputs. Suppose  $S$  has a block with, say  $\approx 2\lg^{2c} n$  keys. Then we could simply spend  $O(\lg n)$  extra bits to encode which block it is, the number of keys in it, as well as a pointer to a separate rank data structure for this block. The space usage for such inputs is still  $\text{OPT} - \Omega(\lg^2 n)$ .

Similarly, we can show that the probability that  $S$  has  $N$  blocks with too many or too few keys is at most  $\exp(-\Theta(N \lg^2 n))$ , suggesting that we can afford to use  $O(N \lg^2 n)$  extra bits (which we will use to store a perfect hash table for these “bad” blocks); the probability that  $S$  has a block with  $m$  keys for  $m > \lg^{3c} n$  is  $\exp(-\Theta(m \lg m))$ , suggesting that we may at least use  $O(m)$  extra bits for such a block (which is sufficient to store the previously known membership data structure). By computing the probability that every particular case happens, we estimate how many extra bits we can afford. The more “non-typical” the input is, the more extra bits we may use. We then construct a data structure within the allowed extra bits. It turns out that overall, the total space usage for inputs with *at least* one “bad” block is at most  $\text{OPT} - \Theta(\lg^2 n)$ . Finally, we apply the fusion operation to combine the “random-looking” inputs and these

“non-typical” inputs: if every block pair has number of keys close to the expectation, we set  $b := 1$  and construct the data structure as in the previous subsection using  $\approx \text{OPT}$  bits; if at least one block pair has too many or too few keys, we set  $b := 2$  and construct a data structure using  $\text{OPT} - \Theta(\lg^2 n)$  bits; then we fuse  $b$  into the data structure. By the guarantee of the fusion operation, the final space is bounded by

$$\approx \lg(2^{\text{OPT}} + 2^{\text{OPT} - \Theta(\lg^2 n)}) = \text{OPT} + \lg(1 + 2^{-\Theta(\lg^2 n)}) = \text{OPT} + o(1).$$

Generalizing the data structure to  $\sigma > 1$  (associating each key with a value) can be done by generalizing Lemma 4. Since the underlying data structure supports rank, it naturally maps  $m$  keys to  $\{1, \dots, m\}$ . Then we use [DPT10], or simply concatenation, to store the list of  $m$  values. To retrieve the value of a key  $x$ , we first find its rank  $x_r$ . Then retrieve the  $x_r$ -th value in the list. This generalizes the data structure to the dictionary problem at essentially no extra cost.

## 2.5 Organization

In Section 3, we define notations and the model of computation. In Section 4, we formally define fractional-length strings, and state the black-box operations (the proofs are deferred to the appendix). In Section 5, we show how to construct the succinct dictionary and locally decodable arithmetic codes using perfect hashing. In Section 6, we design the data structure for the case where  $U = \text{poly } n$  using the main technical lemma. Then we prove the main technical lemma in Section 7, and generalize to all  $n$  and  $U$  in Section 8.

# 3 Preliminaries and Notations

## 3.1 Random access machine

A random access machine (RAM) has a memory divided into  $w$ -bit *words*, where  $w$  is called the *word-size*. Typically, we assume the number of words in the memory is at most  $2^w$ , and they are indexed by  $\{0, \dots, 2^w - 1\}$ . In each time step, an algorithm may load one memory word to one of its  $O(1)$  CPU registers, write the content of a CPU register to one memory word, or compute (limited) word operations on the CPU registers.

The standard word operations are the four basic arithmetic operations (addition, subtraction, multiplication and division) on  $w$ -bit integers, bit-wise operators (AND, OR, XOR), and comparison. In this paper, we also assume that the machine supports *floating-point* numbers. A floating-point number is represented by two registers in the form of  $a \cdot 2^b$ , and the arithmetic operations extend to these numbers as well (possibly with rounding errors). This is without loss of generality, as they can be simulated using the standard operations. Finally, we assume it is possible to compute  $2^x$  up to  $1 \pm 2^{-w}$  multiplicative error, and  $\lg_2 x$  up to additive  $\pm 2^{-w}$  error. We further assume that the error can be arbitrary but has to be deterministic, i.e., for any given  $x$ ,  $2^x$  and  $\lg_2 x$  always compute to the same result within the desired range. By expanding into the Taylor series, the two can be computed in  $O(w)$  time using only arithmetic operations, which is already sufficient for our application. On the other hand, using a lookup table of size  $2^{\epsilon w}$  (we already have lookup tables of this size), the computational time can be reduced to constant.

## 3.2 Notations

In this paper, let  $X \text{ div } Y$  denote  $\lfloor X/Y \rfloor$ ,  $X \bmod Y$  denote  $X - Y \cdot (X \text{ div } Y)$ . Let  $[R]$  denote the set  $\{0, 1, \dots, R - 1\}$ . Let  $\text{frac}(x)$  denote  $x - \lfloor x \rfloor$ . Throughout the paper,  $\lg x$  is the binary logarithm  $\lg_2 x$ ,  $\tilde{O}(f) = f \cdot \text{poly } \lg f$ .

## 4 Fractional-length Strings

In this subsection, we formally define binary strings with fractional lengths using the spillover representation of [Pät08], and state block-box operations. Throughout the paper, let  $\kappa$  be the *fineness parameter*, which characterizes the gaps between adjacent valid lengths, and determines the space loss when doing the operations. It is an integer parameter that is specified by the algorithm designer and will be hardwired to the preprocessing and query algorithms. In the following, we will see that each operation loses  $O(2^{-\kappa})$  bits; and on the other hand, the algorithms will have to perform arithmetic operations on  $\kappa$ -bit integers. In our data structure construction,  $\kappa$  will be set to  $\Theta(\lg U) = O(w)$ , so that each operation loses negligible space, and  $\kappa$ -bit arithmetic operations can still be performed in constant time.

**Definition 5** (fractional-length strings). Let  $\mathcal{S} = (M, K)$  be a pair such that  $M \in \{0, 1\}^m$  and  $K \in [R]$ , where  $m$  is a nonnegative integer and  $R \in [2^\kappa, 2^{\kappa+1})$  is an integer, or  $m = 0$  and integer  $R \in [1, 2^\kappa)$ . Then  $\mathcal{S}$  is a *binary string of length  $m + \lg R$* , and

$$\mathcal{S}[i] := \begin{cases} M[i] & i \in [m], \\ K & i = m. \end{cases}$$

Let  $|\mathcal{S}|$  denote the length of  $\mathcal{S}$ . Let  $\mathcal{S}[i, j]$  denote the sequence (substring)  $(\mathcal{S}[i], \dots, \mathcal{S}[j])$ . Let  $\text{range}(K) := R$  be the size of range of  $K$ .

*Remark.* Note the following facts about fractional-length strings:

- When  $s$  is an integer, by writing  $K$  in its binary representation, a binary string of length  $s$  from Definition 5 is a standard binary string of  $s$  bits;
- A uniformly random string of length  $s$  has entropy exactly  $s$ ;
- $|\mathcal{S}|$  uniquely determines  $|M|$  and  $\text{range}(K)$ , i.e., if  $|\mathcal{S}| < \kappa$ ,  $|M| = 0$  and  $\text{range}(K) = 2^s$ , otherwise,  $|M| = \lfloor |\mathcal{S}| \rfloor - \kappa$  and  $\text{range}(K) = 2^{\kappa + \text{frac}(|\mathcal{S}|)}$ ;
- The length of a string may be an irrational number, but it can always be succinctly encoded, e.g., by encoding  $|M|$  and  $\text{range}(K)$ ;
- One should *not* view  $\text{range}(K)$  as a function of  $K$ , it is indeed a parameter of the *variable*  $K$ .

Since the word-size is  $\Omega(\kappa)$ , any  $O(\kappa)$  consecutive bits of a string can be retrieved using  $O(1)$  memory accesses, which suggests how a fractional-length string is accessed. Formally, we define an access as follows.

**Definition 6** (access). Let  $\mathcal{S}$  be a (fractional-length) string, an *access* to  $\mathcal{S}$  is to retrieve  $\mathcal{S}[i, j]$  for  $j - i \leq O(\kappa)$ .

When  $j < |M|$ , an access is to retrieve  $j - i + 1$  bits of  $M$ . When  $j = |M|$ , it is to retrieve  $j - i$  bits and the integer  $K$ .

In the following, we show how to operate on fractional-length strings. All the proofs are deferred to Appendix A. Firstly, the strings can be *concatenated*.

**Proposition 7** (concatenation). Let  $s_1, \dots, s_B \geq \kappa$ . Suppose for any given  $i$ ,  $s_1 + \dots + s_i$  can be approximated (deterministically) in  $O(t)$  time with an additive error of at most  $2^{-\kappa}$ . Then given  $B$  strings  $\mathcal{S}_1, \dots, \mathcal{S}_B$ , where  $\mathcal{S}_i = (M_i, K_i)$  has length  $s_i$ , they can be concatenated into one string  $\mathcal{S} = (M, K)$  of length at most

$$s_1 + \dots + s_B + (B - 1) \cdot 2^{-\kappa+4},$$

so that each  $M_i$  is a (consecutive) substring of  $M$ . Moreover, for any given  $i$ ,  $\mathcal{S}_i$  can be decoded using  $O(t)$  time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $K_i$ , and finds the starting location of  $M_i$  using  $O(t)$  time and two accesses to  $\mathcal{S}$ .

In particular, by storing approximations of all  $B$  prefix sums in a lookup table of size  $O(B)$ , each  $\mathcal{S}_i$  can be decoded in  $O(1)$  time. Note that this lookup table *does not* depend on the  $B$  strings.

**Proposition 8** (concatenation). *Let  $s_1, \dots, s_B \geq 0$ . There is a lookup table of size  $O(B)$ . Given  $B$  strings  $\mathcal{S}_1, \dots, \mathcal{S}_B$ , where  $\mathcal{S}_i = (M_i, K_i)$  has length  $s_i$ , they can be concatenated into one string  $\mathcal{S} = (M, K)$  of length at most*

$$s_1 + \dots + s_B + (B-1)2^{-\kappa+4},$$

so that each  $M_i$  is a (consecutive) substring of  $M$ . Moreover, assuming we can make random accesses to the lookup table,  $\mathcal{S}_i$  can be decoded using constant time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $K_i$ , and finds the starting location of  $M_i$  using constant time and two accesses to  $\mathcal{S}$ .

Next, an integer  $i \in [C]$  can be fused into a string.

**Proposition 9** (fusion). *Let  $s_1, \dots, s_C \geq 0$ . Suppose for any given  $i$ ,  $2^{s_1} + \dots + 2^{s_i}$  can be approximated (deterministically) in  $O(t)$  time with an additive error of at most  $(2^{s_1} + \dots + 2^{s_C}) \cdot 2^{-\kappa-3}$ . Then given  $i \in \{1, \dots, C\}$  and string  $\mathcal{S}_i = (M_i, K_i)$  of length  $s_i$ , the pair  $(i, \mathcal{S}_i)$  can be stored in  $\mathcal{S} = (M, K)$  of length at most*

$$\lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+4},$$

so that  $M_i$  is a (consecutive) substring of  $M$ . Moreover, we can recover the value of  $i$  and decode  $\mathcal{S}_i$  using  $O(t \lg C)$  time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $i$ ,  $K_i$ , and finds the starting location of  $M_i$  using  $O(t \lg C)$  time and two accesses to  $\mathcal{S}$ .

Note that the error term is always proportional to  $2^{s_1} + \dots + 2^{s_C}$ , for every  $i$ . In particular, it is possible that for some (small)  $i$ , the error term dominates the value, making the assumption easy to satisfy (for that  $i$ ). The decoding algorithm can take constant time if we use a lookup table of size  $O(C)$ . Again, the lookup table does not depend on the string.

**Proposition 10** (fusion). *Let  $s_1, \dots, s_C \geq 0$ . There is a lookup table of size  $O(C)$ . Given  $i \in \{1, \dots, C\}$  and string  $\mathcal{S}_i = (M_i, K_i)$  of length  $s_i$ , the pair  $(i, \mathcal{S}_i)$  can be stored in  $\mathcal{S} = (M, K)$  of length*

$$\lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2},$$

so that  $M_i$  is a (consecutive) substring of  $M$ . Moreover, assuming we can make random accesses to the lookup table, the value of  $i$  can be recovered and  $\mathcal{S}_i$  can be decoded using constant time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $i$ ,  $K_i$ , and finds the starting location of  $M_i$  using constant time and two accesses to  $\mathcal{S}$ .

Fractional-length strings have “one of its ends” encoded using an integer. For technical reasons, we also need the following notion of *double-ended* binary strings.

**Definition 11** (double-ended strings). Let  $\mathcal{S} = (K_h, M, K_t)$  be a triple such that  $M \in \{0, 1\}^m$ ,  $K_h \in [R_h]$  and  $K_t \in [R_t]$ , where  $m$  is a nonnegative integer and  $R_h, R_t \in [2^\kappa, 2^{\kappa+1})$  are integers. Then  $\mathcal{S}$  is a *double-ended binary string of length  $m + \lg R_h + \lg R_t$* , and

$$\mathcal{S}[i] := \begin{cases} K_h & i = -1, \\ M[i] & i \in [m], \\ K_t & i = m. \end{cases}$$

Let  $|\mathcal{S}|$  denote the length of  $\mathcal{S}$ . Let  $\mathcal{S}[i, j]$  denote the substring  $(\mathcal{S}[i], \dots, \mathcal{S}[j])$ . Let  $\text{range}(K_h) := R_h$ ,  $\text{range}(K_t) := R_t$  be the sizes of ranges of  $K_h$  and  $K_t$  respectively.

*Remark.* Note the following facts:

- Unlike the (single-ended) fraction-length strings, the length of a double-ended string does not necessarily determine  $\text{range}(K_h)$ ,  $\text{range}(K_t)$ , or even  $|M|$ ;
- For  $s \geq 2\kappa$ , any  $s$ -bit string  $(M, K)$  can be viewed as a double-ended string by taking the first  $\kappa$  bits of  $M$  as  $K_h$  and letting  $K_t$  be  $K$ ;
- For simplicity, in this paper, we do not define double-ended strings with length shorter than  $2\kappa$ ;

Double-ended strings are accessed in the same way.

**Definition 12** (access). Let  $\mathcal{S}$  be a double-ended string, an *access* to  $\mathcal{S}$  is to retrieve  $\mathcal{S}[i, j]$  for  $j - i \leq O(\kappa)$ .

Prefixes and suffixes of a double-ended string are defined in the natural way, as follows.

**Definition 13** (prefix/suffix). Let  $\mathcal{S} = (K_h, M, K_t)$  be a double-ended string. Then  $\mathcal{S}[-1, j]$  is a *prefix* of  $\mathcal{S}$  for any  $j \leq |M|$ ,  $\mathcal{S}[i, |M|]$  is a *suffix* of  $\mathcal{S}$  for any  $i \geq -1$ .

Using double-ended strings, it is possible to divide a binary string into two substrings.

**Proposition 14** (divide). Let  $s_1, s_2, s \geq 3\kappa$  and  $s \leq s_1 + s_2 - 2^{-\kappa+2}$ . Then given a double-ended string  $\mathcal{S} = (K_h, M, K_t)$  of length  $s$ , a division algorithm outputs two double-ended strings  $\mathcal{S}_1 = (K_{1,h}, M_1, K_{1,t})$  and  $\mathcal{S}_2 = (K_{2,h}, M_2, K_{2,t})$  of lengths at most  $s_1$  and  $s_2$  respectively. Moreover,  $(K_{1,h}, M_1)$  is a prefix of  $\mathcal{S}$ ,  $(M_2, K_{2,t})$  is a suffix of  $\mathcal{S}$ , and  $K_{1,t}$  and  $K_{2,h}$  together determine  $M[|M_1|, |M| - |M_2| - 1]$ , i.e., the remaining bits of  $M$ .  $\text{range}(K_{i,h})$ ,  $\text{range}(K_{i,t})$  and  $|M_i|$  can be computed in  $O(1)$  time given  $\text{range}(K_h)$ ,  $\text{range}(K_t)$ ,  $|M|$  and  $s_1, s_2$ , for  $i = 1, 2$ .

*Remark.* Proposition 14 guarantees that each access to  $\mathcal{S}$  can be implemented using at most two accesses to  $\mathcal{S}_1$  and  $\mathcal{S}_2$ . Moreover, accessing a (short) prefix of  $\mathcal{S}$  requires only accessing the prefix of  $\mathcal{S}_1$  of the same length. Likewise, accessing a suffix of  $\mathcal{S}$  requires only accessing the suffix of  $\mathcal{S}_2$  of the same length.

Finally, the “inverse” of fusion can also be done efficiently.

**Proposition 15** (extraction). Let  $s_1, \dots, s_C \geq 0$ ,  $R_h, R_t \in [2^\kappa, 2^{\kappa+1})$  and  $m \geq \kappa$ , let  $s = m + \lg R_h + \lg R_t$ , and  $s \leq \lg(2^{s_1} + \dots + 2^{s_C}) - C \cdot 2^{-\kappa+2}$ , there is a lookup table of size  $O(C)$ . Given a double-ended string  $\mathcal{S} = (K_h, M, K_t)$  such that  $\text{range}(K_h) = R_h$ ,  $\text{range}(K_t) = R_t$  and  $|M| = m$ , there is an extraction algorithm that generates a pair  $(i, \mathcal{S}_i)$  such that  $i \in \{1, \dots, C\}$ , and  $\mathcal{S}_i = (K_{i,h}, M_i, K_{i,t})$  has length at most  $s_i$ . Moreover,  $(M_i, K_{i,t})$  is a suffix of  $\mathcal{S}$ , and given  $i$  and  $K_{i,h}$ , the rest of  $\mathcal{S}$  (i.e.,  $\mathcal{S}[-1, |M| - |M_i| - 1]$ ) can be recovered in constant time, assuming random access to the lookup table.  $\text{range}(K_{i,h})$ ,  $\text{range}(K_{i,t})$  and  $|M_i|$  does not depend on  $\mathcal{S}$ , and can be stored in the lookup table.

*Remark.* We can safely omit any  $i$  with  $|M| - |M_i| > \kappa + 1$ , since removing such  $s_i$  from the list  $(s_1, \dots, s_C)$  (and decrease  $C$  by one) could only increase the upper bound on  $s$ ,  $\lg(2^{s_1} + \dots + 2^{s_C}) - C \cdot 2^{-\kappa+2}$ . That is, the extraction algorithm may never generate a pair with this  $i$ . Therefore, we may assume that  $\mathcal{S}[-1, |M| - |M_i| - 1]$  has length at most  $O(\kappa)$ , taking constant time to output.

## 5 Reductions to Perfect Hashing

Now, we show how to design succinct dictionary and compress low entropy sequence with local decodability using the membership and perfect hashing data structure in Theorem 2.

**Succinct dictionary.** For dictionary, we shall use the following lemma by Dodis, Pătraşcu and Thorup [DPT10].

**Lemma 16** ([DPT10]). *There is an algorithm that preprocesses a given sequence  $(x_1, \dots, x_n) \in [\sigma]^n$  for  $\sigma \leq 2^\kappa$  into a data structure of length at most  $n \lg \sigma + (n-1)2^{-\kappa+5}$ , such that given any  $i$ ,  $x_i$  can be retrieved in constant time.*

The data structure in [DPT10] requires a lookup table of  $O(\lg n)$  words. We show that by using the view of fractional-length strings, we can completely remove the lookup table.

*Proof.* Let  $b = \lceil 2\kappa / \lg \sigma \rceil$ . We partition the sequence into  $n/b$  chunks of  $b$  symbols each, then combine each chunk into one single character in  $[\sigma^b]$  (if  $n$  is not a multiple of  $b$ , then the last group will have more than  $b$  symbols). Since  $\sigma^b = 2^{O(\kappa)} = 2^{O(w)}$ , each  $x_i$  can be decode in constant time given the character. Then compute  $m = \lfloor \lg \sigma^b \rfloor - \kappa$  and  $R = \lceil \sigma^b \cdot 2^{-m} \rceil$ , and view each character in  $[\sigma^b]$  as a data structure of size  $m + \lg R$ . Note that  $m + \lg R - b \lg \sigma \leq \lg(\sigma^b + 2^m) - b \lg \sigma \leq \lg(1 + 2^{-\kappa}) \leq 2^{-\kappa+1}$ . Then we apply Proposition 7 to concatenate all  $n/b$  data structure. Since  $m$  and  $R$  can both be computed in constant time,  $m + \lg R$  can be approximated in constant time, hence Proposition 7 guarantees that there is a data structure of size

$$(m + \lg R) \cdot (n/b) + (n/b - 1) \cdot 2^{-\kappa+4} \leq n \lg \sigma + (n-1) \cdot 2^{-\kappa+5},$$

supporting symbol retrieval in constant time. This proves the lemma.  $\square$

To store a set of  $n$  key-value pairs for keys in  $[U]$  and values in  $[\sigma]$ , we first apply Theorem 2 on the set of keys. It produces a data structure of size  $\lg \binom{U}{n} + \text{poly} \lg n + O(\lg \lg U)$  bits, which defines a bijection  $h$  between the keys and  $[n]$ . Next, we apply Lemma 16 to store the values. Specifically, we construct the sequence  $(v_1, \dots, v_n)$  such that if  $(x, u)$  is an input key-value pair, then  $v_{h(x)+1} = u$ . This sequence can be stored in space  $n \lg \sigma + O(1)$  by Lemma 16. Hence, the total space of the data structure is

$$\lg \binom{U}{n} + n \lg \sigma + \text{poly} \lg n + O(\lg \lg U) = \text{OPT} + \text{poly} \lg n + O(\lg \lg U),$$

as claimed in Theorem 1.

To answer a query  $\text{valRet}(x)$ , we first query the membership data structure. If  $x$  is not a key, we return  $\perp$ . Otherwise, we retrieve and return the  $(h(x) + 1)$ -th value in the sequence. The total query time is constant in expectation and with high probability. This proves Theorem 1.

**Compression to zeroth order entropy with local decodability.** Given a sequence  $(x_1, \dots, x_n) \in \Sigma^n$  such that each  $\sigma \in \Sigma$  appears  $f_\sigma$  times, we construct a data structure *recursively on*  $\Sigma$ . We first arbitrarily partition  $\Sigma$  into  $\Sigma_1 \cup \Sigma_2$  such that  $|\Sigma_1| = \lfloor |\Sigma|/2 \rfloor$  and  $|\Sigma_2| = \lceil |\Sigma|/2 \rceil$ . For any set  $\Gamma \subseteq \Sigma$ , define  $S_\Gamma := \{i \in [n] : x_i \in \Gamma\}$ . Then we apply Theorem 2 to construct a perfect hashing for  $S_{\Sigma_1}$ , which uses space

$$\lg \binom{n}{|S_{\Sigma_1}|} + \text{poly} \lg n = \lg \left( \frac{n!}{|S_{\Sigma_1}|! \cdot |S_{\Sigma_2}|!} \right) + \text{poly} \lg n,$$

and defines a bijection  $h$  that maps all coordinates in  $S_{\Sigma_1}$  to  $[[S_{\Sigma_1}]]$ , and a bijection  $\bar{h}$  that maps all  $S_{\Sigma_2}$  to  $[[S_{\Sigma_2}]]$ . We recursively construct a data structure for  $\Sigma_1$  over  $h(S_{\Sigma_1})$ , and a data structure for  $\Sigma_2$  over  $\bar{h}([S_{\Sigma_2}])$ .

In general, each node in the recursion tree corresponds to a subset  $\Gamma$  of the alphabet such that  $\Gamma_1$  and  $\Gamma_2$  are the subsets corresponding to the left and the right child respectively. In this node, we store

- the size of subset of its left child  $|\Gamma_1|$ ,
- the perfect hashing data structure for  $S_{\Gamma_1}$ ,
- two pointers to the data structures in its left and its right children.

For  $\Gamma$  of size one, we store nothing. Thus, we obtain a final data structure of size

$$\lg \left( \frac{n!}{f_{\sigma_1}! f_{\sigma_2}! \dots} \right) + |\Sigma| \cdot \text{poly } \lg n.$$

To answer a query  $i$ , we first retrieve the size of  $\Sigma_1$  and query if  $i \in S_{\Sigma_1}$ . If  $i \in S_{\Sigma_1}$ , we go to the left child and recursively query  $h(i)$ . If  $i \notin S_{\Sigma_1}$ , we go to the right child and recursively query  $\bar{h}(i)$ . Finally, when the current subset  $|\Gamma| = 1$ , we return the only element in  $\Gamma$ . Since in each level of the recursion, the perfect hashing data structure takes constant query time in expectation, and the size of  $\Gamma$  reduces by a factor two, the total query time is  $O(\lg |\Sigma|)$  in expectation. This proves Theorem 3.

## 6 Perfect Hashing for Medium-Sized Sets

In this section, we present the minimal perfect hashing and membership data structure when the number of keys  $n$  is neither too large nor too small, focusing on the case where  $n \geq U^{1/12}$  and  $n \leq U - U^{1/12}$ . Generalizing to all  $n$  involves less new ideas, and we defer the proof of the main theorem to Section 8.

Recall that we wish to preprocess a set of  $n$  keys  $S \subseteq [U]$ , such that the data structure defines a bijection  $h$  between  $S$  and  $[n]$  and a bijection  $\bar{h}$  between  $[U] \setminus S$  and  $[U - n]$ . A query  $\text{hash}(x)$  returns a pair  $(b, v)$  such that

- if  $x \in S$ , then  $b = 1$  and  $v = h(x)$ ;
- if  $x \notin S$ , then  $b = 0$  and  $v = \bar{h}(x)$ .

We partition the universe  $[U]$  into pairs of blocks. For each pair, we construct a *main* data structure and an *auxiliary* data structure, such that the main data structure contains “most” of the information in the block and has fixed length, and the auxiliary data structure stores the remaining information (which unavoidably has variable length). Finally, we concatenate all data structures for all blocks.

Our main technical lemma is to construct such two (fractional-length) data structures for a pair of blocks of sizes  $V$  and  $V_{\text{sc}}$ .

**Lemma 17** (main technical lemma). *Let  $\kappa$  be the fineness parameter for fractional-length strings, and  $c$  be a constant positive integer. Let  $V \in [2\kappa^{2c-3}, 2^{\kappa/2}]$  and  $V_{\text{sc}} \geq 4\kappa^{c+1}$ . For any constant  $\epsilon > 0$ , there is a preprocessing algorithm  $\text{perfHashBlk}$ , query algorithms  $\text{qalgBlk}_{\text{main}}$ ,  $\text{qalgBlk}$  and lookup tables  $\text{tableBlk}_{V, V_{\text{sc}}}$  of size  $\tilde{O}(2^{\epsilon\kappa})$ . Given*

- a set  $S \subseteq [V]$  such that  $m := |S| \in [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3]$ ,
- a set  $S_{\text{sc}} \subseteq \{V, \dots, V + V_{\text{sc}} - 1\}$  and  $m_{\text{sc}} := |S_{\text{sc}}| \in [\kappa^{c+1}, 3\kappa^{c+1}]$ ,
- a random string  $\mathcal{R}$  of  $\kappa^{c+1}$  bits,

$\text{perfHashBlk}$  preprocesses  $S$  and  $S_{\text{sc}}$  into a pair of two (fractional-length) data structures  $\mathcal{D}_{\text{main}}$  and  $\mathcal{D}_{\text{aux}}$ , such that

(i)  $\mathcal{D}_{\text{main}}$  has length at most

$$\lg \binom{V}{\kappa^{2c-3}, \kappa^c} + \kappa^{2c-3} \cdot 2^{-\kappa/2+1};$$

(ii)  $\mathcal{D}_{\text{aux}}$  has length at most

$$\lg \binom{V}{m} + \lg \binom{V_{\text{sc}}}{m_{\text{sc}}} - \lg \binom{V}{\kappa^{2c-3}, \kappa^c} + \kappa^{c+1} 2^{-\kappa/2+2};$$

(iii)  $\mathcal{D}_{\text{main}}$  and  $\mathcal{D}_{\text{aux}}$  together define a bijection  $h$  between

$$S \cup S_{\text{sc}} \quad \text{and} \quad [m + m_{\text{sc}}],$$

and a bijection  $\bar{h}$  between

$$[V + V_{\text{sc}}] \setminus (S \cup S_{\text{sc}}) \quad \text{and} \quad [(V + V_{\text{sc}}) - (m + m_{\text{sc}})],$$

such that  $h(S) \supset [\kappa^{2c-3}]$  and  $\bar{h}([V] \setminus S) \supset [V - \kappa^{2c-3} - \kappa^c]$ ;

- (iv) given any  $x \in [V]$ ,  $\text{qalgBlk}_{\text{main}}(V, x)$  outputs  $\text{hash}(x)$  when  $x \in S$  and  $h(x) \in [\kappa^{2c-3}]$ , or when  $x \notin S$  and  $\bar{h}(x) \in [V - \kappa^{2c-3} - \kappa^c]$ , otherwise it outputs “unknown”; moreover, it only accesses  $\mathcal{D}_{\text{main}}$ ,  $\mathcal{R}$  and the lookup table  $\text{tableBlk}_{V, V_{\text{sc}}}$ , and it runs in constant time in the worst case;
- (v) for any  $x \in [V]$ , the probability that  $\text{qalgBlk}_{\text{main}}(V, x)$  outputs “unknown” is at most  $O(\kappa^{-c+3})$  over the randomness of  $\mathcal{R}$ ;
- (vi) given any  $x \in [V + V_{\text{sc}}]$ ,  $\text{qalgBlk}(V, m, V_{\text{sc}}, m_{\text{sc}}, x)$  computes  $\text{hash}(x)$ ; it accesses  $\mathcal{D}_{\text{main}}$ ,  $\mathcal{D}_{\text{aux}}$ ,  $\mathcal{R}$  and the lookup table  $\text{tableBlk}_{V, V_{\text{sc}}}$ , and it runs in  $O(\kappa^4)$  time.

*Remark.* Note that the size of  $\mathcal{D}_{\text{main}}$  does not depend on  $m$  or  $m_{\text{sc}}$ , and  $\text{qalgBlk}_{\text{main}}$  also does not need to know  $m$  or  $m_{\text{sc}}$  to answer a query. The total size of  $\mathcal{D}_{\text{main}}$  and  $\mathcal{D}_{\text{aux}}$  is approximately  $\lg \binom{V}{m} + \lg \binom{V_{\text{sc}}}{m_{\text{sc}}}$ , close to the optimal space given  $m$  and  $m_{\text{sc}}$ . The size of  $\mathcal{D}_{\text{main}}$  is also close to the optimum, as  $\text{qalgBlk}_{\text{main}}$  has to identify a set of  $\kappa^{2c-3}$  keys and  $V - \kappa^{2c-3} - \kappa^c$  non-keys by only accessing  $\mathcal{D}_{\text{main}}$ , which takes exactly  $\lg \binom{V}{\kappa^{2c-3}, V - \kappa^{2c-3} - \kappa^c} = \lg \binom{V}{\kappa^{2c-3}, \kappa^c}$  bits.

The proof of the lemma is deferred to Section 7. In this following, we present our data structure assuming this lemma. For simplicity of the notations, let

$$\text{OPT}_{V, m} := \lg \binom{V}{m}$$

be the information theoretical *optimal* space when storing a set of  $m$  keys over key space of size  $V$ .

**Theorem 18.** For any constant  $\epsilon > 0$  and constant integer  $c > 0$ , there is a preprocessing algorithm  $\text{perfHash}$ , a query algorithm  $\text{qAlg}$  and lookup tables  $\text{table}_{U, n}$  of size  $n^\epsilon$ , such that given

- a set  $S$  of  $n$  keys over the key space  $[U]$ , where  $n \geq U^{1/12}$  and  $n \leq U - U^{1/12}$ ,
- a uniformly random string  $\mathcal{R}$  of length  $O(\lg^{c+1} n)$ ,

$\text{perfHash}$  preprocesses  $S$  into a data structure  $\mathcal{D}$  of (worst-case) length

$$\text{OPT}_{U, n} + U^{-1},$$

such that  $\mathcal{D}$  defines a bijection  $h$  between  $S$  and  $[n]$  and a bijection  $\bar{h}$  between  $[U] \setminus S$  and  $[U - n]$ . Given access to  $\mathcal{D}$ ,  $\mathcal{R}$  and  $\text{table}_{U, n}$ , for any key  $x \in [U]$ ,  $\text{qAlg}(U, n, x)$  outputs  $\text{hash}(x)$  on a RAM with word-size  $w \geq \Omega(\lg U)$ , in time

- $O(1)$  with probability  $1 - O(\lg^{-c+4} U)$  and
- $O(\lg^7 U)$  in worst-case,

where the probability is taken over the random  $\mathcal{R}$ . In particular, the query time is constant in expectation and with high probability by setting  $c = 11$ .

*Remark.* When  $n < U^{1/12}$ , we could use a hash function to map the keys to  $n^2$  buckets with no collisions. We could apply this theorem with the new key space being all buckets, and the keys being the non-empty buckets. By further storing for each non-empty bucket, the key within it, it extends the membership query to  $n < U^{1/12}$ , using  $O(\lg n + \lg \lg U)$  extra bits. We will see a more generic approach in Section 8 (which works for perfect hashing and improves the  $\lg \lg U$  term).

Without loss of generality, we may assume  $n \leq U/2$ , since otherwise, we could simply take the complement of  $S$ . Let  $\kappa := \lceil 4 \lg U \rceil$  be the fineness parameter, and  $c$  be a (large) constant positive integer to be specified later. We partition the universe  $[U]$  into pairs of blocks. Each block pair consists of a larger *primary* block containing roughly  $\kappa^{2c-3} + \kappa^c/2$  keys, and a smaller *secondary* block containing roughly  $2\kappa^{c+1}$  keys. Formally, let

$$V_{\text{pr}} := \lfloor \frac{(\kappa^{2c-3} + \kappa^c/2)U}{n} \rfloor,$$

$$V_{\text{sc}} := \lfloor \frac{2\kappa^{c+1}U}{n} \rfloor$$

and  $V_{\text{bl}} = V_{\text{pr}} + V_{\text{sc}}$ . Each primary block has size  $V_{\text{pr}}$  and each secondary block has size  $V_{\text{sc}}$ . Every block pair has size  $V_{\text{bl}}$ . For simplicity, let us first consider the case where  $U$  is a multiple of  $V_{\text{bl}}$ , and  $U = V_{\text{bl}} \cdot N_{\text{bl}}$ . We will show how to handle general  $U$  later.

Thus, we partition  $U$  into  $N_{\text{bl}}$  block pairs in the natural way, where the  $i$ -th primary block

$$\mathcal{B}_{\text{pr}}^i := \{x \in [U] : (i-1)V_{\text{bl}} \leq x < V_{\text{pr}} + (i-1)V_{\text{bl}}\}$$

and the  $i$ -th secondary block

$$\mathcal{B}_{\text{sc}}^i := \{x \in [U] : V_{\text{pr}} + (i-1)V_{\text{bl}} \leq x < iV_{\text{bl}}\}.$$

We call the  $i$ -th block pair *good*, if the numbers of keys in the primary and secondary blocks are close to the average:

$$|S \cap \mathcal{B}_{\text{pr}}^i| \in [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3],$$

and

$$|S \cap \mathcal{B}_{\text{sc}}^i| \in [\kappa^{c+1}, 3\kappa^{c+1}].$$

The pair is *bad* if at least one of the two blocks has the number of keys outside the range.

In the following, we show that for inputs  $S$  with no bad blocks, we can construct a good data structure. The goal is to design a data structure using space close to  $\text{OPT}_{N_{\text{bl}}V_{\text{bl}},n}$ .

## 6.1 No bad block pair

**Lemma 19.** *If  $U$  is a multiple of  $V_{\text{bl}}$ , then there is a data structure with the guarantees in Theorem 18, for all sets  $S$  with no bad block pair. Moreover, the size of the data structure is*

$$\text{OPT}_{U,n} + n \cdot 2^{-\kappa/2+2}.$$

Before starting the preprocessing, we view the last  $O(\lg U)$  bits of the random bits  $\mathcal{R}$  as a random number  $\Delta \in [U]$ . We shift the entire universe according to  $\Delta$ , i.e.,  $x \mapsto (x + \Delta) \bmod U$ . It is applied to input  $S$ , and will be applied to the queries too (which guarantees that the query is in a primary block with good probability).

The preprocessing algorithm is based on recursion. The following algorithm `dict_rec` preprocesses  $S$  restricted to the  $i$ -th to  $j$ -th blocks  $\mathcal{B}_{\text{pr}}^i, \mathcal{B}_{\text{sc}}^i, \dots, \mathcal{B}_{\text{pr}}^j, \mathcal{B}_{\text{sc}}^j$ , and outputs  $j-i+2$  data structures  $\mathcal{D}_{\text{main}}^i, \dots, \mathcal{D}_{\text{main}}^j$  and  $\mathcal{D}_{\text{aux}}$ . We will inductively prove upper bounds on the sizes of the data structures: the length of each  $\mathcal{D}_{\text{main}}^i$  is at most

$$\text{SIZE}_{\text{main}} := \lg \left( \frac{V_{\text{pr}}}{\kappa^{2c-3}}, \kappa^c \right) + \kappa^{2c-3} \cdot 2^{-\kappa/2+1}, \quad (1)$$

and the length of  $\mathcal{D}_{\text{aux}}$  generated from  $i, \dots, j$ -th block pair is at most

$$\text{OPT}_{(j-i+1)V_{\text{bl}}, m} - (j-i+1)\text{SIZE}_{\text{main}} + (m-1)2^{-\kappa/2+2}, \quad (2)$$

where  $m$  is the number of keys in blocks  $i$  to  $j$ . In the base case with only one block pair, we simply apply Lemma 17.

**preprocessing algorithm** `dict_rec`( $i, j, m, S, \mathcal{R}$ ):

1. if  $i = j$
2. let  $S_{\text{pr}} \subseteq S$  be the set of keys in the  $i$ -th primary block
3. let  $S_{\text{sc}} \subseteq S$  be the set of keys in the  $i$ -th secondary block
4.  $m_{\text{pr}} := |S_{\text{pr}}|$  and  $m_{\text{sc}} := |S_{\text{sc}}|$
5.  $(\mathcal{D}_{\text{main}}^i, \mathcal{D}'_{\text{aux}}) := \text{perfHashBlk}(V_{\text{pr}}, m_{\text{pr}}, V_{\text{sc}}, m_{\text{sc}}, S_{\text{pr}}, S_{\text{sc}}, \mathcal{R})$  (from Lemma 17)
6. apply Proposition 10 to fuse  $m_{\text{pr}}$  into  $\mathcal{D}'_{\text{aux}}$ , and obtain  $\mathcal{D}_{\text{aux}}$
7. return  $(\mathcal{D}_{\text{main}}^i, \mathcal{D}_{\text{aux}})$  (to be cont'd)

**Claim 20.** If  $i = j$ ,  $|\mathcal{D}_{\text{main}}^i| \leq \text{SIZE}_{\text{main}}$  and  $|\mathcal{D}_{\text{aux}}| \leq \text{OPT}_{V_{\text{bl}}, m} - \text{SIZE}_{\text{main}} + (m-1) \cdot 2^{-\kappa/2+2}$ .

To prove the claim, note that the premises of Lemma 17 are satisfied: since  $2n \leq U$ ,  $V_{\text{pr}} \geq 2\kappa^{2c-3}$  and  $V_{\text{pr}} \leq U \leq 2^{\kappa/2}$ ;  $V_{\text{sc}} \geq 4\kappa^{c+1}$ ; by assumption, every primary block has between  $\kappa^{2c-3} + \kappa^c/3$  and  $\kappa^{2c-3} + 2\kappa^c/3$  keys, and every secondary block has between  $\kappa^{c+1}$  and  $3\kappa^{c+1}$  keys. Therefore, by Lemma 17, the size of  $\mathcal{D}_{\text{main}}^i$  is at most  $\text{SIZE}_{\text{main}}$ , and the size of  $\mathcal{D}'_{\text{aux}}$  is at most

$$\lg \left( \frac{V_{\text{pr}}}{m_{\text{pr}}} \right) + \lg \left( \frac{V_{\text{sc}}}{m - m_{\text{pr}}} \right) - \text{SIZE}_{\text{main}} + \kappa^{2c-3} \cdot 2^{-\kappa/2+2}.$$

By fusing the value of  $m_{\text{pr}}$  into the data structure, the size of  $\mathcal{D}_{\text{aux}}$  is at most

$$\text{OPT}_{V_{\text{bl}}, m} - \text{SIZE}_{\text{main}} + (m-1) \cdot 2^{-\kappa/2+2},$$

due to the fact that  $m \geq \kappa^{2c-3} + \kappa^{c+1}$  and  $\sum_{m_{\text{pr}}} \binom{V_{\text{pr}}}{m_{\text{pr}}} \binom{V_{\text{sc}}}{m - m_{\text{pr}}} = \lg \binom{V_{\text{bl}}}{m} = \text{OPT}_{V_{\text{bl}}, m}$ .  $\mathcal{D}_{\text{main}}^i$  and  $\mathcal{D}_{\text{aux}}$  both have sizes as claimed in (1) and (2). Also, note that we give the same random string  $\mathcal{R}$  to all block pairs. Thus, the total number of random bits needed is  $\kappa^{c+1}$  by Lemma 17.

Next, when  $i < j$ , we recurse on the two halves of the block pairs, and merge them.

8.  $k := \lfloor (i+j)/2 \rfloor$
9. let  $m_1$  be the number of keys in the  $i$ -th,  $\dots, k$ -th block pair
10. let  $m_2$  be the number of keys in the  $(k+1)$ -th,  $\dots, j$ -th block pair
11. recurse on the two halves:
  - $(\mathcal{D}_{\text{main}}^i, \dots, \mathcal{D}_{\text{main}}^k, \mathcal{D}_{\text{aux},1}) := \text{dict\_rec}(i, k, m_1, S, \mathcal{R})$
  - $(\mathcal{D}_{\text{main}}^{k+1}, \dots, \mathcal{D}_{\text{main}}^j, \mathcal{D}_{\text{aux},2}) := \text{dict\_rec}(k+1, j, m_2, S, \mathcal{R})$
12. apply Proposition 7 to concatenate  $\mathcal{D}_{\text{aux},1}$  and  $\mathcal{D}_{\text{aux},2}$ , and obtain  $\mathcal{D}'_{\text{aux}}$
13. apply Proposition 9 to fuse the value of  $m_1$  into  $\mathcal{D}'_{\text{aux}}$  for  $m_1 \in \{0, \dots, m\}$ , and obtain  $\mathcal{D}_{\text{aux}}$
14. return  $(\mathcal{D}_{\text{main}}^i, \dots, \mathcal{D}_{\text{main}}^j, \mathcal{D}_{\text{aux}})$

**Claim 21.**  $|\mathcal{D}_{\text{main}}^i| \leq \text{SIZE}_{\text{main}}$  for all  $i$ , and  $|\mathcal{D}_{\text{aux}}| \leq \text{OPT}_{(j-i+1)V_{\text{bl}},m} - \text{SIZE}_{\text{main}} + (j-i+1)(m-1) \cdot 2^{-\kappa/2+2}$ .

We have already showed that each  $|\mathcal{D}_{\text{main}}^i| \leq \text{SIZE}_{\text{main}}$  above. To prove the bound on  $|\mathcal{D}_{\text{aux}}|$ , by inductive hypothesis, we know that  $\mathcal{D}_{\text{aux},1}$  has size at most

$$\text{OPT}_{(k-i+1)V_{\text{bl}},m_1} - (k-i+1)\text{SIZE}_{\text{main}} + (m_1-1)2^{-\kappa/2+2}$$

and  $\mathcal{D}_{\text{aux},2}$  has size at most

$$\text{OPT}_{(j-k)V_{\text{bl}},m_2} - (j-k)\text{SIZE}_{\text{main}} + (m_2-1)2^{-\kappa/2+2}.$$

To apply Proposition 7 in line 12, it requires us to approximate the data structure sizes. The following claim implies that the premises can be satisfied.

**Claim 22.** Both  $\text{OPT}_{(k-i+1)V_{\text{bl}},m_1} - (k-i+1)\text{SIZE}_{\text{main}} + (m_1-1)2^{-\kappa/2+2}$  and  $\text{OPT}_{(j-k)V_{\text{bl}},m_2} - (j-k)\text{SIZE}_{\text{main}} + (m_2-1)2^{-\kappa/2+2}$  can be approximated with an additive error of at most  $2^{-\kappa}$  in  $O(1)$  time.

Assuming Claim 22, Proposition 7 concatenates  $\mathcal{D}_{\text{aux},1}$  and  $\mathcal{D}_{\text{aux},2}$  into a data structure  $\mathcal{D}'_{\text{aux}}$  of length at most

$$\text{OPT}_{(k-i+1)V_{\text{bl}},m_1} + \text{OPT}_{(j-k)V_{\text{bl}},m-m_1} - (j-i+1)\text{SIZE}_{\text{main}} + (m-2)2^{-\kappa/2+2} + 2^{-\kappa+4}.$$

The following claim implies that the premises of Proposition 9 from line 13 can be satisfied, because  $-(j-i+1)\text{SIZE}_{\text{main}} + (m-2)2^{-\kappa/2+2} + 2^{-\kappa+4}$  does not depend on  $m_1$ , and can be computed efficiently.

**Claim 23.** For any  $V_1, V_2, m \geq 0$ , and  $0 \leq l \leq m$ ,  $\sum_{i=0}^l 2^{\text{OPT}_{V_1,i} + \text{OPT}_{V_2,m-i}}$  can be approximated up to an additive error of at most  $2^{-\kappa-3}$ .  $\sum_{i=0}^m 2^{\text{OPT}_{V_1,i} + \text{OPT}_{V_2,m-i}}$  in  $O(\kappa^5)$  time.

The proofs of both claims are deferred to Appendix B. Assuming Claim 23, Proposition 9 fuses  $m_1$  into  $\mathcal{D}'_{\text{aux}}$ , and obtains  $\mathcal{D}_{\text{aux}}$  of length at most

$$\begin{aligned} & \lg \left( \sum_{l=0}^m 2^{\text{OPT}_{(k-i+1)V_{\text{bl}},l} + \text{OPT}_{(j-k)V_{\text{bl}},m-l} - (j-i+1)\text{SIZE}_{\text{main}} + (m-2)2^{-\kappa/2+2} + 2^{-\kappa+4}} \right) + (m+1) \cdot 2^{-\kappa+4} \\ & \leq \lg \left( \sum_{l=0}^m 2^{\text{OPT}_{(k-i+1)V_{\text{bl}},l} + \text{OPT}_{(j-k)V_{\text{bl}},m-l}} \right) - (j-i+1)\text{SIZE}_{\text{main}} + (m-1)2^{-\kappa/2+2} \\ & = \text{OPT}_{(j-i+1)V_{\text{bl}},m} - (j-i+1)\text{SIZE}_{\text{main}} + (m-1)2^{-\kappa/2+2} \end{aligned}$$

This proves Claim 21.

Thus, by induction, `dict_rec` outputs  $\mathcal{D}_{\text{main}}^1, \dots, \mathcal{D}_{\text{main}}^{N_{\text{bl}}}$  of length  $\text{SIZE}_{\text{main}}$  and  $\mathcal{D}_{\text{aux}}$  of length

$$\text{OPT}_{N_{\text{bl}}V_{\text{bl}},n} - N_{\text{bl}} \cdot \text{SIZE}_{\text{main}} + (n-1)2^{-\kappa/2+2}.$$

Finally, we apply Proposition 7 again to concatenate all  $N_{\text{bl}} + 1$  data structures. By storing approximations of sizes of  $\mathcal{D}_{\text{main}}^i$  and  $\mathcal{D}_{\text{aux}}$  in the lookup table, we obtain a data structure of length at most

$$\begin{aligned} & N_{\text{bl}} \cdot \text{SIZE}_{\text{main}} + (\text{OPT}_{N_{\text{bl}}V_{\text{bl}},n} - N_{\text{bl}} \cdot \text{SIZE}_{\text{main}} + (n-1)2^{-\kappa/2+2}) + N_{\text{bl}} \cdot 2^{-\kappa+3} \\ & \leq \text{OPT}_{N_{\text{bl}}V_{\text{bl}},n} + n \cdot 2^{-\kappa/2+2}. \end{aligned}$$

This proves the space bound in Lemma 19.

**Hash functions.** Let  $h_i$  and  $\bar{h}_i$  be the bijections obtained by Lemma 17 for blocks  $\mathcal{B}_{\text{pr}}^i$  and  $\mathcal{B}_{\text{sc}}^i$ . We define the bijections  $h$  and  $\bar{h}$  as follows:

- for key  $x \in S \cap (\mathcal{B}_{\text{pr}}^i \cup \mathcal{B}_{\text{sc}}^i)$ , if  $h_i(x) < \kappa^{2c-3}$ , let  $h(x) := (i-1)\kappa^{2c-3} + h_i(x)$ , otherwise, let  $h(x) := (N_{\text{bl}} - i) \cdot \kappa^{2c-3} + \sum_{j < i} |(\mathcal{B}_{\text{pr}}^j \cup \mathcal{B}_{\text{sc}}^j) \cap S| + h_i(x)$ ;
- for non-key  $x \notin S$ , if  $\bar{h}_i(x) < V_{\text{pr}} - \kappa^{2c-3} - \kappa^c$ , let  $\bar{h}(x) := (i-1)(V_{\text{pr}} - \kappa^{2c-3} - \kappa^c) + \bar{h}_i(x)$ , otherwise, let  $\bar{h}(x) := (N_{\text{bl}} - i) \cdot (V_{\text{pr}} - \kappa^{2c-3} - \kappa^c) + \sum_{j < i} |(\mathcal{B}_{\text{pr}}^j \cup \mathcal{B}_{\text{sc}}^j) \setminus S| + \bar{h}_i(x)$ .

Essentially, the smallest hash values will be those with  $h_i(x) < \kappa^{2c-3}$  or  $\bar{h}_i(x) < V_{\text{pr}} - \kappa^{2c-3} - \kappa^c$ , ordered according to  $i$  and  $h_i(x)$  or  $\bar{h}_i(x)$ . Then the rest take larger values ordered according to  $i$  and  $h_i(x)$  or  $\bar{h}_i(x)$ . By definition, they are both bijections.

**Lookup tables.** We store the following information in the lookup table.

**lookup table tbl:**

1. `tableBlk $_{V_{\text{pr}}, V_{\text{sc}}}$`  from Lemma 17
2. the lookup table for line 6 from Proposition 10 for all valid values  $m_{\text{pr}}$  and  $m_{\text{sc}}$
3. approximated value of  $\text{SIZE}_{\text{main}}$  and the (final) size of  $\mathcal{D}_{\text{aux}}$ , up to  $O(\kappa)$  bits of precision

By Lemma 17, the lookup table size is  $2^{\epsilon\kappa}$ . Since  $\kappa = O(\lg U)$  and  $n \geq U^{1/12}$ , by readjusting the constant  $\epsilon$ , the lookup table size is at most  $n^\epsilon$ .

**Query algorithm.** Now, we show how to answer hash queries. Given a query  $x \in [U]$ , we first shift it according to  $\Delta$ , as we did at preprocessing,  $x \mapsto (x + \Delta) \bmod U$ . If  $x$  is in a primary block, we query the corresponding main data structure. If the main data structure does not return the answer, or  $x$  is not in a primary block, we recursively decode the corresponding auxiliary data structure, and run `qalgBlk`.

**query algorithm qalgG( $U, n, x$ ):**

1. if  $x$  is in the  $i$ -th primary block
2. apply Proposition 7 to decode  $\mathcal{D}_{\text{main}}^i$
3. if  $(b, v) := \mathcal{D}_{\text{main}}^i.\text{qalgBlk}_{\text{main}}(V_{\text{pr}}, x) \neq \text{"unknown"}$  (from Lemma 17)
4. if  $b = 1$ , return  $(1, (i-1)\kappa^{2c-3} + v)$
5. if  $b = 0$ , return  $(0, (i-1)(V_{\text{pr}} - \kappa^{2c-3} - \kappa^c) + v)$
6. decode  $\mathcal{D}_{\text{aux}}$  and return  $\mathcal{D}_{\text{aux}}.\text{qalg\_rec}(1, N_{\text{bl}}, 0, n, x)$

Since  $V_{\text{pr}}/V_{\text{sc}} = O(\kappa^{c-4})$  and we randomly shifted the universe,  $x$  is in a primary block with probability  $1 - O(\kappa^{-c+4})$ . Also, by Lemma 17, `qalgBlk $_{\text{main}}$`  runs in constant time. It returns “unknown” with probability at most  $O(\kappa^{-c+3})$  for a uniformly random  $\mathcal{R}$ , and returns `hash( $x$ )` otherwise. Therefore, the probability that `qalgG` terminates before reaching the last line is  $1 - O(\kappa^{-c+4})$ . Since  $\kappa = \Theta(\lg U)$ , it computes `hash( $x$ )` in constant time with probability  $1 - O(\lg^{-c+4} U)$ .

Next, we show how to implement `qalg_rec( $i, j, s, m, x$ )`, which takes as parameters

- $(i, j)$ : a range of blocks,
- $s$ : the total number of keys before block  $i$ ,
- $m$ , the total number of keys in blocks  $i$  to  $j$ , and
- $x$ , the element being queried.

We will prove that its worst-case running time is  $O(\lg^7 U)$ .

<p><b>query algorithm</b> <code>qalg_rec</code>(<math>i, j, s, m, x</math>):</p> <ol style="list-style-type: none"> <li>1. if <math>i = j</math></li> <li>2.   apply Proposition 10 to decode <math>m_{\text{pr}}</math> and <math>\mathcal{D}'_{\text{aux}}</math></li> <li>3.   <math>(b, v) := (\mathcal{D}_{\text{main}}^i, \mathcal{D}'_{\text{aux}}).\text{qalgBlk}(V_{\text{pr}}, m_{\text{pr}}, V_{\text{sc}}, m - m_{\text{pr}}, x - (i - 1)(V_{\text{pr}} + V_{\text{sc}}))</math>  <span style="float: right;">(from Lemma 17)</span></li> <li>4.   if <math>b = 1</math>, return <math>(1, (N_{\text{bl}} - i) \cdot \kappa^{2c-3} + s + v)</math></li> <li>5.   if <math>b = 0</math>, return <math>(0, (N_{\text{bl}} - i) \cdot (V_{\text{pr}} - \kappa^{2c-3} - \kappa^c) + ((i - 1) \cdot V_{\text{bl}} - s) + v)</math> <span style="float: right;">(to be cont'd)</span></li> </ol>
--

In the base case with only one block, we simply decode the value of  $m_{\text{pr}}$  as well as the corresponding  $\mathcal{D}'_{\text{aux}}$  from the  $i$ -th block pair. By running `qalgBlk` from Lemma 17 to query within the block pair, we compute  $\text{hash}(x)$  according to its definition in  $O(\kappa^4)$  time.

<ol style="list-style-type: none"> <li>6. <math>k := \lfloor (i + j)/2 \rfloor</math></li> <li>7. apply Proposition 9 and Claim 23 to decode <math>m_1</math> and <math>\mathcal{D}'_{\text{aux}}</math></li> <li>8. apply Proposition 7 and Claim 22 to decode <math>\mathcal{D}_{\text{aux},1}</math> and <math>\mathcal{D}_{\text{aux},2}</math></li> <li>9. if <math>x</math> is in <math>i</math>-th, <math>\dots</math>, <math>k</math>-th block pair</li> <li>10.   return <math>\mathcal{D}_{\text{aux},1}.\text{qalg_rec}(i, k, s, m_1, x)</math></li> <li>11. else</li> <li>12.   return <math>\mathcal{D}_{\text{aux},2}.\text{qalg_rec}(k + 1, j, s + m_1, m - m_1, x)</math></li> </ol>
--

In general, we decode  $m_1$ , the number of elements in the first half of the blocks. Then we decode the data structures for the two halves. Depending on where the query is, we recurse into one of the two data structures. Proposition 7, Proposition 9, Claim 22 and Claim 23 guarantee that the decoding takes  $O(\kappa^6)$  time. The recursion has at most  $O(\lg n) \leq \kappa$  levels. Thus, the total running time of `qalg_rec` is at most  $O(\kappa^7)$ . This proves the claim on the query time, and hence, it proves Lemma 19.

## 6.2 At least one bad block pair

Now, let us show how to handle sets with at least one bad block. We will show that the space usage for such sets is  $\text{OPT}_{U,n} - \Omega(\kappa^3)$ .

**Lemma 24.** *If  $U$  is a multiple of  $V_{\text{bl}}$ , then there is a data structure with guarantees in Theorem 18, for all sets with at least one bad block pair. Moreover, the size of the data structure is at most*

$$\text{OPT}_{U,n} - \Omega(\kappa^3).$$

Note that this is possible, because by Chernoff bound, there are only at most  $2^{-\Omega(\kappa^3)}$  fraction such sets, and we can even afford to spend at least  $O(\kappa^3) = O(\lg^3 U)$  extra bits. The first  $\lceil \lg N_{\text{bl}} \rceil$  bits are used to encode the number of bad block pairs  $N_{\text{bad}}$ . It turns out that the fraction of input sets with  $N_{\text{bad}}$  bad pairs is  $2^{-\Omega(\kappa^3 N_{\text{bad}})}$ , as we mentioned in Section 2. By the argument there, we can afford to use  $O(\kappa^3 N_{\text{bad}})$  extra bits.

The idea is to construct a mapping which maps all good block pairs to the first  $N_{\text{bl}} - N_{\text{bad}}$  pairs, construct a data structure using the above algorithm for good blocks, and finally handle the bad pairs separately.

To construct such a mapping, observe that the following two numbers are equal:

- (a) the number of *good* pairs among the last  $N_{\text{bad}}$  pairs, and
- (b) the number of *bad* pairs among the first  $N_{\text{bl}} - N_{\text{bad}}$  pairs.

Hence, in the mapping, we map all the good pairs among the last  $N_{\text{bad}}$  to all bad pairs among the first  $N_{\text{bl}} - N_{\text{bad}}$ . The good pairs in the first  $N_{\text{bl}} - N_{\text{bad}}$  pairs will be mapped to themselves. To store such a mapping, we spend  $O(N_{\text{bad}} \cdot \lg N_{\text{bl}})$  bits to store a hash table of all the bad pairs using the FKS hashing.

Then we spend  $O(N_{\text{bad}} \cdot \lg N_{\text{bl}})$  bits to store for each pair in the last  $N_{\text{bad}}$  pairs, whether it is a good pair and if it is, which bad pair it will be mapped to. The mapping takes  $O(N_{\text{bad}} \cdot \lg N_{\text{bl}})$  bits to store in total (which is much smaller than  $\kappa^3 N_{\text{bad}}$ ). It takes constant time to evaluate.

This mapping maps all good pairs to the first  $N_{\text{bl}} - N_{\text{bad}}$  pairs. Then we apply `dict_rec` for good pairs from Lemma 19 to construct a data structure using

$$\text{OPT}_{(N_{\text{bl}} - N_{\text{bad}})V_{\text{bl}}, n - n_{\text{bad}}} + (n - n_{\text{bad}}) \cdot 2^{-\kappa/2+2} \leq \lceil \text{OPT}_{(N_{\text{bl}} - N_{\text{bad}})V_{\text{bl}}, n - n_{\text{bad}}} \rceil + 1$$

bits, where  $n_{\text{bad}}$  is the number of keys in the bad block pairs.

Next, we construct data structures for the bad pairs. Consider a bad pair with  $m_{\text{pr}}$  keys in the primary block and  $m_{\text{sc}}$  keys in the secondary block. Thus, either  $m_{\text{pr}} \notin [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3]$ , or  $m_{\text{sc}} \notin [\kappa^{c+1}, 3\kappa^{c+1}]$ . We construct two separate data structures, one for the primary block and one for secondary block (note that it might be the case that the number of keys in the primary block is within the above range, but the block pair is bad due to the secondary block, or vice versa, we still construct two separate data structures for *both* of them using the following argument). It turns out that if the number of keys in the block is at most  $\kappa^{O(1)}$ , then there is a data structure using only  $O(1)$  extra bit, answering queries in constant time.

**Lemma 25.** *Let  $c$  be any constant positive integer and  $\epsilon$  be any positive constant. There is a preprocessing algorithm `perfHashS`, query algorithm `qalgS` and lookup tables `tableSV,m` of sizes  $\tilde{O}(2^{\epsilon\kappa})$ , such that for any  $V \leq 2^{\kappa/2}$  and  $m \leq \kappa^c$ , given a set  $S \subset [V]$  of  $m$  keys, `perfHashS` preprocesses  $S$  into a data structure of size at most*

$$\text{OPT}_{V,m} + (m - 1) \cdot 2^{-\kappa/2+1},$$

*such that it defines a bijection  $h$  between  $S$  and  $[m]$  and a bijection  $\bar{h}$  between  $[V] \setminus S$  and  $[V - m]$ . Given any  $x \in [V]$ , `qalgS` answers `hash(x)` in constant time, by accessing the data structure and `tableSV,m`.*

In particular, the size is at most  $\text{OPT}_{V,m} + O(1)$ . The lemma is an immediate corollary of Lemma 28 in Section 7.2. Its proof is deferred to Section 7.2.

On the other hand, sets that have a block with more than  $\kappa^{3c}$  keys are even more rare. By Chernoff bound, we can estimate that the fraction of sets with at least one block with  $m > \kappa^{3c}$  keys is at most  $2^{-\Omega(m \lg m)}$ . This suggests that for every (bad) block with  $m > \kappa^{3c}$  keys, we can afford to spend  $O(m \lg m)$  extra bits. A simple modification to [Pag01a] gives such a data structure.

**Lemma 26.** *Given a set  $S \subset [V]$  of  $m$  keys, there is a data structure of size*

$$\text{OPT}_{V,m} + O(m + \lg \lg V),$$

*such that it defines a bijection  $h$  between  $S$  and  $[m]$  and a bijection  $\bar{h}$  between  $[V] \setminus S$  and  $[V - m]$ . It supports `hash` queries in constant time.*

Note that  $\lg \lg V \leq \lg \kappa$ , and  $m \geq \kappa^{3c}$ . The number of extra bits is simply  $O(m)$ . We prove this lemma in Appendix C.

For each bad block pair, we write down the two numbers  $m_{\text{pr}}$  and  $m_{\text{sc}}$  using  $O(\lg n)$  bits. Then if  $m_{\text{pr}} \leq \kappa^{3c}$ , we apply Lemma 25, and obtain a data structure with

$$\text{OPT}_{V_{\text{pr}}, m_{\text{pr}}} + O(1)$$

bits. If  $m_{\text{pr}} > \kappa^{3c}$ , we apply Lemma 26, and obtain a data structure with

$$\text{OPT}_{V_{\text{pr}}, m_{\text{pr}}} + O(m_{\text{pr}})$$

bits. Likewise for the secondary block, we obtain a data structure with

$$\text{OPT}_{V_{\text{sc}}, m_{\text{sc}}} + O(1)$$

bits if  $m_{\text{sc}} \leq \kappa^{3c}$ , and

$$\text{OPT}_{V_{\text{sc}}, m_{\text{sc}}} + O(m_{\text{sc}})$$

bits if  $m_{\text{sc}} > \kappa^{3c}$ .

Finally, we concatenate all data structures (which now all have integer lengths), and for each bad pair, we further store a pointer pointing to its corresponding data structure, as well as the total number of keys in all *bad* blocks prior to it (which helps us compute the hash values).

**Space usage.** Let us first bound the space usage for the data structures for a bad block pair. Consider the  $i$ -th bad block pair, suppose it has  $m_{\text{pr},i}$  keys in the primary block and  $m_{\text{sc},i}$  keys in the secondary block. Then note that we have

$$\lg \binom{V_{\text{pr}}}{m_{\text{pr},i}} = \lg \frac{V_{\text{pr}}!}{m_{\text{pr},i}!(V_{\text{pr}} - m_{\text{pr},i})!}$$

which by Sterling's formula, is at most

$$\begin{aligned} &\leq \lg \left( \frac{V_{\text{pr}}}{e} \right)^{V_{\text{pr}}} - \lg \left( \frac{m_{\text{pr},i}}{e} \right)^{m_{\text{pr},i}} - \lg \left( \frac{V_{\text{pr}} - m_{\text{pr},i}}{e} \right)^{V_{\text{pr}} - m_{\text{pr},i}} + O(\lg V_{\text{pr}}) \\ &= V_{\text{pr}} \lg V_{\text{pr}} - m_{\text{pr},i} \lg m_{\text{pr},i} - (V_{\text{pr}} - m_{\text{pr},i}) \lg (V_{\text{pr}} - m_{\text{pr},i}) + O(\lg V_{\text{pr}}). \end{aligned} \quad (3)$$

In the following, we are going to compare (3) with

$$V_{\text{pr}} \lg V_{\text{pr}} - m_{\text{pr},i} \lg \overline{m_{\text{pr}}} - (V_{\text{pr}} - m_{\text{pr},i}) \lg (V_{\text{pr}} - \overline{m_{\text{pr}}}) + O(\lg V_{\text{pr}}), \quad (4)$$

where  $\overline{m_{\text{pr}}} = \kappa^{2c-3} + \kappa^c/2$  is the average number of keys in a primary block. First observe that  $f(x) = m \lg x + (V - m) \lg(V - x)$  achieves its maximum at  $x = m$ , thus, (3)  $\leq$  (4). On the other hand, if  $m_{\text{pr},i} \notin [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3]$ , i.e.,  $m_{\text{pr},i}$  is far from  $\overline{m_{\text{pr}}}$ , (3) is even smaller.

**Claim 27.** We have (4)  $\geq$  (3). Moreover, if  $m_{\text{pr},i} \notin [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3]$ ,

$$(4) - (3) \geq \begin{cases} \Omega(\kappa^3) & m_{\text{pr},i} \leq \kappa^{3c}, \\ m_{\text{pr},i} \lg \kappa & m_{\text{pr},i} > \kappa^{3c}. \end{cases}$$

*Proof.*

$$\begin{aligned} (4) - (3) &= -m_{\text{pr},i} \lg \frac{\overline{m_{\text{pr}}}}{m_{\text{pr},i}} - (V_{\text{pr}} - m_{\text{pr},i}) \lg \frac{V_{\text{pr}} - \overline{m_{\text{pr}}}}{V_{\text{pr}} - m_{\text{pr},i}} \\ &= -m_{\text{pr},i} \lg \left( 1 + \frac{\overline{m_{\text{pr}}} - m_{\text{pr},i}}{m_{\text{pr},i}} \right) - (V_{\text{pr}} - m_{\text{pr},i}) \lg \left( 1 + \frac{m_{\text{pr},i} - \overline{m_{\text{pr}}}}{V_{\text{pr}} - m_{\text{pr},i}} \right). \end{aligned}$$

By the facts that  $\ln(1+x) \leq x$  for  $x > -1$ ,  $\ln(1+x) \leq x - \frac{1}{4}x^2$  for  $|x| \leq 1$ , and  $\ln(1+x) \leq 3x/4$  for  $x > 1$ , when  $m_{\text{pr},i} \leq \overline{m_{\text{pr}}}/2$ , we have

$$(4) - (3) = -m_{\text{pr},i} \lg \left( 1 + \frac{\overline{m_{\text{pr}}} - m_{\text{pr},i}}{m_{\text{pr},i}} \right) - (V_{\text{pr}} - m_{\text{pr},i}) \lg \left( 1 + \frac{m_{\text{pr},i} - \overline{m_{\text{pr}}}}{V_{\text{pr}} - m_{\text{pr},i}} \right)$$

$$\begin{aligned}
&\geq -m_{\text{pr},i} \cdot \frac{3}{4} \cdot \frac{\overline{m}_{\text{pr}} - m_{\text{pr},i}}{m_{\text{pr},i}} \lg e - (V_{\text{pr}} - m_{\text{pr},i}) \cdot \frac{m_{\text{pr},i} - \overline{m}_{\text{pr}}}{V_{\text{pr}} - m_{\text{pr},i}} \lg e \\
&= \frac{\lg e}{4} (\overline{m}_{\text{pr}} - m_{\text{pr},i}) \\
&\geq \Omega(\kappa^{2c-3}).
\end{aligned}$$

When  $m_{\text{pr},i} \geq \overline{m}_{\text{pr}}/2$  and  $m_{\text{pr},i} \leq \kappa^{3c}$ , we have

$$\begin{aligned}
(4) - (3) &= -m_{\text{pr},i} \lg \left( 1 + \frac{\overline{m}_{\text{pr}} - m_{\text{pr},i}}{m_{\text{pr},i}} \right) - (V_{\text{pr}} - m_{\text{pr},i}) \lg \left( 1 + \frac{m_{\text{pr},i} - \overline{m}_{\text{pr}}}{V_{\text{pr}} - m_{\text{pr},i}} \right) \\
&\geq -m_{\text{pr},i} \left( \frac{\overline{m}_{\text{pr}} - m_{\text{pr},i}}{m_{\text{pr},i}} - \frac{1}{4} \left( \frac{\overline{m}_{\text{pr}} - m_{\text{pr},i}}{m_{\text{pr},i}} \right)^2 \right) \lg e \\
&\quad - (V_{\text{pr}} - m_{\text{pr},i}) \cdot \frac{m_{\text{pr},i} - \overline{m}_{\text{pr}}}{V_{\text{pr}} - m_{\text{pr},i}} \lg e \\
&= \frac{\lg e}{4} \cdot \frac{(m_{\text{pr},i} - \overline{m}_{\text{pr}})^2}{m_{\text{pr},i}} \\
&\geq \Omega(\kappa^3),
\end{aligned}$$

where the last inequality uses the fact that  $m_{\text{pr},i} \notin [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3]$ .

When  $m_{\text{pr},i} \geq \kappa^{3c}$ , we have

$$\begin{aligned}
(4) - (3) &= -m_{\text{pr},i} \lg \left( 1 + \frac{\overline{m}_{\text{pr}} - m_{\text{pr},i}}{m_{\text{pr},i}} \right) - (V_{\text{pr}} - m_{\text{pr},i}) \lg \left( 1 + \frac{m_{\text{pr},i} - \overline{m}_{\text{pr}}}{V_{\text{pr}} - m_{\text{pr},i}} \right) \\
&\geq -m_{\text{pr},i} \lg \frac{\overline{m}_{\text{pr}}}{m_{\text{pr},i}} - (m_{\text{pr},i} - \overline{m}_{\text{pr}}) \lg e \\
&\geq m_{\text{pr},i} \lg \frac{m_{\text{pr},i}}{e \cdot \overline{m}_{\text{pr}}} \\
&\geq m_{\text{pr},i} \lg \kappa.
\end{aligned}$$

Combining the three cases, we conclude that

$$(4) - (3) \geq \begin{cases} \Omega(\kappa^3) & m_{\text{pr},i} \leq \kappa^{3c}, \\ m_{\text{pr},i} \lg \kappa & m_{\text{pr},i} > \kappa^{3c}, \end{cases}$$

proving the claim. □

On the other hand, since  $V_{\text{pr}} = U \cdot \frac{\overline{m}_{\text{pr}}}{n} \pm O(1)$ , we have

$$(4) = V_{\text{pr}} \lg U - m_{\text{pr},i} \lg n - (V_{\text{pr}} - m_{\text{pr},i}) \lg(U - n) + O(\lg V_{\text{pr}}).$$

Thus, Claim 27 implies that if  $m_{\text{pr},i} \notin [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3]$ , i.e., the primary block is bad, then the data structure size for the primary block is at most

$$V_{\text{pr}} \lg U - m_{\text{pr},i} \lg n - (V_{\text{pr}} - m_{\text{pr},i}) \lg(U - n) - \Omega(\kappa^3)$$

(since  $O(m_{\text{pr},i}) \ll m_{\text{pr},i} \lg \kappa$  and  $O(1) \ll \kappa^3$ ), and otherwise, it is at most

$$V_{\text{pr}} \lg U - m_{\text{pr},i} \lg n - (V_{\text{pr}} - m_{\text{pr},i}) \lg(U - n) + O(\lg V_{\text{pr}}).$$

By applying the same argument to the secondary blocks, we conclude that if  $m_{sc,i} \notin [\kappa^{c+1}, 3\kappa^{c+1}]$ , i.e., the secondary block is bad, then the data structure size for the secondary block is at most

$$V_{sc} \lg U - m_{sc,i} \lg n + (V_{sc} - m_{sc,i}) \lg(U - n) - \Omega(\kappa^3),$$

and otherwise, it is at most

$$V_{sc} \lg U - m_{sc,i} \lg n + (V_{sc} - m_{sc,i}) \lg(U - n) + O(\lg V_{pr}).$$

Summing up the two bounds, and by the fact that at least one of the primary and secondary block is bad, the data structure size for the  $i$ -th bad block pair is at most

$$V_{bl} \lg U - (m_{pr,i} + m_{sc,i}) \lg n - (V_{bl} - m_{pr,i} - m_{sc,i}) \lg(U - n) - \Omega(\kappa^3),$$

since  $\kappa^3 \gg \lg V_{pr}$ .

Now we sum up the size for all  $N_{bad}$  bad blocks, which in total contain  $n_{bad}$  keys, the total size is at most

$$N_{bad} \cdot V_{bl} \lg U - n_{bad} \lg n - (N_{bad} \cdot V_{bl} - n_{bad}) \lg(U - n) - \Omega(N_{bad} \cdot \kappa^3).$$

Therefore, the total size of the data structure is at most

$$\begin{aligned} & \text{OPT}_{(N_{bl}-N_{bad})V_{bl}, n-n_{bad}} + O(N_{bad} \lg N_{bl}) + N_{bad} \cdot V_{bl} \lg U \\ & - n_{bad} \lg n - (N_{bad} \cdot V_{bl} - n_{bad}) \lg(U - n) - \Omega(N_{bad} \cdot \kappa^3) \\ = & (N_{bl} - N_{bad})V_{bl} \lg((N_{bl} - N_{bad})V_{bl}) - (n - n_{bad}) \lg(n - n_{bad}) \\ & - ((N_{bl} - N_{bad})V_{bl} - (n - n_{bad})) \lg((N_{bl} - N_{bad})V_{bl} - (n - n_{bad})) \\ & + N_{bad} \cdot V_{bl} \lg U - n_{bad} \lg n - (N_{bad} \cdot V_{bl} - n_{bad}) \lg(U - n) - \Omega(N_{bad} \cdot \kappa^3), \end{aligned}$$

which by Claim 27 and the fact that  $N_{bad} \geq 1$ , is at most

$$\begin{aligned} & = (N_{bl} - N_{bad})V_{bl} \lg U - (n - n_{bad}) \lg n - ((N_{bl} - N_{bad})V_{bl} - (n - n_{bad})) \lg(U - n) \\ & + N_{bad} \cdot V_{bl} \lg U - n_{bad} \lg n - (N_{bad} \cdot V_{bl} - n_{bad}) \lg(U - n) - \Omega(\kappa^3), \\ = & U \lg U - n \lg n - (U - n) \lg(U - n) - \Omega(\kappa^3) \\ \leq & \lg \binom{U}{n} - \Omega(\kappa^3) \\ = & \text{OPT}_{U,n} - \Omega(\kappa^3), \end{aligned}$$

as we claimed.

**Hash functions.** For all  $x$  in the good blocks, we simply use their hash value according to Lemma 19, for which,  $h$  takes values in  $[n - n_{bad}]$  and  $\bar{h}$  takes values in  $[V_{bl} \cdot (N_{bl} - N_{bad}) - (n - n_{bad})]$ . For  $x$  in the  $i$ -th bad pair with hash value  $v$ , let  $s_i$  be the total number of keys in first  $i - 1$  bad pairs (which is explicitly stored in the data structure), then if  $x \in S$ , we set  $h(x) := n - n_{bad} + s_i + v$ ; if  $x \notin S$ , we set  $\bar{h}(x) := V_{bl} \cdot (N_{bl} - N_{bad} + (i - 1)) - (n - n_{bad} + s_i) + v$ .

That is, all elements in good blocks take the smallest values, and elements in bad blocks take the rest according to the order of the blocks. By definition, they are both bijections.

**Lookup tables.** We include the lookup table from the data structure for no bad blocks, as well as all tables  $\text{tableS}_{V,m}$  from Lemma 25 for  $V = V_{\text{pr}}$  or  $V = V_{\text{sc}}$ , and  $1 \leq m \leq \kappa^{3c}$ . The total lookup table size is  $\tilde{O}(2^{\epsilon\kappa})$ . It is at most  $n^\epsilon$  by readjusting the constant  $\epsilon$ .

**Query algorithm.** Given a query  $x$ , suppose  $x$  is in the  $i$ -th block pair. We first query the hash table to check if it is one of the bad pairs. If the block pair is bad, we follow the pointer and query the data structure for the primary block or the secondary block depending on which block  $x$  is in. Its hash value can be computed according to the definition. It takes constant time.

If the block pair is good, we spend constant time to find out where  $i$ -th block pair is mapped to, in the first  $N_{\text{bl}} - N_{\text{bad}}$  pairs. Then we run  $\text{qa1gG}$  for good blocks, which takes constant time in expectation. This proves Lemma 24.

### 6.3 Final data structure for medium size sets

Consider the following preprocessing algorithm for general  $U$  and  $U^{1/12} \leq n \leq U/2$  ( $U$  not necessarily a multiple of  $V_{\text{bl}}$ ). We first construct a data structure for the block pairs, and fuse the two cases (with or without bad blocks) together.

**preprocessing algorithm**  $\text{perfHash}(U, n, S, \mathcal{R})$ :

1. compute  $V_{\text{pr}}$ ,  $V_{\text{sc}}$  and  $\kappa$
2. compute  $N_{\text{bl}} := U \text{ div } V_{\text{bl}}$  and  $V := U \text{ mod } V_{\text{bl}}$
3. divide the universe  $[U]$  into  $N_{\text{bl}}$  block pairs and a last block of size  $V$
4. if all  $N_{\text{bl}}$  block pairs are good
5.   set  $i := 0$ , apply Lemma 19 on the  $N_{\text{bl}}$  block pairs, and obtain a data structure  $\mathcal{D}_0$
6. else
7.   set  $i := 1$ , apply Lemma 24 on the  $N_{\text{bl}}$  block pairs, and obtain a data structure  $\mathcal{D}_1$
8. apply Proposition 10 to fuse  $i$  into  $\mathcal{D}_i$ , and obtain a data structure  $\mathcal{D}_{\text{bl}}$  (to be cont'd)

Suppose there are  $n_{\text{bl}}$  keys in the first  $N_{\text{bl}}$  block pairs. By Lemma 19,  $\mathcal{D}_0$  has length at most

$$\text{OPT}_{N_{\text{bl}}V_{\text{bl}}, n_{\text{bl}}} + n_{\text{bl}} \cdot 2^{-\kappa/2+2}.$$

By Lemma 24,  $\mathcal{D}_1$  has length at most

$$\text{OPT}_{N_{\text{bl}}V_{\text{bl}}, n_{\text{bl}}} - \Omega(\kappa^3).$$

Thus, by Proposition 10,  $\mathcal{D}_{\text{bl}}$  has length at most

$$\begin{aligned} & \text{OPT}_{N_{\text{bl}}V_{\text{bl}}, n_{\text{bl}}} + \lg(2^{n_{\text{bl}} \cdot 2^{-\kappa/2+2}} + 2^{-\Omega(\kappa^3)}) + 2^{-\kappa+2} \\ & \leq \text{OPT}_{N_{\text{bl}}V_{\text{bl}}, n_{\text{bl}}} + \lg(1 + n_{\text{bl}} \cdot 2^{-\kappa/2+2} + 2^{-\Omega(\kappa^3)}) + 2^{-\kappa+2} \\ & \leq \text{OPT}_{N_{\text{bl}}V_{\text{bl}}, n_{\text{bl}}} + n_{\text{bl}} \cdot 2^{-\kappa/2+3} \\ & = \text{OPT}_{U-V, n_{\text{bl}}} + n_{\text{bl}} \cdot 2^{-\kappa/2+3}. \end{aligned}$$

Then we construct a separate data structure for the last block using Lemma 25 or Lemma 26 based on the number of keys in it.

9. if  $n - n_{\text{bl}} \leq \kappa^{3c}$
10. construct  $\mathcal{D}_{\text{last}}$  for the last block using Lemma 25
11. apply Proposition 8 to concatenate  $\mathcal{D}_{\text{bl}}$  and  $\mathcal{D}_{\text{last}}$ , and obtain  $\mathcal{D}'$
12. let  $n'_{\text{bl}} = n_{\text{bl}}$
13. else
14. construct  $\mathcal{D}_{\text{last}}$  for the last block using Lemma 26
15. spend  $\lceil \lg n \rceil$  bits to store  $n_{\text{bl}}$ ,
16. round both  $\mathcal{D}_{\text{bl}}$  and  $\mathcal{D}_{\text{last}}$  to integral lengths and concatenate them
17. spend  $\lceil \lg n \rceil$  bits to store a point to  $\mathcal{D}_{\text{last}}$ , let the resulting data structure be  $\mathcal{D}'$
18. let  $n'_{\text{bl}} = n - \kappa^{3c} - 1$
19. apply Proposition 10 to fuse the value of  $n'_{\text{bl}}$  into  $\mathcal{D}'$  for  $n'_{\text{bl}} \in [n - \kappa^{3c} - 1, n]$ , and obtain  $\mathcal{D}$
20. return  $\mathcal{D}$

We do not fuse the whole value of  $n_{\text{bl}}$  into  $\mathcal{D}'$ , as its range is large and Proposition 10 requires a large lookup table to do this. However, we only fuse its value if  $n_{\text{bl}} \geq n - \kappa^{3c}$ , and otherwise only indicate that it is smaller than  $n - \kappa^{3c}$  (by setting  $n'_{\text{bl}}$  to  $n - \kappa^{3c} - 1$ ). This is fine because we spend extra  $\lg n$  bits to explicitly store its value in this case. We will show in the following that the total space is close the optimum.

There are  $n - n_{\text{bl}}$  keys in the last block. If  $n - n_{\text{bl}} \leq \kappa^{3c}$ ,  $\mathcal{D}_{\text{last}}$  has size at most

$$\text{OPT}_{V, n - n_{\text{bl}}} + (n - n_{\text{bl}} - 1) \cdot 2^{-\kappa/2+1}.$$

In this case, the length of  $\mathcal{D}'$  is at most

$$\begin{aligned} & \text{OPT}_{U-V, n_{\text{bl}}} + \text{OPT}_{V, n - n_{\text{bl}}} + n_{\text{bl}} \cdot 2^{-\kappa/2+3} + (n - n_{\text{bl}} - 1) \cdot 2^{-\kappa/2+1} + 2^{-\kappa+2} \\ & \leq \text{OPT}_{U-V, n_{\text{bl}}} + \text{OPT}_{V, n - n_{\text{bl}}} + n \cdot 2^{-\kappa/2+3} \\ & = \lg \binom{U-V}{n_{\text{bl}}} \binom{V}{n - n_{\text{bl}}} + n \cdot 2^{-\kappa/2+3}. \end{aligned} \quad (5)$$

If  $n - n_{\text{bl}} > \kappa^{3c}$ ,  $\mathcal{D}_{\text{last}}$  has size at most

$$\text{OPT}_{V, n - n_{\text{bl}}} + O(n - n_{\text{bl}} + \lg n).$$

The length of  $\mathcal{D}'$  is at most

$$\lg \binom{U-V}{n_{\text{bl}}} \binom{V}{n - n_{\text{bl}}} + O(n - n_{\text{bl}} + \lg n). \quad (6)$$

By Sterling's formula, the first term is at most

$$\begin{aligned} \lg \binom{U-V}{n_{\text{bl}}} \binom{V}{n - n_{\text{bl}}} &= \lg \frac{(U-V)!V!}{n_{\text{bl}}!(U-V-n_{\text{bl}})!(n-n_{\text{bl}})!(V-n+n_{\text{bl}})!} \\ &\leq (U-V) \lg(U-V) - n_{\text{bl}} \lg n_{\text{bl}} - (U-V-n_{\text{bl}}) \lg(U-V-n_{\text{bl}}) \\ &\quad + V \lg V - (n-n_{\text{bl}}) \lg(n-n_{\text{bl}}) - (V-n+n_{\text{bl}}) \lg(V-n+n_{\text{bl}}) + O(\lg U) \end{aligned}$$

which again by the fact that  $f(x) = m \lg x + (V-m) \lg(V-x)$  is maximized at  $x = m$ , is at most

$$\begin{aligned} & \leq (U-V) \lg(U-V) - n_{\text{bl}} \lg \frac{(U-V)n}{U} - (U-V-n_{\text{bl}}) \lg(U-V - \frac{(U-V)n}{U}) \\ & \quad + V \lg V - (n-n_{\text{bl}}) \lg(n-n_{\text{bl}}) - (V-n+n_{\text{bl}}) \lg(V-n+n_{\text{bl}}) + O(\lg U) \end{aligned}$$

$$\begin{aligned}
&= (U - V) \lg U - n_{\text{bl}} \lg n - (U - V - n_{\text{bl}}) \lg(U - n) \\
&\quad + V \lg V - (n - n_{\text{bl}}) \lg \frac{Vn}{U} - (V - n + n_{\text{bl}}) \left( V - \frac{Vn}{U} \right) \\
&\quad - (n - n_{\text{bl}}) \lg \frac{U(n - n_{\text{bl}})}{Vn} - (V - n + n_{\text{bl}}) \lg \frac{V - n + n_{\text{bl}}}{V - \frac{Vn}{U}} + O(\lg U) \\
&= U \lg U - n \lg n - (U - n) \lg(U - n) \\
&\quad - (n - n_{\text{bl}}) \lg \frac{U(n - n_{\text{bl}})}{Vn} + (V - n + n_{\text{bl}}) \lg \left( 1 + \frac{n - n_{\text{bl}} - \frac{Vn}{U}}{V - n + n_{\text{bl}}} \right) + O(\lg U)
\end{aligned}$$

which by the fact that  $\ln(1 + x) \leq x$ , is at most

$$\leq \lg \binom{U}{n} - (n - n_{\text{bl}}) \lg \frac{U(n - n_{\text{bl}})}{Vn} + (n - n_{\text{bl}}) \lg e + O(\lg U)$$

which by the fact that  $n - n_{\text{bl}} \geq \kappa^{3c}$ , is at most

$$\leq \lg \binom{U}{n} - (n - n_{\text{bl}}) \lg \frac{U \kappa^{3c}}{eVn} + O(\lg U)$$

which by the fact that  $U/(eVn) \geq \kappa^{-2c}$  and  $c \geq 1$ , is at most

$$\leq \lg \binom{U}{n} - (n - n_{\text{bl}}) \lg \kappa + O(\lg U).$$

Therefore, when  $n - n_{\text{bl}} > \kappa^{3c}$ , the size of  $\mathcal{D}'$  is at most

$$\begin{aligned}
(6) &\leq \lg \binom{U}{n} - \Omega((n - n_{\text{bl}}) \lg \kappa) + O(\lg U + \lg n) \\
&\leq \text{OPT}_{U,n} - \Omega(\kappa^3).
\end{aligned}$$

Finally together with Equation (5), by Proposition 10, the size of  $\mathcal{D}$  is at most

$$\begin{aligned}
&\lg \left( 2^{\text{OPT}_{U,n} - \Omega(\kappa^3)} + \sum_{n_{\text{bl}}=n-\kappa^{3c}}^n \binom{U-V}{n_{\text{bl}}} \binom{V}{n-n_{\text{bl}}} \cdot 2^{n \cdot 2^{-\kappa/2+3}} \right) + n \cdot 2^{-\kappa+2} \\
&\leq \lg \left( 2^{\text{OPT}_{U,n} - \Omega(\kappa^3)} + \sum_{n_{\text{bl}}=0}^n \binom{U-V}{n_{\text{bl}}} \binom{V}{n-n_{\text{bl}}} \cdot 2^{n \cdot 2^{-\kappa/2+3}} \right) + n \cdot 2^{-\kappa+2} \\
&= \lg \left( 2^{\text{OPT}_{U,n} - \Omega(\kappa^3)} + \binom{U}{n} \cdot 2^{n \cdot 2^{-\kappa/2+3}} \right) + n \cdot 2^{-\kappa+2} \\
&= \text{OPT}_{U,n} + \lg \left( 2^{-\Omega(\kappa^3)} + 2^{n \cdot 2^{-\kappa/2+3}} \right) + n \cdot 2^{-\kappa+2} \\
&\leq \text{OPT}_{U,n} + \lg \left( 2^{-\Omega(\kappa^3)} + 1 + n \cdot 2^{-\kappa/2+3} \right) + n \cdot 2^{-\kappa+2} \\
&\leq \text{OPT}_{U,n} + 2^{-\Omega(\kappa^3)} + n \cdot 2^{-\kappa/2+4} + n \cdot 2^{-\kappa+2} \\
&\leq \text{OPT}_{U,n} + U^{-1}.
\end{aligned}$$

**Hash functions.** For all  $x$  in the first  $N_{\text{bl}}$  block pairs, we simply use their hash values defined by  $\mathcal{D}_{\text{bl}}$  (from Lemma 19 or Lemma 24), for which,  $h$  takes values from  $[n_{\text{bl}}]$  and  $\bar{h}$  takes values from  $[V_{\text{bl}} \cdot N_{\text{bl}} - n_{\text{bl}}]$ . For all  $x$  in the last block, let  $v$  be its hash value defined by  $\mathcal{D}_{\text{last}}$ . If  $x \in S$ , let  $h(x) := n_{\text{bl}} + v$ ; if  $x \notin S$ , let  $\bar{h}(x) := V_{\text{bl}} \cdot N_{\text{bl}} - n_{\text{bl}} + v$ . By definition,  $h$  and  $\bar{h}$  are both bijections.

**Lookup table.** We include the lookup tables in Section 6.1 and in Section 6.2, which both have size  $n^\epsilon$ . Then we include the lookup tables needed by Proposition 8 and Proposition 10 in line 8, 11 and 19. The total size is  $n^\epsilon$ .

**Query algorithm.** Given a query  $x$ , we decode all the components, and query the part based on the value of  $x$ .

**query algorithm**  $\text{qAlg}(x)$ :

1. compute  $V_{\text{pr}}, V_{\text{sc}}, N_{\text{bl}}$  and  $V$
2. apply Proposition 10 to recover  $n'_{\text{bl}}$  and decode  $\mathcal{D}'$  from  $\mathcal{D}$
3. if  $n'_{\text{bl}} > n - \kappa^{3c} - 1$
4.   let  $n_{\text{bl}} := n'_{\text{bl}}$
5.   apply Proposition 8 to decode  $\mathcal{D}_{\text{bl}}$  and  $\mathcal{D}_{\text{last}}$
6. else
7.   recover  $n_{\text{bl}}$  from  $\mathcal{D}'$
8.   decode  $\mathcal{D}_{\text{bl}}$  and  $\mathcal{D}_{\text{last}}$
9. if  $x \geq U - V$
10.   query  $x$  in  $\mathcal{D}_{\text{last}}$ , and obtain  $(b, v)$
11.   if  $b = 1$ , then return  $(1, n_{\text{bl}} + v)$
12.   if  $b = 0$ , then return  $(0, V_{\text{bl}} \cdot N_{\text{bl}} - n_{\text{bl}} + v)$
13. else
14.   apply Proposition 10 to recover  $i$  and decode  $\mathcal{D}_i$  from  $\mathcal{D}_{\text{bl}}$
15.   query  $x$  in  $\mathcal{D}_i$  using the corresponding query algorithm, and return the outcome

All subroutines run in constant time, the overall query time is constant. This proves Theorem 18.

## 7 Data Structure Pair for Block Pair

In this section, we prove our main technical lemma (Lemma 17), which constructs a pair of data structures for a pair of blocks.

**Lemma 17 (restated).** *Let  $\kappa$  be the fineness parameter for fractional-length strings, and  $c$  be a constant positive integer. Let  $V \in [2\kappa^{2c-3}, 2^{\kappa/2}]$  and  $V_{\text{sc}} \geq 4\kappa^{c+1}$ . For any constant  $\epsilon > 0$ , there is a preprocessing algorithm  $\text{perfHashBlk}$ , query algorithms  $\text{qalgBlk}_{\text{main}}$ ,  $\text{qalgBlk}$  and lookup tables  $\text{tableBlk}_{V, V_{\text{sc}}}$  of size  $\tilde{O}(2^{\epsilon\kappa})$ . Given*

- *a set  $S \subseteq [V]$  such that  $m := |S| \in [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3]$ ,*
- *a set  $S_{\text{sc}} \subseteq \{V, \dots, V + V_{\text{sc}} - 1\}$  and  $m_{\text{sc}} := |S_{\text{sc}}| \in [\kappa^{c+1}, 3\kappa^{c+1}]$ ,*
- *a random string  $\mathcal{R}$  of  $\kappa^{c+1}$  bits,*

*$\text{perfHashBlk}$  preprocesses  $S$  and  $S_{\text{sc}}$  into a pair of two (fractional-length) data structures  $\mathcal{D}_{\text{main}}$  and  $\mathcal{D}_{\text{aux}}$ , such that*

- (i)  $\mathcal{D}_{\text{main}}$  has length at most

$$\lg \binom{V}{\kappa^{2c-3}, \kappa^c} + \kappa^{2c-3} \cdot 2^{-\kappa/2+1};$$

(ii)  $\mathcal{D}_{\text{aux}}$  has length at most

$$\lg \binom{V}{m} + \lg \binom{V_{\text{sc}}}{m_{\text{sc}}} - \lg \binom{V}{\kappa^{2c-3}, \kappa^c} + \kappa^{c+1} 2^{-\kappa/2+2};$$

(iii)  $\mathcal{D}_{\text{main}}$  and  $\mathcal{D}_{\text{aux}}$  together define a bijection  $h$  between

$$S \cup S_{\text{sc}} \quad \text{and} \quad [m + m_{\text{sc}}],$$

and a bijection  $\bar{h}$  between

$$[V + V_{\text{sc}}] \setminus (S \cup S_{\text{sc}}) \quad \text{and} \quad [(V + V_{\text{sc}}) - (m + m_{\text{sc}})],$$

such that  $h(S) \supset [\kappa^{2c-3}]$  and  $\bar{h}([V] \setminus S) \supset [V - \kappa^{2c-3} - \kappa^c]$ ;

- (iv) given any  $x \in [V]$ ,  $\text{qalgBlk}_{\text{main}}(V, x)$  outputs  $\text{hash}(x)$  when  $x \in S$  and  $h(x) \in [\kappa^{2c-3}]$ , or when  $x \notin S$  and  $\bar{h}(x) \in [V - \kappa^{2c-3} - \kappa^c]$ , otherwise it outputs “unknown”; moreover, it only accesses  $\mathcal{D}_{\text{main}}$ ,  $\mathcal{R}$  and the lookup table  $\text{tableBlk}_{V, V_{\text{sc}}}$ , and it runs in constant time in the worst case;
- (v) for any  $x \in [V]$ , the probability that  $\text{qalgBlk}_{\text{main}}(V, x)$  outputs “unknown” is at most  $O(\kappa^{-c+3})$  over the randomness of  $\mathcal{R}$ ;
- (vi) given any  $x \in [V + V_{\text{sc}}]$ ,  $\text{qalgBlk}(V, m, V_{\text{sc}}, m_{\text{sc}}, x)$  computes  $\text{hash}(x)$ ; it accesses  $\mathcal{D}_{\text{main}}$ ,  $\mathcal{D}_{\text{aux}}$ ,  $\mathcal{R}$  and the lookup table  $\text{tableBlk}_{V, V_{\text{sc}}}$ , and it runs in  $O(\kappa^4)$  time.

As mentioned in the overview, we first improve the rank data structure of Pătraşcu [Păt08]. We show that if a block has  $\kappa^c$  keys, then there is a rank data structure with constant query time and negligible extra space. The rank problem asks to preprocess a set  $S$  of  $m$  keys into a data structure, supporting

- $\text{rank}_S(x)$ : return the number of keys that are at most  $x$ .

In particular, by computing both  $\text{rank}_S(x)$  and  $\text{rank}_S(x-1)$ , one can decide if  $x \in S$ .

**Lemma 28.** *Let  $c$  be any constant positive integer and  $\epsilon$  be any positive constant. There is a preprocessing algorithm  $\text{prepRank}$ , query algorithm  $\text{qAlgRank}$  and lookup tables  $\text{tableRank}_{V, m}$  of sizes  $\tilde{O}(2^{\epsilon\kappa})$ , such that for any integers  $V \leq 2^{\kappa/2}$ ,  $m \leq \kappa^c$ , given a set  $S \subset [V]$  of size  $m$ ,  $\text{prepRank}(V, m, S)$  outputs a data structure  $\mathcal{D}$  of length*

$$\lg \binom{V}{m} + (m-1) \cdot 2^{-\kappa/2}.$$

Given  $x \in [V]$ ,  $\text{qAlgRank}(V, m, x)$  computes  $\text{rank}_S(x)$  in constant time, by accessing  $\mathcal{D}$  and  $\text{tableRank}_{V, m}$ . In particular, by computing both  $\text{rank}_S(x)$  and  $\text{rank}_S(x-1)$ , one can decide if  $x \in S$  in constant time. The algorithms run on a random access machine with word-size  $w = \Theta(\kappa)$ .

We prove this lemma in Section 7.2. Note that a rank data structure easily defines hash functions that map the keys to  $[m]$ , and the non-keys to  $[V - m]$ : For each key  $x$ , we set  $h(x) := \text{rank}(x) - 1$ ; for each non-key  $x$ , we set  $\bar{h}(x) := x - \text{rank}(x) - 1$ . Lemma 17 designs a pair of two data structure, where the size of the main data structure *does not* depend on  $m$  (it only assumes that  $m$  is in a small range). Most queries can be answered by accessing only the main data structure, also *without* knowing the value of  $m$ .

To construct the two data structures, we first apply Lemma 28 to construct a rank data structure for the secondary block  $S_{\text{sc}} \subset [V_{\text{sc}}]$ . Denote this data structure by  $\mathcal{D}_{\text{sc}}$ . Then we pick a set of  $\kappa^{2c-3}$  keys from  $S$ , as well as  $V - \kappa^{2c-3} - \kappa^c$  non-keys from  $[V] \setminus S$ , and construct the above rank data structure, which will be the main data structure. The remaining  $\kappa^c$  elements in  $[V]$  will correspond to the “unknowns.” We pick the two sets based on the bits in  $\mathcal{D}_{\text{sc}}$ . That is, we apply Proposition 14 and divide  $\mathcal{D}_{\text{sc}}$  into a string of length  $\approx \lg \binom{m}{\kappa^{2c-3}}$ , a string of length  $\approx \lg \binom{V-m}{V-\kappa^{2c-3}-\kappa^c}$  and the remaining bits. Then we apply the following lemma to *interpret* the first string as a subset of  $S$  and the second string as a subset of  $[V] \setminus S$ . The final auxiliary data structure consists of the remaining bits of  $\mathcal{D}_{\text{sc}}$ , as well as a data structure for the “unknowns.”

**Lemma 29.** Let  $c \geq 2$  be a constant positive integer. There is a preprocessing algorithm `prepIntoSet`, a query algorithm `request` and lookup tables `tableIntV,m` of sizes  $O(\kappa^{c+2})$ , such that for any integers  $V$  and  $m$  where  $V \leq 2^{\kappa/2}$  and  $m \leq \kappa^c$ , given a (double-ended) string  $\mathcal{D} = (K_h, M, K_t)$  of length

$$s \leq \lg \binom{V}{m} - m(V-1)2^{-\kappa+2},$$

`prepIntoSet`( $V, m, \mathcal{D}$ ) outputs a set  $S \subseteq [V]$  of size  $m$ . For any  $-1 \leq a_1 \leq a_2 \leq |M|$  and  $a_2 < a_1 + \kappa$ , `request`(0,  $V, m, \text{range}(K_h), |M|, \text{range}(K_t), a_1, a_2$ ) computes  $\mathcal{D}[a_1, a_2]$  in  $O(\kappa^4)$  time using  $O(\kappa^2)$  rank queries to  $S$ , assuming it can make random accesses to the lookup table `tableIntV,m`.

The lemma guarantees that even if the original string is not stored explicitly as part of the final data structure, it can still be accessed implicitly assuming rank queries to the sets, which the above dictionary data structures support naturally. We prove the above lemma in Section 7.1.

In the remainder of the section, we show how to construct the pair of data structures in detail.

*Proof of Lemma 17.* We begin by presenting the preprocessing algorithm.

**Preprocessing algorithm.** In the preprocessing algorithm, we first construct a rank data structure for the secondary block, and divide it into three substrings.

**Preprocessing algorithm** `perfHashBlk`( $V, m, V_{sc}, m_{sc}, S, S_{sc}, \mathcal{R}$ ):

1. run  $\mathcal{D}_{sc} := \text{prepRank}(V_{sc}, m_{sc}, S_{sc})$  (from Lemma 28) to construct a rank data structure for  $S_{sc}$  using
$$s_{sc} \leq \lg \binom{V_{sc}}{m_{sc}} + (m_{sc} - 1) \cdot 2^{-\kappa/2}$$

bits
2. apply Proposition 14 twice to divide  $\mathcal{D}_{sc}$  into
  - $\mathcal{D}_{sc,1}$ : length  $\leq \lg \binom{m}{m - \kappa^{2c-3}} - \kappa^c(m-1)2^{-\kappa+2}$
  - $\mathcal{D}_{sc,2}$ : length  $\leq \lg \binom{V-m}{\kappa^{2c-3} + \kappa^c - m} - \kappa^c(V-m-1)2^{-\kappa+2}$
  - $\mathcal{D}_{sc,3}$ : length  $\leq s_{sc} - \lg \binom{m}{m - \kappa^{2c-3}} - \lg \binom{V-m}{\kappa^{2c-3} + \kappa^c - m} + \kappa^c V 2^{-\kappa+2}$  (to be cont'd)

Note that  $m - \kappa^{2c-3} \geq \kappa^c/3$  and  $\kappa^{2c-3} + \kappa^c - m \geq \kappa^c/3$ , therefore, both  $\mathcal{D}_{sc,1}$  and  $\mathcal{D}_{sc,2}$  have length at least  $4\kappa$ . For  $\mathcal{D}_{sc,3}$ , we have  $s_{sc} \geq \lg \binom{V_{sc}}{m_{sc}} \geq \kappa^{c+1}$ , and thus,

$$\begin{aligned} & s_{sc} - \lg \binom{m}{m - \kappa^{2c-3}} - \lg \binom{V-m}{\kappa^{2c-3} + \kappa^c - m} + \kappa^c V 2^{-\kappa+2} \\ & \geq \kappa^{c+1} - (m - \kappa^{2c-3}) \lg V - (\kappa^{2c-3} + \kappa^c - m) \lg V \\ & \geq \kappa^{c+1} - \kappa^c \lg V \\ & \geq 4\kappa. \end{aligned}$$

The premises of Proposition 14 are satisfied. In order to store two *random* subsets in the main data structure, we “XOR”  $\mathcal{D}_{sc,1}$  and  $\mathcal{D}_{sc,2}$  with the random string  $\mathcal{R}$ .

3. compute  $\mathcal{D}_{sc,1} \oplus \mathcal{R}$  and  $\mathcal{D}_{sc,2} \oplus \mathcal{R}$  (to be cont'd)

For double-ended string  $\mathcal{D} = (K_h, M, K_t)$ ,  $\mathcal{D} \oplus \mathcal{R}$  is defined as follows: compute the bitwise XOR of  $M$  and  $\mathcal{R}[1, |M|]$ , treat  $\mathcal{R}[|M|+1, |M|+2\kappa]$  and  $\mathcal{R}[|M|+2\kappa+1, |M|+4\kappa]$  as two  $2\kappa$ -bit integers, and compute  $(K_h + \mathcal{R}[|M|+1, |M|+2\kappa]) \bmod \text{range}(K_h)$  and  $(K_t + \mathcal{R}[|M|+2\kappa+1, |M|+4\kappa]) \bmod \text{range}(K_t)$ ;

$\mathcal{D} \oplus \mathcal{R}$  is the double-ended string (with the same length as  $\mathcal{D}$ ), formed by the outcomes. In particular, since  $\text{range}(K_h)$  and  $\text{range}(K_t)$  are both smaller than  $2^{\kappa+1} \ll 2^{2\kappa}$ , when  $\mathcal{R}$  is uniformly random,  $\mathcal{D} \oplus \mathcal{R}$  is very close to uniform. We have the following claim by standard information theory.

**Claim 30.** *For any fixed  $\mathcal{D} = (K_h, M, K_t)$  and uniformly random  $\mathcal{R}$ , we have*

$$\begin{aligned} H(\mathcal{D} \oplus \mathcal{R}) &\geq (\lg(\text{range}(K_h)) + |M| + \lg(\text{range}(K_t)))(1 - 2^{-\kappa+2}) \\ &\geq (|\mathcal{D}| - 2^{-\kappa+2})(1 - 2^{-\kappa+2}). \end{aligned}$$

Also,  $K_h, K_t$  and any  $O(\kappa)$  consecutive bits of  $M$  can be computed in constant time, given random access to  $\mathcal{D} \oplus \mathcal{R}$  and  $\mathcal{R}$ . Next, we interpret the  $\mathcal{D}_{\text{sc},1} \oplus \mathcal{R}$  and  $\mathcal{D}_{\text{sc},2} \oplus \mathcal{R}$  as two subsets using Lemma 29.

4. run  $S_1 := \text{prepIntoSet}(m, m - \kappa^{2c-3}, \mathcal{D}_{\text{sc},1} \oplus \mathcal{R})$  (from Lemma 29) to interpret  $\mathcal{D}_{\text{sc},1} \oplus \mathcal{R}$  as a set  $S_1 \subseteq [m]$  of size  $m - \kappa^{2c-3}$   
run  $S_2 := \text{prepIntoSet}(V - m, \kappa^{2c-3} + \kappa^c - m, \mathcal{D}_{\text{sc},2} \oplus \mathcal{R})$  to interpret  $\mathcal{D}_{\text{sc},2} \oplus \mathcal{R}$  as a set  $S_2 \subseteq [V - m]$  of size  $\kappa^{2c-3} + \kappa^c - m$
5. compute  $S_{\text{unk}} \subseteq S$  according to  $S_1$   
compute  $\overline{S}_{\text{unk}} \subseteq [V] \setminus S$  according to  $S_2$  (to be cont'd)

More specifically, for each  $i \in [m]$ ,  $S_{\text{unk}}$  contains the  $(i+1)$ -th smallest element in  $S$  if and only if  $i \in S_1$ . Similarly,  $\overline{S}_{\text{unk}}$  contains the  $(i+1)$ -th smallest element in  $[V] \setminus S$  if and only if  $i \in S_2$ . They are the keys and non-keys that are *not* to be stored in the main data structure, i.e., the “unknowns.”

Then, we compute the main data structure  $\mathcal{D}_{\text{main}}$ .

6. apply Proposition 8 to concatenate the following two data structures: and obtain  $\mathcal{D}_{\text{main}}$ :
  - $\mathcal{D}_{\text{main},1} := \text{prepRank}(V, \kappa^{2c-3}, S \setminus S_{\text{unk}})$  (from Lemma 28), a rank data structure
  - $\mathcal{D}_{\text{main},2} := \text{prepRank}(V - \kappa^{2c-3}, \kappa^c, “S_{\text{unk}} \cup \overline{S}_{\text{unk}}”) a rank data structure for  $S_{\text{unk}} \cup \overline{S}_{\text{unk}}$  over  $[V] \setminus (S \setminus S_{\text{unk}})$  (see below) (to be cont'd)$

For  $\mathcal{D}_{\text{main},2}$ , before running  $\text{prepRank}$ , we first remove all  $\kappa^{2c-3}$  elements in  $S \setminus S_{\text{unk}}$  from both the universe  $[V]$  and  $S_{\text{unk}} \cup \overline{S}_{\text{unk}}$ , and keep the order of the remaining elements. Thus, the new universe becomes  $[V - \kappa^{2c-3}]$ . In the other words,  $\mathcal{D}_{\text{main},2}$  supports queries of form “return # of elements in  $S_{\text{unk}} \cup \overline{S}_{\text{unk}}$  that are no larger than  $i$ -th smallest element in  $[V] \setminus (S \setminus S_{\text{unk}})$ ”.

Finally, we compute the auxiliary data structure  $\mathcal{D}_{\text{aux}}$ .

7. apply Proposition 8 to concatenate the following two data structures and obtain  $\mathcal{D}_{\text{aux}}$ :
  - $\mathcal{D}_{\text{aux},1} := \text{prepRank}(\kappa^c, m - \kappa^{2c-3}, “S_{\text{unk}}”) a rank data structure for  $S_{\text{unk}}$  over  $S_{\text{unk}} \cup \overline{S}_{\text{unk}}$$
  - $\mathcal{D}_{\text{aux},2} := \mathcal{D}_{\text{sc},3}$

Similarly,  $\mathcal{D}_{\text{aux},1}$  supports queries of form “return # of elements in  $S_{\text{unk}}$  that are no larger than the  $i$ -th smallest element in  $S_{\text{unk}} \cup \overline{S}_{\text{unk}}$ .”

**Space analysis.** Next, we analyze the length of  $\mathcal{D}_{\text{main}}$  and  $\mathcal{D}_{\text{aux}}$ .  $\mathcal{D}_{\text{main}}$  is the concatenation of  $\mathcal{D}_{\text{main},1}$  and  $\mathcal{D}_{\text{main},2}$ . For  $\mathcal{D}_{\text{main},1}$ , its length is at most

$$\lg \binom{V}{\kappa^{2c-3}} + (\kappa^{2c-3} - 1)2^{-\kappa/2}$$

by Lemma 28. For  $\mathcal{D}_{\text{main},2}$ , its length is at most

$$\lg \binom{V - \kappa^{2c-3}}{\kappa^c} + (\kappa^c - 1)2^{-\kappa/2}.$$

By Proposition 8, the length of  $\mathcal{D}_{\text{main}}$  is at most

$$\begin{aligned} & \lg \binom{V}{\kappa^{2c-3}} + (\kappa^{2c-3} - 1)2^{-\kappa/2} + \lg \binom{V - \kappa^{2c-3}}{\kappa^c} + (\kappa^c - 1)2^{-\kappa/2} + 2^{-\kappa+4} \\ & \leq \lg \binom{V}{\kappa^{2c-3}, \kappa^c} + \kappa^{2c-3}2^{-\kappa/2+1}. \end{aligned}$$

$\mathcal{D}_{\text{aux}}$  is the concatenation of  $\mathcal{D}_{\text{aux},1}$  and  $\mathcal{D}_{\text{aux},2}$ . For  $\mathcal{D}_{\text{aux},1}$ , its length is at most

$$\lg \binom{\kappa^c}{m - \kappa^{2c-3}} + (m - \kappa^{2c-3} - 1)2^{-\kappa/2}.$$

For  $\mathcal{D}_{\text{aux},2}$ , which is  $\mathcal{D}_{\text{sc},3}$ , its length is at most

$$\begin{aligned} & s_{\text{sc}} - \lg \binom{m}{m - \kappa^{2c-3}} - \lg \binom{V - m}{\kappa^{2c-3} + \kappa^c - m} + \kappa^c V 2^{-\kappa+2} \\ & \leq \lg \binom{V_{\text{sc}}}{m_{\text{sc}}} - \lg \binom{m}{m - \kappa^{2c-3}} - \lg \binom{V - m}{\kappa^{2c-3} + \kappa^c - m} + 3\kappa^{c+1}2^{-\kappa/2}, \end{aligned}$$

since  $V \leq 2^{\kappa/2}$  and  $m_{\text{sc}} \leq 3\kappa^{c+1}$ . Summing up the lengths and by Proposition 8, the length of  $\mathcal{D}_{\text{aux}}$  is at most

$$\begin{aligned} & \lg \binom{\kappa^c}{m - \kappa^{2c-3}} + (m - \kappa^{2c-3} - 1)2^{-\kappa/2} + \lg \binom{V_{\text{sc}}}{m_{\text{sc}}} \\ & \quad - \lg \binom{m}{m - \kappa^{2c-3}} - \lg \binom{V - m}{\kappa^{2c-3} + \kappa^c - m} + 3\kappa^{c+1}2^{-\kappa/2} + 2^{-\kappa+4} \\ & \leq \lg \frac{\kappa^c! \kappa^{2c-3}! (V - \kappa^{2c-3} - \kappa^c)!}{m! (V - m)!} + \lg \binom{V_{\text{sc}}}{m_{\text{sc}}} + (m - \kappa^{2c-3} + 3\kappa^{c+1})2^{-\kappa/2} \\ & \leq \lg \binom{V}{m} + \lg \binom{V_{\text{sc}}}{m_{\text{sc}}} - \lg \binom{V}{\kappa^{2c-3}, \kappa^c} + \kappa^{c+1}2^{-\kappa/2+2}, \end{aligned}$$

as we claimed. This proves item (i) and (ii) in the statement.

**Hash functions.** For  $x \in S \cup S_{\text{sc}}$ , we define  $h(x)$  as follows.

- For  $x \in S \setminus S_{\text{unk}}$ , let  $h(x) := \text{rank}_{S \setminus S_{\text{unk}}}(x) - 1$ ; they are mapped to  $[\kappa^{2c-3}]$ .
- For  $x \in S_{\text{unk}}$ , let  $h(x) := \kappa^{2c-3} + \text{rank}_{S_{\text{unk}}}(x) - 1$ ; they are mapped to  $\{\kappa^{2c-3}, \dots, m - 1\}$ .
- For  $x \in S_{\text{sc}}$ , let  $h(x) := m + \text{rank}_{S_{\text{sc}}}(x) - 1$ ; they are mapped to  $\{m, \dots, m + m_{\text{sc}} - 1\}$ .

Similarly, for  $x \notin S \cup S_{\text{sc}}$ , we define  $\bar{h}$  as follows.

- For  $x \in ([V] \setminus S) \setminus \bar{S}_{\text{unk}}$ , let  $\bar{h}(x) := \text{rank}_{[V] \setminus S \setminus \bar{S}_{\text{unk}}}(x) - 1$ ; they are mapped to  $[V - \kappa^{2c-3} - \kappa^c]$ .
- For  $x \in \bar{S}_{\text{unk}}$ , let  $\bar{h}(x) := V - \kappa^{2c-3} - \kappa^c + \text{rank}_{\bar{S}_{\text{unk}}}(x) - 1$ ; they are mapped to  $\{V - \kappa^{2c-3} - \kappa^c, \dots, V - m - 1\}$ .
- For  $x \in \{V, \dots, V + V_{\text{sc}} - 1\} \setminus S_{\text{sc}}$ , let  $\bar{h}(x) := V - m + \text{rank}_{\{V, \dots, V + V_{\text{sc}} - 1\} \setminus S_{\text{sc}}}(x) - 1$ ; they are mapped to  $\{V - m, \dots, V + V_{\text{sc}} - m - m_{\text{sc}} - 1\}$ .

Overall,  $h$  is a bijection between  $S \cup S_{\text{sc}}$  and  $[m + m_{\text{sc}}]$ , and  $\bar{h}$  is a bijection between  $[V + V_{\text{sc}}] \setminus (S \cup S_{\text{sc}})$  and  $[V + V_{\text{sc}} - m - m_{\text{sc}}]$ . Moreover,  $h(S) \supset [\kappa^{2c-3}]$  and  $\bar{h}([V] \setminus S) \supset [V - \kappa^{2c-3} - \kappa^c]$ . This proves item (iii) in the statement.

**Lookup table.** We store the following information in the lookup table.

<b>lookup table</b> $\text{tableBlk}_{V, V_{sc}}$ : 1. lookup table for line 6 from Proposition 8 2. $\text{tableRank}_{V, \kappa^{2c-3}}$ , $\text{tableRank}_{V - \kappa^{2c-3}, \kappa^c}$ from Lemma 28 3. for all $m \in [\kappa^{2c-3} + \kappa^c/3, \kappa^{2c-3} + 2\kappa^c/3]$ and $m_{sc} \in [\kappa^{c+1}, 2\kappa^{c+1}]$ <ul style="list-style-type: none"> <li>lookup tables for line 7 from Proposition 8</li> <li><math>\text{tableRank}_{\kappa^c, m - \kappa^{2c-3}}</math> and <math>\text{tableRank}_{V_{sc}, m_{sc}}</math> from Lemma 28</li> <li><math>\text{tableInt}_{m, m - \kappa^{2c-3}}</math> and <math>\text{tableInt}_{V - m, \kappa^{2c-3} + \kappa^c - m}</math> from Lemma 29</li> </ul>
--

Each  $\text{tableRank}$  has size  $\tilde{O}(2^{\epsilon\kappa})$  and each  $\text{tableInt}$  has size  $O(\kappa^{c+2})$ . The total size of  $\text{tableBlk}_{V, V_{sc}}$  is  $\tilde{O}(2^{\epsilon\kappa})$ .

**The main query algorithm.** We show how to answer each query  $x$  in constant time with high probability, by querying only the main data structure (and without knowing  $m$ ). We begin by decoding the two data structures  $\mathcal{D}_{\text{main},1}$  and  $\mathcal{D}_{\text{main},2}$  from  $\mathcal{D}_{\text{main}}$ , and query  $\mathcal{D}_{\text{main},1}$ .

<b>query algorithm</b> $\text{qalgBlk}_{\text{main}}(V, x)$ : 1. decode $\mathcal{D}_{\text{main},1}$ and $\mathcal{D}_{\text{main},2}$ from $\mathcal{D}_{\text{main}}$ using Proposition 8 2. $x_r := \mathcal{D}_{\text{main},1}.\text{qAlgRank}(V, \kappa^{2c-3}, x)$ (from Lemma 28) 3. if $x_r > \mathcal{D}_{\text{main},1}.\text{qAlgRank}(V, \kappa^{2c-3}, x - 1)$ 4. return $(1, x_r - 1)$	(to be cont'd)
---	----------------

$x_r$  is the number of elements in  $S \setminus S_{\text{unk}}$  that are at most  $x$ . Line 3 checks if  $x \in S \setminus S_{\text{unk}}$ . If it is, then  $x$  is the  $x_r$ -th element in  $S \setminus S_{\text{unk}}$ , and we return its hash value according by the definition of  $h$ .

5. $x_{\text{unk}} := \mathcal{D}_{\text{main},2}.\text{qAlgRank}(V - \kappa^{2c-3}, \kappa^c, x - x_r)$ 6. if $x_{\text{unk}} > \mathcal{D}_{\text{main},2}.\text{qAlgRank}(V - \kappa^{2c-3}, \kappa^c, x - x_r - 1)$ 7. return “unknown” 8. return $(0, x - x_r - x_{\text{unk}})$
--

If  $x \notin S \setminus S_{\text{unk}}$ , we query  $\mathcal{D}_{\text{main},2}$  to check if  $x \in S_{\text{unk}} \cup \bar{S}_{\text{unk}}$  in line 6. Note that  $x$  is the  $(x - x_r + 1)$ -th element in  $[V] \setminus (S \setminus S_{\text{unk}})$ . If it is, we return “unknown”. Otherwise, we know that  $x \notin S$ , and it is the  $(x - x_r - x_{\text{unk}} + 1)$ -th element in  $[V] \setminus (S \cup \bar{S}_{\text{unk}})$ , we return its  $\bar{h}$ -value. Since  $\text{qAlgRank}$  has constant query time,  $\text{qalgBlk}$  also runs in constant time. Clearly,  $\text{qalgBlk}$  outputs  $\text{hash}(x)$  when  $x \in S$  and  $h(x) \in [\kappa^{2c-3}]$ , or  $x \notin S$  and  $\bar{h}(x) \in [V - \kappa^{2c-3} - \kappa^c]$ , and otherwise it outputs “unknown”. This proves item (iv) in the statement.

Next, we show that the probability that it outputs “unknown” is small. To this end, let us fix the input data  $S, S_{sc}$  and query  $x$ , and let  $\mathcal{R}$  be uniformly random. We will show that  $S_{\text{unk}}$  is close to a uniformly random subset of  $S$  of size  $m - \kappa^{2c-3}$ , and  $\bar{S}_{\text{unk}}$  is close to a uniformly random subset of  $[V] \setminus S$  of size  $\kappa^{2c-3} + \kappa^c - m$ . By Claim 30, we have  $H(\mathcal{D}_{sc,1} \oplus \mathcal{R}) \geq (|\mathcal{D}_{sc,1}| - 2^{-\kappa+2})(1 - 2^{-\kappa+2})$ . Since the division operation in Proposition 14 is an injection, we have

$$\begin{aligned}
 |\mathcal{D}_{sc,1}| &\geq s_{sc} - |\mathcal{D}_{sc,2}| - |\mathcal{D}_{sc,3}| \\
 &\geq \lg \binom{m}{m - \kappa^{2c-3}} - \kappa^c V 2^{-\kappa+3}.
 \end{aligned}$$

Therefore,  $H(\mathcal{D}_{\text{sc},1} \oplus \mathcal{R}) \geq \lg \binom{m}{m - \kappa^{2c-3}} - \kappa^c 2^{-\kappa/2+4}$ . Furthermore, since  $\text{prepIntoSet}(m, m - \kappa^{2c-3}, \cdot)$  is an injection, we have  $H(S_1) \geq \lg \binom{m}{m - \kappa^{2c-3}} - \kappa^c 2^{-\kappa/2+4}$ , which in turn implies that

$$H(S_{\text{unk}}) \geq \lg \binom{m}{m - \kappa^{2c-3}} - \kappa^c 2^{-\kappa/2+4},$$

for any fixed  $S$  and  $S_{\text{sc}}$ . By Pinsker's inequality, the  $\ell_1$  distance between  $S_{\text{unk}}$  and a uniformly random subset of  $S$  of size  $m - \kappa^{2c-3}$  is at most  $O(\kappa^{c/2} 2^{-\kappa/4})$ . In particular, it implies that for any fixed  $x \in S$ , the probability that  $x \in S_{\text{unk}}$  is at most

$$\frac{m - \kappa^{2c-3}}{m} + O(\kappa^{c/2} 2^{-\kappa/4}) \leq O(\kappa^{-c+3}).$$

By applying the same argument to  $\mathcal{D}_{\text{sc},2}$ ,  $S_2$  and  $\overline{S}_{\text{unk}}$ , we conclude that for any fixed  $x \notin S$ , the probability that  $x \in \overline{S}_{\text{unk}}$  is at most

$$\frac{\kappa^{2c-3} + \kappa^c - m}{V - m} + O(\kappa^{c/2} 2^{-\kappa/4}) \leq O(\kappa^{-c+3}).$$

This proves item (v) in the statement.

**The general query algorithm.** Finally, we describe the query algorithm for all  $x \in [V + V_{\text{sc}}]$ . We use two different algorithms for  $x \in [V]$  and  $x \in \{V, \dots, V + V_{\text{sc}} - 1\}$ . We begin by the  $x \in [V]$  case ( $x$  is in the primary block).

**query algorithm**  $\text{qalgBlk}(V, m, V_{\text{sc}}, m_{\text{sc}}, x)$ :

1. (if  $x < V$ )
2.  $(b, v) := \text{qalgBlk}_{\text{main}}(V, x)$
3. if  $(b, v)$  is not “unknown”
4. return  $(b, v)$
5. let  $x_{\text{unk}} := \text{rank}_{S_{\text{unk}} \cup \overline{S}_{\text{unk}}}(x)$  (already computed in  $\text{qalgBlk}_{\text{main}}(V, x)$ ) (to be cont'd)

When  $\text{qalgBlk}_{\text{main}}$  returns “unknown”,  $x$  is the  $x_{\text{unk}}$ -th element in  $S_{\text{unk}} \cup \overline{S}_{\text{unk}}$ . Next, we query  $\mathcal{D}_{\text{aux},1}$  to find out whether  $x \in S_{\text{unk}}$  or  $x \in \overline{S}_{\text{unk}}$  and its rank in the corresponding set. Then we return its  $h$  or  $\overline{h}$  value according to the definition.

6. apply Proposition 8 on  $\mathcal{D}_{\text{aux}}$  to decode  $\mathcal{D}_{\text{aux},1}$
7.  $x_{\text{unk},r} := \mathcal{D}_{\text{aux},1}.\text{qAlgRank}(\kappa^c, m - \kappa^{2c-3}, x_{\text{unk}} - 1)$
8. if  $x_{\text{unk},r} > \mathcal{D}_{\text{aux},1}.\text{qAlgRank}(\kappa^c, m - \kappa^{2c-3}, x_{\text{unk}} - 2)$
9. return  $\kappa^{2c-3} + x_{\text{unk},r} - 1$
10. else
11. return  $V - \kappa^{2c-3} - \kappa^c + (x_{\text{unk}} - x_{\text{unk},r}) - 1$

Similarly to  $\text{qalgBlk}_{\text{main}}$ , we check if the  $x_{\text{unk}}$ -th element is in  $S_{\text{unk}}$ . if it is, then it is the  $x_{\text{unk},r}$ -th element in  $S_{\text{unk}}$ . Otherwise, it is the  $(x_{\text{unk}} - x_{\text{unk},r})$ -th element in  $\overline{S}_{\text{unk}}$ . In this case ( $x \in [V]$ ), the query algorithm runs in constant time.

Next, we show how to handle  $x \in \{V, \dots, V + V_{\text{sc}} - 1\}$ . To this end, let us first assume that we can make random access to  $\mathcal{D}_{\text{sc}}$ .

12. (if  $x \geq V$ )
13. apply Proposition 8 on  $\mathcal{D}_{\text{aux}}$  to decode  $\mathcal{D}_{\text{aux},2}$
14.  $x_r := \mathcal{D}_{\text{sc}}.\text{qAlgRank}(V_{\text{sc}}, m_{\text{sc}}, x - V)$  (from Lemma 28)
15. if  $x_r > \mathcal{D}_{\text{sc}}.\text{qAlgRank}(V_{\text{sc}}, m_{\text{sc}}, x - V - 1)$
16.     return  $m + x_r - 1$
17. else
18.     return  $V - m + (x - V - x_r)$

If we had access to  $\mathcal{D}_{\text{sc}}$ , then the query algorithm would be similar to the previous cases, and it runs in constant time. However,  $\mathcal{D}_{\text{sc}}$  is not stored in the data structure explicitly. In the following, we show how `qalgBlk` accesses  $\mathcal{D}_{\text{sc}}$  from its implicit representation.

More specifically, `qalgBlk` only needs to access  $\mathcal{D}_{\text{sc}}$  when it runs the query algorithm `qAlgRank` on  $\mathcal{D}_{\text{sc}}$ . By Lemma 28, `qAlgRank` runs on a RAM with word-size  $\Theta(\kappa)$ , i.e., it may request  $\Theta(\kappa)$  consecutive bits of the data structure  $\mathcal{D}_{\text{sc}}$  during its runtime. To implement such access requests, we first apply Proposition 14 to reduce each access to  $O(1)$  accesses to  $\mathcal{D}_{\text{sc},1}$ ,  $\mathcal{D}_{\text{sc},2}$  and  $\mathcal{D}_{\text{sc},3}$ .  $\mathcal{D}_{\text{sc},3}$  is stored as  $\mathcal{D}_{\text{aux},2}$ , which has been decoded. Each access to it can be implemented in constant time. For  $\mathcal{D}_{\text{sc},1}$ ,  $\mathcal{D}_{\text{sc},1} \oplus \mathcal{R}$  is interpreted as a set  $S_1 \subseteq [m]$  of size  $m - \kappa^{2c-3}$ . Lemma 29 guarantees that each access to  $\mathcal{D}_{\text{sc},1} \oplus \mathcal{R}$  can be implemented in  $O(\kappa^4)$  time and  $O(\kappa^2)$  rank queries to  $S_1$ , which by the previous argument, implies that each access to  $\mathcal{D}_{\text{sc},1}$  can also be implemented in the same time and number of rank queries.

On the other hand, the way the preprocessing algorithm “encodes”  $S_1$  guarantees that  $\text{rank}_{S_1}(k)$  queries can be implemented efficiently. To see this, recall that  $S_{\text{unk}} \subset S$  is determined according to  $S_1$ .  $\text{rank}_{S_1}(k)$  is exactly the number of elements in  $S_{\text{unk}}$  that are no larger than the  $(k+1)$ -th smallest element in  $S$ . We first do a binary search to find the  $(k+1)$ -th smallest element in  $S$ .

**implementing rank queries on  $S_1$   $\text{rank}_{S_1}(k)$ :**

1. decode  $\mathcal{D}_{\text{main},1}$ ,  $\mathcal{D}_{\text{main},2}$  and  $\mathcal{D}_{\text{aux},1}$
2. binary search for  $(k+1)$ -th element  $x^*$  in  $S$ : given  $x \in [V]$ ,
  - (i)  $x_r := \mathcal{D}_{\text{main},1}.\text{qAlgRank}(V, \kappa^{2c-3}, x)$
  - (ii)  $x_{\text{unk}} := \mathcal{D}_{\text{main},2}.\text{qAlgRank}(V - \kappa^{2c-3}, \kappa^c, x - x_r)$
  - (iii)  $\text{rank}_S(x) := x_r + \mathcal{D}_{\text{aux},1}.\text{qAlgRank}(\kappa^c, m - \kappa^{2c-3}, x_{\text{unk}} - 1)$

$x_r$  is the number of elements in  $S \setminus S_{\text{unk}}$  that are at most  $x$ .  $x_{\text{unk}}$  is the number of elements in  $S_{\text{unk}} \cup \bar{S}_{\text{unk}}$  that are at most  $x$ .  $\mathcal{D}_{\text{aux},1}.\text{qAlgRank}(\kappa^c, m - \kappa^{2c-3}, x_{\text{unk}} - 1)$  computes the number of elements in  $S_{\text{unk}}$  that are at most  $x$ . By summing up  $x_r$  and  $\mathcal{D}_{\text{aux},1}.\text{qAlgRank}(\kappa^c, m - \kappa^{2c-3}, x_{\text{unk}})$ , we compute  $\text{rank}_S(x)$ , the number of elements in  $S$  that are at most  $x$ , in constant time. Being able to compute  $\text{rank}_S(x)$  for any given  $x$  allows us to binary search for the  $(k+1)$ -th smallest element  $x^*$  in  $S$  in  $O(\lg V) = O(\kappa)$  time, which then allows us to compute  $\text{rank}_{S_1}(k)$ .

3.  $x_r^* := \mathcal{D}_{\text{main},1}.\text{qAlgRank}(V, \kappa^{2c-3}, x^*)$
4. return  $\text{rank}_{S_1}(k) := k - x_r^* + 1$

This shows that  $\text{rank}_{S_1}(k)$  can be computed in  $O(\kappa)$  time, and thus, each access to  $\mathcal{D}_{\text{sc},1}$  can be implemented in  $O(\kappa^4 + \kappa \cdot \kappa^2) = O(\kappa^4)$  time.

Similarly, each access to  $\mathcal{D}_{\text{sc},2}$  can be implemented in  $O(\kappa^4)$  time: Lemma 29 reduces it to  $O(\kappa^2)$  rank queries to  $S_2$  and  $O(\kappa^4)$  processing time; For  $\text{rank}_{S_2}(k)$ , we do binary search to find the  $(k+1)$ -th element in  $[V] \setminus S$ ; By querying  $\mathcal{D}_{\text{main},1}$ ,  $\mathcal{D}_{\text{main},2}$  and  $\mathcal{D}_{\text{aux},1}$ , we compute  $\text{rank}_{S_2}(k)$ .

Overall, the above algorithms allow us to access  $\mathcal{D}_{\text{sc},1}$ ,  $\mathcal{D}_{\text{sc},2}$  and  $\mathcal{D}_{\text{sc},3}$  in  $O(\kappa^4)$  time, which in turn, allows us to access  $\mathcal{D}_{\text{sc}}$  in  $O(\kappa^4)$ . Thus, `qalgBlk` runs in  $O(\kappa^4)$  time. This proves item (vi) in the statement.  $\square$

## 7.1 Data interpretation

In this subsection, we prove Lemma 29, showing how to represent a string as a set which allows us to locally decode the string given access to a rank oracle of the resulting set.

**Lemma 29 (restated).** *Let  $c \geq 2$  be a constant positive integer. There is a preprocessing algorithm `prepIntoSet`, a query algorithm `request` and lookup tables `tableIntV,m` of sizes  $O(\kappa^{c+2})$ , such that for any integers  $V$  and  $m$  where  $V \leq 2^{\kappa/2}$  and  $m \leq \kappa^c$ , given a (double-ended) string  $\mathcal{D} = (K_h, M, K_t)$  of length*

$$s \leq \lg \binom{V}{m} - m(V-1)2^{-\kappa+2},$$

*`prepIntoSet`( $V, m, \mathcal{D}$ ) outputs a set  $S \subseteq [V]$  of size  $m$ . For any  $-1 \leq a_1 \leq a_2 \leq |M|$  and  $a_2 < a_1 + \kappa$ , `request`(0,  $V, m, \text{range}(K_h), |M|, \text{range}(K_t), a_1, a_2$ ) computes  $\mathcal{D}[a_1, a_2]$  in  $O(\kappa^4)$  time using  $O(\kappa^2)$  rank queries to  $S$ , assuming it can make random accesses to the lookup table `tableIntV,m`.*

To construct set  $S$  from the input string  $\mathcal{D}$ , the main idea is to apply Proposition 15 and 14, and then recurse on the two halves of  $[V]$ . We extract an integer  $m_1 \in [m+1]$  from  $\mathcal{D}$  using Proposition 15, which encodes the number of elements in the first half of  $[V]$ . Then we divide  $\mathcal{D}$  into two data structure  $\mathcal{D}_1$  and  $\mathcal{D}_2$  such that the length of  $\mathcal{D}_1$  is approximately  $\lg \binom{V/2}{m_1}$  and the length of  $\mathcal{D}_2$  is approximately  $\lg \binom{V/2}{m-m_1}$ . Then the set  $S$  is recursively constructed in the two halves. When the  $m$  gets sufficiently small and the  $\mathcal{D}$  has length  $O(\kappa^2)$ , we continue the recursion without applying the two propositions about fraction-length strings. Instead, we take the whole string as an integer smaller than  $2^{O(\kappa^2)}$ , and use the integer to encode a set (also decode the whole integer at the decoding time). See below for the formal argument.

**Encoding and decoding an integer using a set.** We first show that given an integer  $Z \leq \binom{V}{m}$ , how to turn it into a set of size  $m$  in  $[V]$ , such that  $Z$  can be recovered using rank queries.

**encoding algorithm** `encSet`( $V, m, Z$ ):

1. if  $m = 0$ , return  $\emptyset$
2. if  $m = V$ , return  $[V]$
3.  $V_1 := \lfloor V/2 \rfloor$  and  $V_2 := \lceil V/2 \rceil$
4. compute the largest  $0 \leq j \leq m$  such that  $Z \geq \sum_{i=0}^{j-1} \binom{V_1}{i} \binom{V_2}{m-i}$
5.  $Y := Z - \sum_{i=0}^{j-1} \binom{V_1}{i} \binom{V_2}{m-i}$
6.  $Z_1 := Y \text{ div } \binom{V_2}{m-j}$  and  $Z_2 := Y \text{ mod } \binom{V_2}{m-j}$
7. return `encSet`( $V_1, j, Z_1$ )  $\cup$  (`encSet`( $V_2, m-j, Z_2$ ) +  $V_1$ )

To construct the set, the algorithm is a standard recursion. All possible sets are listed in the increasing order of the number of elements in the  $[V_1]$ . We compute this number, and then recurse into the two halves.  $Z$  can be recovered by the following algorithm, assuming the set generated is in the universe  $[X, X+V]$ . For technical reasons that we will encounter later, sometimes we may only have access to the *complement* of the set. The bit  $b$  indicates whether we should take the complement.

**decoding algorithm** `decSet`( $X, V, m, b$ ):

1. if  $m = 0$  or  $m = V$ , return 0
2.  $V_1 := \lfloor V/2 \rfloor$  and  $V_2 := \lceil V/2 \rceil$
3.  $j := \text{rank}_S(X + V_1 - 1) - \text{rank}_S(X - 1)$
4. if  $b$ , then  $j := V_1 - j$
5.  $Z_1 := \text{decSet}(X, V_1, j, b)$  and  $Z_2 := \text{decSet}(X + V_1, V_2, m-j, b)$
6. return  $Z := \sum_{i=0}^{j-1} \binom{V_1}{i} \binom{V_2}{m-i} + Z_1 \cdot \binom{V_2}{m-j} + Z_2$

We will store the sum  $\sum_{i=0}^{j-1} \binom{V_1}{i} \binom{V_2}{m-i}$  and the binomial coefficient  $\binom{V_2}{m-j}$  in the lookup table. Since the recursion terminates when  $m = 0$  and the value of  $V$  decreases by a factor of two each time, the size of the recursion tree is  $O(m \lg V)$ . Thus, we have the following claim.

**Claim 31.** *decSet uses  $O(m \lg V)$  arithmetic operations on  $O(\lg \binom{V}{m})$ -bit integers, as well as  $O(m \lg V)$  rank queries.*

**Preprocessing into a set.** Given a string  $\mathcal{D} = (K_h, M, K_t)$  of length at most  $\lg \binom{V}{m} - m(V-1)2^{-\kappa+2}$ , we preprocess it into a set  $S \subseteq [V]$  of size  $m$ .

**preprocessing algorithm** `prepIntoSet`( $V, m, \mathcal{D} = (K_h, M, K_t)$ ):

1. if  $2m > V$
2.   return  $[V] \setminus \text{prepIntoSet}(V, V-m, \mathcal{D})$
3. if  $m \leq 24\kappa$
4.   rewrite  $\mathcal{D}$  as a nonnegative integer  $Z < \text{range}(K_h) \cdot \text{range}(K_t) \cdot 2^{|M|}$
5.   return `encSet`( $V, m, Z$ )
6. if  $|\mathcal{D}| \leq 24\kappa$
7.   return `prepIntoSet`( $48\kappa, 24\kappa, \mathcal{D}$ )  $\cup \{V - (m - 24\kappa), \dots, V - 1\}$  (to be cont'd)

If  $m$  is larger than  $V/2$ , we simply work on the complement. If  $m$  is  $O(\kappa)$ , we view the entire string  $\mathcal{D}$  as an integer, and call `encSet`. If the string is too short while  $m$  (and  $V$ ) are still large, we manually decrease  $m$  and  $V$ , and reduce it to the  $m = O(\kappa)$  case. Note that  $\binom{V}{m}$  may be at most  $2^{O(\kappa^2)}$ ,  $Z$  occupies  $O(\kappa)$  words (as  $\kappa = \Theta(w)$ ).

Otherwise, we extract an integer  $j$  from  $\mathcal{D}$ .

8.  $V_1 := \lfloor V/2 \rfloor$  and  $V_2 := \lceil V/2 \rceil$
9. compute  $s_j := \lg \binom{V_1}{j} + \lg \binom{V_2}{m-j} - m(V-2)2^{-\kappa+2}$
10. apply Proposition 15 for  $j \in \{\lfloor m/3 \rfloor + 1, \dots, 2\lfloor m/3 \rfloor\}$  and  $C = \lfloor m/3 \rfloor$ , encode  $\mathcal{D}$  using a pair  $(j, \mathcal{D}_j)$  such that  $\mathcal{D}_j$  has length at most  $s_j$
11. let  $(S_1, S_2) := \text{prepIntoTwo}(V_1, V_2, j, m-j, \mathcal{D}_j)$
12. return  $S_1 \cup (S_2 + V_1)$

`prepIntoTwo` preprocesses  $\mathcal{D}_j$  into two sets of sizes  $j$  and  $m-j$  over the two halves of the universe (see below). Then we return their union. Proposition 15 requires that the length of  $\mathcal{D}$  is at least  $3\kappa + 2$  and at most  $\lg(\sum_j 2^{s_j}) - (C-1)2^{-\kappa+2}$ . This is true, because on one hand, the length of  $\mathcal{D}$  is at least  $24\kappa$ ; on the other hand,

$$\begin{aligned}
2^{s_1} + \dots + 2^{s_C} &= \sum_{j=\lfloor m/3 \rfloor + 1}^{2\lfloor m/3 \rfloor} \binom{V_1}{j} \cdot \binom{V_2}{m-j} \cdot 2^{-m(V-2)2^{-\kappa+2}} \\
&\geq 2^{-m(V-2)2^{-\kappa+2}} \cdot \left( \binom{V}{m} - \frac{2m}{3} \binom{V_1}{\lfloor m/3 \rfloor} \binom{V_2}{\lceil 2m/3 \rceil} \right) \\
&= 2^{-m(V-2)2^{-\kappa+2}} \cdot \left( \binom{V}{m} - \frac{2m}{3} \binom{V_1}{\lfloor m/2 \rfloor} \binom{V_2}{\lceil m/2 \rceil} \cdot \prod_{j=\lfloor m/3 \rfloor + 1}^{\lfloor m/2 \rfloor} \frac{j(V_2 - m + j)}{(V_1 - j + 1)(m - j + 1)} \right) \\
&\geq 2^{-m(V-2)2^{-\kappa+2}} \cdot \binom{V}{m} \cdot \left( 1 - \frac{2m}{3} \cdot \prod_{j=\lfloor m/3 \rfloor + 1}^{\lfloor m/2 \rfloor} \frac{j}{m - j + 1} \right)
\end{aligned}$$

$$\begin{aligned}
&\geq 2^{-m(V-2)2^{-\kappa+2}} \cdot \binom{V}{m} \cdot \left(1 - \frac{2m}{3} \cdot e^{-\sum_{j=\lfloor m/3 \rfloor + 1}^{\lfloor m/2 \rfloor} \frac{m-2j+1}{m-j+1}}\right) \\
&\geq 2^{-m(V-2)2^{-\kappa+2}} \cdot \binom{V}{m} \cdot \left(1 - \frac{2m}{3} \cdot e^{-m/24}\right).
\end{aligned}$$

Therefore, by the fact that  $m \geq 24\kappa$ , we have

$$\begin{aligned}
\lg(2^{s_1} + \dots + 2^{s_C}) &\geq \lg \binom{V}{m} - m(V-2)2^{-\kappa+2} - m2^{-\kappa} \\
&\geq s + m2^{-\kappa+2} - m2^{-\kappa} \\
&\geq s + (C-1) \cdot 2^{-\kappa+2}.
\end{aligned}$$

Thus, the premises of Proposition 15 are also satisfied. Next, we describe `prepIntoTwo`.

**preprocessing algorithm** `prepIntoTwo`( $V_1, V_2, m_1, m_2, \mathcal{D}$ ):

1. let  $s_1 := \lg \binom{V_1}{m_1} - m_1(V_1-1)2^{-\kappa+2}$  and  $s_2 := \lg \binom{V_2}{m_2} - m_2(V_2-1)2^{-\kappa+2}$
2. apply Proposition 14, divide  $\mathcal{D}$  into  $\mathcal{D}_1$  and  $\mathcal{D}_2$  of lengths at most  $s_1$  and  $s_2$  respectively
3. let  $S_1 := \text{prepIntoSet}(V_1, m_1, \mathcal{D}_1)$  and  $S_2 := \text{prepIntoSet}(V_2, m_2, \mathcal{D}_2)$
4. return  $(S_1, S_2)$

Proposition 14 requires that the length of  $\mathcal{D}$  is at most  $s_1 + s_2 - 2^{-\kappa+2}$  (and at least  $3\kappa$ ), and  $s_1, s_2 \geq 3\kappa$ . It is easy to verify the former. For the latter, because  $m_1 + m_2 \leq V/2$ ,  $m_2/2 \leq m_1 \leq 2m_2$  and  $m_1 + m_2 > 24\kappa$ , and in particular, we have  $V_1 \geq 24\kappa$  and  $m_1 \in [V_1/3, 2V_1/3]$ . Hence, we have

$$\binom{V_1}{m_1} \geq 3^{8\kappa},$$

and it implies  $s_1 \geq 8\kappa$ . Similarly, we also have  $s_2 \geq 8\kappa$ .

**Lookup table.** We store the following lookup table.

**lookup table** `tableIntV,m`:

1. if  $m \leq 24\kappa$
2.  $\sum_{i=0}^{j-1} \binom{V_1}{i} \binom{V_2}{m-i}$  for  $j = 0, \dots, m$
3.  $\binom{V_2}{j}$  for all  $j = 0, \dots, m$
4. else
5. lookup table from Proposition 15 for line 10 of `prepIntoSet`
6. include all tables `tableIntV',m'` for  $V' = \lfloor V/2^i \rfloor$  or  $V' = \lceil V/2^i \rceil$  for  $i \geq 1$ , and  $0 \leq m' \leq m$

The lookup table `tableIntV,m` itself has size at most  $O(\kappa)$  words for  $m > 24\kappa$  and  $O(\kappa^2)$  words for  $m \leq 24\kappa$ . Including the smaller tables, its total size is at most  $O(\kappa^2 m + \kappa^4) \leq O(\kappa^{c+2})$  words for  $m \leq \kappa^c$  and  $c \geq 2$ .

**Access the string.** Suppose  $S$  is the set generated from a string  $\mathcal{D}$  using the above preprocessing algorithm. In the following, we show how to access  $\mathcal{D}[a_1, a_2]$  for  $a_2 - a_1 < \kappa$ , assuming rank queries can be computed efficiently on  $S$ . Assuming the set  $S$  restricted to  $[X, X+V)$  (with  $m$  elements in this range) is generated from a string  $\mathcal{D} = (K_h, M, K_t)$ , `request`( $X, V, m, \text{range}(K_h), |M|, \text{range}(K_t), a_1, a_2, b$ ) recovers  $\mathcal{D}[a_1, a_2]$ , where  $b$  indicates if we take the complement of  $S$ .

**accessing algorithm** `request`( $X, V, m, \text{range}(K_h), |M|, \text{range}(K_t), a_1, a_2, b$ ):

1. if  $2m > V$
2.    $m := V - m$  and  $b := \neg b$
3. if  $m \leq 24\kappa$
4.    $Z := \text{decSet}(x, V, m, b)$
5.   rewrite  $Z$  as a string  $\mathcal{D} = (K_h, M, K_t)$
6.   return  $\mathcal{D}[a_1, a_2]$
7. if  $\lg(\text{range}(K_h)) + |M| + \lg(\text{range}(K_t)) \leq 24\kappa$
8.   return `request`( $X, 48\kappa, 24\kappa, \text{range}(K_h), |M|, \text{range}(K_t), a_1, a_2, b$ ) (to be cont'd)

If  $S$  has small size, we recover the whole data structure using `decSet`. If  $\mathcal{D}$  is too short, we reduce  $m$  and  $V$ .

9.  $V_1 := \lfloor V/2 \rfloor, V_2 := \lceil V/2 \rceil$
10. ask `rank` queries and compute  $j := \text{rank}_S(X + V_1 - 1) - \text{rank}_S(X - 1)$
11. if  $b$ , then  $j := V_1 - j$
12. find the size of  $\mathcal{D}_j = (K_{j,h}, M_j, K_{j,t})$  in the lookup table

We compute  $j$ , the integer extracted from  $\mathcal{D}$ , which encodes the number of elements in the first half. We recover the size of  $\mathcal{D}_j$ , and use the fact that  $(M_j, K_{j,t})$  is a suffix of  $\mathcal{D}$  (by Proposition 15) to recurse.

13. if  $a_1 \geq |M| - |M_j|$
14.   return `reqFromTwo`( $X, V_1, V_2, j, m - j, \text{range}(K_{j,h}), |M_j|, \text{range}(K_{j,t}), a_1 - (|M| - |M_j|), a_2 - (|M| - |M_j|), b$ )
15. else
16.   recover  $\mathcal{D}_j[-1, a_2 - (|M| - |M_j|)] :=$   
       `reqFromTwo`( $X, V_1, V_2, j, m - j, \text{range}(K_{j,h}), |M_j|, \text{range}(K_{j,t}), -1, a_2 - (|M| - |M_j|), b$ )
17.   compute  $\mathcal{D}[a_1, a_2]$  using Proposition 15

`reqFromTwo` recovers the requested substring of  $\mathcal{D}$  assuming `rank` queries to the set generated from `prepIntoTwo`. Since  $(M_j, K_{j,t})$  is a suffix of  $\mathcal{D}$ , if  $\mathcal{D}[a_1, a_2]$  is entirely contained in this range, we simply recurse on  $\mathcal{D}_j$ . Otherwise, Proposition 15 guarantees that the remaining bits can be recovered from  $j$  and  $K_{j,h}$ . Note that in either case, the difference  $a_2 - a_1$  does not increase.

Next, we describe `reqFromTwo`.

**accessing algorithm** `reqFromTwo`( $X, V_1, V_2, m_1, m_2, \text{range}(K_{j,h}), |M_j|, \text{range}(K_{j,t}), a_1, a_2, b$ ):

1. compute the sizes of  $\mathcal{D}_1 = (K_{1,h}, M_1, K_{1,t})$  and  $\mathcal{D}_2 = (K_{2,h}, M_2, K_{2,t})$ , which  $\mathcal{D}$  is divided into (to be cont'd)

Suppose  $S$  restricted to  $[X, X + V_1)$  and  $[X + V_1, X + V_1 + V_2)$  is generated from  $\mathcal{D}$  using `prepIntoTwo`. Then by Proposition 14,  $(K_{1,h}, M_1)$  is a prefix of  $\mathcal{D}$  and  $(M_2, K_{2,t})$  is a suffix.

2. if  $a_1 \geq |M| - |M_2|$
3.   return `request`( $X + V_1, V_2, m_2, \text{range}(K_{2,h}), |M_2|, \text{range}(K_{2,t}), a_1 - (|M| - |M_2|), a_2 - (|M| - |M_2|), b$ )
4. if  $a_2 < |M_1|$
5.   return `request`( $X, V_1, m_1, \text{range}(K_{1,h}), |M_1|, \text{range}(K_{1,t}), a_1, a_2, b$ ) (to be cont'd)

If the requested bits  $\mathcal{D}[a_1, a_2]$  are entirely contained in  $\mathcal{D}_1$  or  $\mathcal{D}_2$ , we simply recurse on the corresponding substring. In this case, the difference  $a_2 - a_1$  does not change either.

6. recover  $\mathcal{D}_1[a_1, |M_1|] := \text{request}(X, V_1, m_1, \text{range}(K_{1,h}), |M_1|, \text{range}(K_{1,t}), a_1, |M_1|, b)$
7. recover  $\mathcal{D}_2[-1, a_2 - (|M| - |M_2|)] :=$   
       `request`( $X + V_1, V_2, m_2, \text{range}(K_{2,h}), |M_2|, \text{range}(K_{2,t}), -1, a_2 - (|M| - |M_2|), b$ )
8. reconstruct  $\mathcal{D}[a_1, a_2]$  using Proposition 14

Finally, if the requested bits  $\mathcal{D}[a_1, a_2]$  split across both substrings, then we make two recursive calls.

**Query time.** Next, we analyze the query time. First observe that `request` has at most  $O(\lg m)$  levels of recursion before we call `decSet`. This is because each time  $m$  is reduced at least by a factor of  $1/3$  by the preprocessing algorithm. The only place that the whole recursion makes more than one recursive calls is line 6 and line 7 in `reqFromTwo`. In all other cases, the algorithm makes at most one recursive call with the same (or smaller) difference  $a_2 - a_1$ . Moreover, we claim that those two lines can only be executed at most once throughout the whole recursion.

**Claim 32.** *Line 6 and line 7 in `reqFromTwo` can at most be executed once throughout the whole recursion.*

*Proof.* When these two lines are executed, the two recursive calls will both request either a prefix or a suffix of the substring. Also, as we observed above, the difference  $a_2 - a_1$  never increases throughout the recursion. The recursive call that requests a prefix will have  $a_1 = -1$  and  $a_2 < \kappa - 1$ . Thereafter, any subsequence recursive calls in this branch will have  $a_1 = -1$  and  $a_2 < \kappa - 1$ . Since Proposition 14 always generates two strings of length at least  $3\kappa$ , line 4 in `reqFromTwo` is always true (as  $|M_1| \geq \kappa - 1$ ). Line 6 and line 7 will hence not be executed in this branch. The recursive branch that requests a suffix is similar, in which line 2 in `reqFromTwo` is always true. This proves the claim.  $\square$

Claim 32 implies that the whole recursion tree has at most  $O(\lg m)$  nodes, and at most two leaves. In each node, the algorithm spends constant time, and makes two rank queries. In each leaf, the algorithm makes one call to `decSet`. As we argued earlier, the integer  $Z \leq \lg \binom{V}{m}$  has at most  $O(\kappa^2)$  bits (and  $O(\kappa)$  words). Since  $m \leq O(\kappa)$  when `decSet` is called, by Claim 31, each `decSet` takes  $O(\kappa^4)$  time ( $O(\kappa)$ -word numbers take  $O(\kappa^2)$  time to multiply or divide), and makes  $O(\kappa^2)$  rank queries. Combining the above facts, we conclude that `request` runs in  $O(\kappa^4)$  time, and it makes at most  $O(\kappa^2)$  rank queries. This proves Lemma 29.

## 7.2 Small sets

In this section, we prove Lemma 28, which constructs a succinct rank data structure for sets of size  $\kappa^{O(1)}$ , with constant query time. We first show that the fusion trees [FW93] can be implemented succinctly. This gives us a data structure for small sets with a sublinear, although large, redundancy.

**Lemma 33.** *Let  $c$  be any constant positive integer and  $\epsilon$  be any positive constant. There is a preprocessing algorithm `prepRankL`, a query algorithm `qAlgRankL` and lookup tables `tableRankL` <sub>$V, m$</sub>  of sizes  $2^{\epsilon\kappa}$  such that for any integers  $V, m$  such that  $V \leq 2^\kappa$  and  $m \leq \kappa^c$ , given a set  $S \subset [V]$  of size  $m$ , `prepRankL` preprocesses it into a data structure using*

$$\lg \binom{V}{m} + \frac{1}{8}m \lg \kappa$$

*bits of space. Given any  $x \in [V]$ , `qAlgRankL` compute `rankS(x)` in constant time, by accessing the data structure and `tableRankL` <sub>$V, m$</sub> . The algorithms run on a random access machine with word-size  $w \geq \Omega(\kappa)$ .*

Since the main ideas are similar, we may omit the proof of a few claims in the construction, and refer the readers to the original fusion trees for details ([FW93]).

*Proof. (sketch)* Let  $S = \{y_1, \dots, y_m\}$  and  $y_1 < y_2 < \dots < y_m$ . Let us first show how to construct such a data structure using

$$m \lceil \lg V \rceil + m \lceil \lg \kappa \rceil$$

bits when  $m \leq \kappa^{1/4}$ . We view each  $y_i$  as a  $\lceil \lg V \rceil$ -bit binary string, and consider the first bit where  $y_i$  and  $y_{i+1}$  differ, for every  $i = 1, \dots, m-1$ . Let  $W$  be this set of bits, i.e.,  $j \in W$  if and only there exists some  $i$  such that  $j$ -th bit is the first bit where  $y_i$  and  $y_{i+1}$  differ. Then  $|W| \leq m-1$ . Similar to fusion trees, let  $\text{sketch}(y)$  be  $y$  restricted to  $W$ . Observe that we must have  $\text{sketch}(y_1) < \text{sketch}(y_2) < \dots < \text{sketch}(y_m)$ .

The data structure first stores  $W$  using  $m \lceil \lg \kappa \rceil$  bits. Then it stores  $\text{sketch}(y_1), \dots, \text{sketch}(y_m)$ . Finally, the data structure stores the remaining bits of each  $y_i$ , for  $i = 1, \dots, m$  and from the top bits to the low bits. It is clear that the data structure occupies  $m \lceil \lg V \rceil + m \lceil \lg \kappa \rceil$  bits of space.

To answer a query  $x \in [V]$ ,  $\text{qAlgRankL}_{V,m}$  first breaks  $x$  into  $\text{sketch}(x)$  and the remaining bits. That is, it generates two strings:  $x$  restricted to  $W$  (a  $|W|$ -bit string), and the remaining bits (a  $(\lceil \lg V \rceil - |W|)$ -bit string). It can be done in constant time using a lookup table of size  $2^{O(\epsilon \kappa)}$ , e.g., we divide the bits of  $x$  into chunks of length  $\epsilon \kappa$ , and store in  $\text{tableRankL}_{V,m}$  for each chunk, every possible set  $W$  and every possible assignment to the bits of  $x$  in the chunk, their contribution to  $\text{sketch}(x)$  and the remaining bits (note that there are only  $2^{o(\kappa)}$  different sets  $W$ ). Summing over all chunks gives us  $\text{sketch}(x)$  and the remaining bits. The query algorithm then finds the unique  $i$  such that  $\text{sketch}(y_i) \leq \text{sketch}(x) < \text{sketch}(y_{i+1})$ . This can be done by storing a lookup table of size at most  $2^{(m+1)|W|} \leq 2^{\kappa^{1/2}}$ , since  $(\text{sketch}(y_1), \dots, \text{sketch}(y_m))$  has only  $m|W|$  bits, and  $\text{sketch}(x)$  has  $|W|$  bits. However, we might not necessarily have  $y_i \leq x < y_{i+1}$ , but similar to the arguments in fusion trees,  $x$  has the longest common prefix (LCP) with either  $y_i$  or  $y_{i+1}$  (among all  $y \in S$ ).  $\text{qAlgRankL}_{V,m}$  next computes the LCP between  $x$  and  $y_i$  and the LCP between  $x$  and  $y_{i+1}$ . Both can be done in constant time, since to compute the LCP between  $x$  and  $y_i$ , it suffices to compute the LCP between  $\text{sketch}(x)$  and  $\text{sketch}(y_i)$  and the LCP between their remaining bits. Suppose  $x$  and  $y_{i^*}$  have a longer LCP ( $i^* = i$  or  $i+1$ ). If  $x = y_{i^*}$ , then  $\text{rank}_S(x) = i^*$ . Otherwise, let their common prefix be  $x'$ . If  $x > y_{i^*}$ , then let  $j$  be the unique index such that  $\text{sketch}(y_j) \leq \text{sketch}(x'111\dots 11) < \text{sketch}(y_{j+1})$ . The argument from fusion trees shows that we must have  $y_j < x < y_{j+1}$ , i.e.,  $\text{rank}_S(x) = j$ . Likewise, if  $x < y_{i^*}$ , then let  $j$  be the unique index such that  $\text{sketch}(y_j) < \text{sketch}(x'000\dots 00) \leq \text{sketch}(y_{j+1})$ . We must have  $y_j < x < y_{j+1}$ . By computing the value of  $j$  using the lookup table again, we find the number of elements in  $S$  that is at most  $x$ . Note that this data structure also allows us to retrieve each  $y_i$  in constant time.

Next, we show that the above data structure generalizes to any  $m \leq \kappa^c$ , and uses space

$$m(\lg V + (c+3)\lg \kappa) \leq \lg \binom{V}{m} + (2c+3)m \lg \kappa.$$

When  $m > \kappa^{1/4}$ , let  $B = \lfloor \kappa^{1/4} \rfloor$ , we take  $B$  evenly spaced elements from  $S$ , i.e.,  $y_{\lceil im/B \rceil}$  for  $i = 1, \dots, B$ . Denote the set of these  $B$  elements by  $S' = \{y'_1, \dots, y'_B\}$ , where  $y'_i = y_{\lceil im/B \rceil}$ . We apply the above data structure to  $S'$ , using space

$$B \lceil \lg V \rceil + B \lceil \lg \kappa \rceil < B(\lg V + \lg \kappa + 2).$$

Then, we recurse on all  $B$  subsets between elements in  $S'$ , where the  $i$ -th subset has  $\lceil im/B \rceil - \lceil (i-1)m/B \rceil - 1$  elements. Then the final data structure stores

- the data structure for  $S'$ ;
- $B$  data structures for all subsets between elements in  $S'$ ;
- an array of  $B$  pointers, pointing to the starting locations of the above  $B$  data structures.

We assign  $(c + 3/2) \lg \kappa$  bits to each pointer.

Suppose for each subset, we are able to (recursively) construct a data structure using

$$(\lceil im/B \rceil - \lceil (i-1)m/B \rceil - 1)(\lg V + (c+3) \lg \kappa)$$

bits of space. The total space usage is

$$B(\lg V + \lg \kappa + 2) + (m - B)(\lg V + (c+3) \lg \kappa) + B(c + 3/2) \lg \kappa \leq m(\lg V + (c+3) \lg \kappa).$$

On the other hand, assigning  $(c + 3/2) \lg \kappa$  bits to each pointer is sufficient, because

$$\lg(m(\lg V + (c+3) \lg \kappa)) \leq \lg(m\kappa + (c+3)m \lg \kappa) \leq (c+1) \lg \kappa + 1.$$

To answer query  $x$ , we first query the data structure for  $S'$ , and find the  $i$  such that  $y'_i \leq x < y'_{i+1}$ . Then we recurse into the  $i$ -th subset. The query time is constant, because the size of the set reduces by a factor of  $B = \Theta(\kappa^{1/4})$  each time. Note that for any given  $i$ , this data structure can also return  $y_i$  in constant time.

Finally, we show that the redundancy  $(2c+3)m \lg \kappa$  can be reduced to  $\frac{1}{8}m \lg \kappa$ . To this end, let  $S'$  be the subset of  $S$  with gap  $16(2c+3)$ , i.e.,  $S' = \{y'_1, y'_2, \dots\}$  such that  $y'_i = y_{16(2c+3) \cdot i}$ . Then  $|S'| = \lfloor \frac{m}{16(2c+3)} \rfloor$ . We construct a data structure for  $S'$  using space

$$|S'|(\lg V + (c+3) \lg \kappa).$$

Naturally,  $S'$  partitions  $S$  into chunks of  $16(2c+3) - 1$  elements. We simply write them down using

$$(16(2c+3) - 1) \lceil \lg(y'_{i+1} - y'_i - 1) \rceil$$

bits for chunk  $i$ . The final data structure consists of

1. the data structure for  $S'$ ,
2. all other elements in  $S$  encoded as above,
3.  $|S'| + 1$  pointers to each chunk.

We assign  $\lceil (c + 3/2) \lg \kappa \rceil$  bits to each pointer. By the concavity of  $\lg x$ , the total space usage is

$$\begin{aligned} & |S'|(\lg V + (c+3) \lg \kappa) + \sum_i (16(2c+3) - 1) \lceil \lg(y'_{i+1} - y'_i - 1) \rceil + (|S'| + 1) \lceil (c + 3/2) \lg \kappa \rceil \\ & \leq |S'| \lg \frac{V}{m} + |S'| (3c+5) \lg \kappa + \sum_i (16(2c+3) - 1) \lg \frac{V}{|S'| + 1} + m \\ & \leq |S'| \lg \frac{V}{m} + \frac{(3c+5)m}{16(2c+3)} \lg \kappa + \sum_i (16(2c+3) - 1) \lg \frac{V}{m} + O(m) \\ & \leq m \lg \frac{V}{m} + \frac{(3c+5)m}{16(2c+3)} \lg \kappa + O(m) \\ & \leq \lg \binom{V}{m} + \frac{m}{8} \lg \kappa. \end{aligned}$$

To answer query  $x$ , we first query the data structure for  $S'$ , and find  $i$  such that  $y'_i \leq x < y'_{i+1}$ . Then we go over the  $16(2c+3)$  elements between  $y'_i$  and  $y'_{i+1}$ , and compare each of them with  $x$ .  $\square$

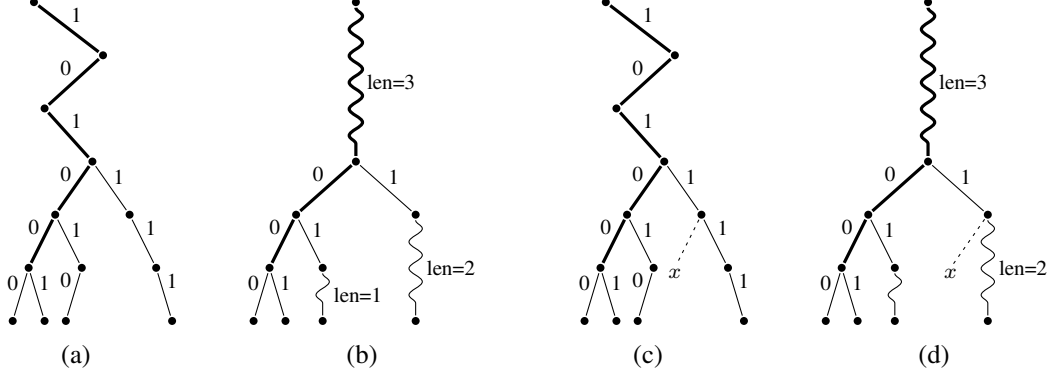


Figure 1: (b) is the *topological structure* of (a), by getting rid of the information that for each single child, whether it is a left or a right child. The thick edges are *shared*. Query  $x$  branches off the tree from the dotted edge.

Next, we show that if the sets are very small ( $m \leq O(\kappa/\lg \kappa)$ ), then there is a data structure with constant query time and negligible extra bits.

**Lemma 34.** *Let  $c \geq 2, \epsilon$  be two positive constants. There is a preprocessing algorithm `prepRankS`, a query algorithm `qAlgRankS` and lookup tables `tableRankSV,m` of sizes  $O(2^{\epsilon\kappa})$ , such that for any integers  $V \leq 2^\kappa$  and  $m \leq c \cdot \kappa/\lg \kappa$ , such that given a set  $S \subset [V]$  of size  $m$ , `prepRankS` preprocesses  $S$  into a data structure using  $\lg \binom{V}{m} + 2^{-\kappa/2}$  bits of space. Given any  $x \in [V]$ , `qAlgRankS` computes  $\text{rank}_S(x)$  in constant time by accessing the data structure and `tableRankSV,m`.*

*Proof.* Consider the binary trie over  $\{0, \dots, V\}$ .<sup>4</sup> Every element in  $\{0, \dots, V\}$  corresponds to a root-to-leaf path. Consider all paths corresponding to an element in  $S \cup \{V\}$  ( $V$  is included for technical reasons). Their union forms a subtree  $T(S)$  of the binary trie with  $m + 1$  leaves. In the following, we construct a data structure assuming the *topological structure* of  $T(S)$  is known, then apply Proposition 10 to fuse the topological structure into the data structure.

Roughly speaking, the *topological structure* of a subtree  $T$  is the tree  $T$  without specifying for each node with only one child, whether it is a left or a right child (see Figure 1a). Formally, it is defined by partitioning the set of such subtrees into equivalence classes, modulo the `flip` operation. Let  $v$  be a node in  $T$  with only a left [resp. right] child, let `flip`( $v, T$ ) be  $T$  relocating  $v$ 's entire left [resp. right] subtree to its right [resp. left] child. We say two trees  $T \sim T'$  if there is a (finite) sequence of `flip` operations that modifies  $T$  to  $T'$ . It is easy to verify that  $\sim$  is an equivalence relation, hence it partitions the set of all  $T$  into equivalence classes.

We call an edge in  $T(S)$  a *shared edge* if it has more than one leaf in its subtree. Equivalently, a shared edge is shared between at least two root-to-leaf paths. Note that if an edge is shared, then all edges on the path from root to it are shared. It turns out that the *number of shared edges* in  $T(S)$  is an important parameter, which is also invariant under `flip`. Thus, for each equivalence class  $\mathcal{T}$ , all  $T \in \mathcal{T}$  have the same number of shared edges (see Figure 1b).

Intuitively, for a typical set  $S$ , the corresponding  $\mathcal{T} \ni T(S)$  should have most of its degree-two nodes close to the root, i.e, it should have very *few* shared edges. Indeed, if we sample a uniformly random  $S$ , the

<sup>4</sup>We write every integer in the set as a  $\lceil \lg(V+1) \rceil$ -bit string, then construct a trie over these  $V+1$  binary strings. Note that  $S$  is a subset of  $\{0, \dots, V-1\}$ , while the trie has  $V+1$  leaves.

number of shared edges is at most  $O(\kappa)$  with probability at least  $1 - 2^{-\Omega(\kappa)}$ . As we will see below, on the inputs with few shared edges, it is relatively easy to construct data structures and answer queries. However, for the rare inputs with more than  $\Omega(\kappa)$  shared edges, we can afford to use a different construction with a larger redundancy. Since they are sufficiently rare, the overall redundancy turns out to be small.

**Few shared edges.** Let us fix an equivalence class  $\mathcal{T}$ , assume  $\mathcal{T}$  is known and consider all inputs  $S$  such that  $T(S) \in \mathcal{T}$ . Furthermore, assume the trees in  $\mathcal{T}$  have at most  $(2c + 1)\kappa$  shared edges. For each such  $\mathcal{T}$ , we construct a lookup table `tableRankSV,m,T`, and preprocess  $S$  into a data structure using about  $\lg |\mathcal{T}|$  bits such that if the query algorithm is given access to this particular lookup table (specific for  $\mathcal{T}$ ), it answers rank queries in constant time.

Since the tree  $T(S)$  uniquely determines  $S$ , to construct the data structure for  $S$ , it suffices to encode for each edge in  $T(S)$  that connects a single child and its parent, whether the child is left or right. The preprocessing algorithm constructs  $T(S)$ , then goes through all such edges in a *fixed* order, and uses one bit to indicate whether the corresponding edge in  $T(S)$  connects to a left child or a right child. To facilitate the queries (which we will describe in the next paragraph), all shared edges are encoded first in the *depth-first-search* order, followed by all other edges encoded in the *depth-first-search* order. This ensures that

1. if a shared edge  $e_1$  is on the path from root to shared edge  $e_2$ , then  $e_1$  is encoded before  $e_2$ ;
2. for each  $y_i$ , its non-shared edges (which is a suffix in the root-to-leaf path) are consecutive in the data structure.

Note that this encoding is a *one-to-one* mapping: Every  $S$  such that  $T(S) \in \mathcal{T}$  is encoded to a different string; Every string has a corresponding  $S$  with  $T(S) \in \mathcal{T}$  encoded to it. Thus, the algorithm constructs a data structure using exactly

$$\lg |\{S : T(S) \in \mathcal{T}\}|$$

bits of space.

Let  $S = \{y_1, \dots, y_m\}$  such that  $y_1 < y_2 < \dots < y_m$ , and let  $y_0 = -1$  and  $y_m = V$ . Given a query  $x \in \{0, \dots, V - 1\}$ , the goal is to compute  $i$  such that  $y_i \leq x < y_{i+1}$ . Let us consider the process of walking down the tree  $T(S)$  following the bits in  $x$ . That is, we write  $x$  also as a  $\lceil \lg(V + 1) \rceil$ -bit string, and walk down the tree from the root: if the current bit in  $x$  is 0, we follow the left child, otherwise we follow the right child. The process stops when either the current node in  $T(S)$  does not have a left (or right) child to follow, or we have reached a leaf. The location where it stops determines the answer to the query, in the same way for *all*  $T \in \mathcal{T}$ . See Figure 1c and 1d for a concrete example. Note that in the example,  $x$  branches off the tree from left, which may not be possible at the same location for all  $T \in \mathcal{T}$ , as some  $T$  may have a left child there. But *given* that  $x$  branches off the tree at this location from left, all  $T(S) \in \mathcal{T}$  must have the same answer to `rankS(x)`. Thus, we store in `tableRankSV,m,T`, for all nodes  $v$  in the tree, `ansv,0` and `ansv,1`, the answer to the query when the process branches off the tree from  $v$  due to the lack of its left child (i.e., from left), and the answer when it branches off from  $v$  due to the lack of its right child (i.e., from right) respectively. It takes  $O(\kappa^2)$  words, since  $m \leq \kappa$ .

Now the task is reduced to efficiently simulating this walk. To this end, the query algorithm needs to compare the bits in  $x$  with the corresponding bits of  $T(S)$ , which are stored in the data structure. It turns out that the difficult part is to compare  $x$  with the shared edges, which are stored in the first (at most)  $(2c + 1)\kappa$  bits. The first step is to simulate the walk, and check if  $x$  branches off  $T(S)$  at a shared edge. We create lookup tables of size  $2^{\epsilon\kappa}$  to compare  $\epsilon\kappa$  bits at once. For now, let us focus on the first  $\epsilon\kappa$  bits  $x_{\leq \epsilon\kappa}$ . These bits determine for all the degree-two nodes in the first  $\epsilon\kappa$  levels, which child  $x$  follows (note we have fixed  $\mathcal{T}$ ). Thus, it determines for all other bits, which bits in the data structure they should compare with. In the lookup table, we store for each of the  $2^{\epsilon\kappa}$  possible values,

- a  $(2c + 1)\kappa$ -bit string, which permutes  $x_{\leq \epsilon\kappa}$  to the same location as the bits they are comparing with;
- a  $(2c + 1)\kappa$ -bit string, indicating for each shared edge in the data structure, whether they are being compared.

With these two strings, the query algorithm is able to compare  $x_{\leq \epsilon\kappa}$  with the first  $\epsilon\kappa$  levels of  $T(S)$ . If they do not match, we could find the first edge where they differ (since edges are encoded in the DFS order), which is the location where  $x$  branches off  $T(S)$ . If they all equal, we proceed and compare the next  $\epsilon\kappa$  bits. Note that we may start the next chunk of the walk from different nodes depending on the value of  $x_{\leq \epsilon\kappa}$ , and we will need a different lookup for each starting location. However,  $\mathcal{T}$  can have at most  $m$  nodes in each level, thus, only  $m$  tables are needed for each chunk. We repeat the above process until we find a different bit, or we find out that  $x$  matches all shared edges from the root. In the former case, as we argued above, the answer to the query can be found in the lookup table. In the latter case, by the definition of shared edges, we identified one  $y_i$  which is the only element in  $S$  that matches the prefix of  $x$ . Thus, it suffices to retrieve the remaining bits of  $y_i$ , which are stored consecutively in the data structure and take constant retrieval time, and compare  $y_i$  with  $x$ . If  $y_i \leq x$ , then the query algorithm returns  $i$ , otherwise, it returns  $i - 1$ . The query time is constant.

So far for every  $\mathcal{T}$  with at most  $(2c + 1)\kappa$  shared edges, we have designed a data structure that works for all inputs  $S$  such that  $S \in \mathcal{T}$  using space  $\lg |\{S : T(S) \in \mathcal{T}\}|$  bits, constant query time and lookup table of size  $2^{\epsilon\kappa}$ . Next, we fuse  $\mathcal{T}$  into the data structure and merge all lookup tables, obtaining a single data structure that works for all  $S$  such that  $T(S)$  has at most  $(2c + 1)\kappa$  shared edge, which uses lookup table `tableRankSV,m,few`. To this end, we fix an arbitrary ordering of all such equivalence classes  $\mathcal{T}: \mathcal{T}_1, \dots, \mathcal{T}_C$ , where  $C$  is the number of equivalence classes. Let  $s_i = \lg |\{S : T(S) \in \mathcal{T}_i\}|$  be the size of the data structure for  $\mathcal{T}_i$ . Then,  $C \leq 2^{2m} \cdot \binom{(2c+1)\kappa+1}{m-1} \leq 2^{m \lg(\kappa/m) + O(m)}$ . This is because there are at most  $2^{2m}$  rooted binary trees with  $m + 1$  nodes (corresponding to the degree-two nodes). Each such tree can be extended to a class  $\mathcal{T}$  by specifying the distance from each child to its parent (adding the degree-one nodes). However, there are only  $(2c + 1)\kappa$  shared edges, thus, the sum of distances of all internal edges is at most  $(2c + 1)\kappa$ , and there are  $m - 1$  internal edges.<sup>5</sup> Hence, it is at most  $\binom{\leq (2c+1)\kappa}{m-2} \leq \binom{(2c+1)\kappa+1}{m-1}$  choices. Once the distances on all internal edges are determined, the distance on each edge connecting to a leaf is also fixed, because all leaves are at depth  $\lceil \lg(V + 1) \rceil$ .

Given an input set  $S$  such that  $T(S)$  has at most  $(2c + 1)\kappa$  shared edges, the preprocessing algorithm computes  $T(S)$  and finds the index  $i$  such that  $\mathcal{T}_i \ni T(S)$ . Then it runs the preprocessing algorithm for class  $\mathcal{T}_i$  on  $S$ , and computes a data structure  $\mathcal{D}_i$  of at most  $s_i$  bits. Next, we use Proposition 10 to store the pair  $(i, \mathcal{D}_i)$ , using space at most

$$\begin{aligned} \lg \sum_{i=1}^C 2^{s_i} + C \cdot 2^{-\kappa+2} &\leq \lg \left( \sum_{i=1}^C |\{S : T(S) \in \mathcal{T}_i\}| \right) + 2^{m \lg(\kappa/m) + O(m) - \kappa + 2} \\ &< \lg \binom{V}{m} + 2^{m \lg(\kappa/m) + O(m) - \kappa + 2} \\ &< \lg \binom{V}{m} + 2^{-\frac{3}{4}\kappa}. \end{aligned}$$

The lookup table `tableRankSV,m,few` is simply the concatenation of all tables `tableRankSV,m,Ti` for  $i = 1, \dots, C$ , as well as the  $O(C)$ -sized table from Proposition 10. Thus, the total size is at most  $2^{\epsilon\kappa} \cdot C + O(C) = 2^{(\epsilon+o(1))\kappa}$ .

---

<sup>5</sup>An edge is internal if it does not connect to a leaf.

To answer a query  $x$ , Proposition 10 allows us to decode  $i$  and  $\mathcal{D}_i$  in constant time by storing a lookup table of size  $O(C)$ . Then, we find the corresponding lookup table  $\text{tableRankS}_{V,m,\mathcal{T}_i}$  and run the query algorithm for  $\mathcal{T}_i$  on query  $x$  and data structure  $\mathcal{D}_i$ . The query time is constant.

**Many shared edges.** Next, we construct a data structure that works for all  $S$  such that  $T(S)$  has more than  $(2c+1)\kappa$  shared edges, using

$$\lg \binom{V}{m} - \kappa$$

bits of space. Note that this is possible, because there are very few such sets  $S$  (a tiny fraction of all  $\binom{V}{m}$  sets). We find the largest  $k$  such that  $T(S_{\leq k})$  has at most  $(2c+1)\kappa$  shared edges, where  $S_{\leq k} = \{y_1, \dots, y_k\}$ . Note that every element can introduce no more than  $\kappa$  shared edges, thus,  $T(S_{\leq k})$  has at least  $2c\kappa$  shared edges. The data structure stores the (index of) equivalence class  $\mathcal{T} \ni T(S_{\leq k})$ , then we run the preprocessing algorithm on  $S_{\leq k}$ . This encodes the first  $k$  elements of  $S$ . For the next  $m-k$  elements, we simply apply Lemma 33.

More specifically, for  $k$  elements, there are at most  $2^{k \lg(\kappa/k) + O(k)}$  equivalence classes, as we showed earlier. We construct the data structure as follows:

1. write down the index  $k$  using  $\lceil \lg m \rceil$  bits;
2. write down the index  $i$  such that  $\mathcal{T}_i \ni T(S_{\leq k})$  using  $\lceil k \lg(\kappa/k) + O(k) \rceil$  bits;
3. run the preprocessing algorithm on  $S_{\leq k}$  and obtain a data structure of size

$$\lg |\{S_{\leq k} : T(S_{\leq k}) \in \mathcal{T}_i\}|;$$

4. run  $\text{prepRankL}$  on  $\{y_{k+1}, \dots, y_m\}$  and obtain a data structure of size

$$\lg \binom{V}{m-k} + \frac{1}{8}(m-k) \lg \kappa.$$

Observe that Step 3 uses at most

$$k \lceil \lg V \rceil - 2c\kappa$$

bits, because for any such  $\mathcal{T}_i$ ,

- by construction, each bit of the data structure stores an input bit, i.e., one of the bits representing  $\{y_1, \dots, y_k\}$ ;
- each of the  $\geq 2c\kappa$  shared edges corresponds to at least two input bits (since given  $\mathcal{T}$ , these two input bits are always the same);
- each input bit is stored only once.

Therefore, the preprocessing algorithm outputs a data structure using

$$\begin{aligned} & \lg m + k \lg(\kappa/k) + O(k) + (k \lg V - 2c\kappa) + \left( \lg \binom{V}{m-k} + \frac{1}{8}(m-k) \lg \kappa \right) + k + 2 \\ & \leq \lg m + k \lg(\kappa/k) + (k \lg V - 2c\kappa) + (m-k) \lg V + \frac{1}{8}m \lg \kappa + O(k) \\ & \leq m \lg V - 2c\kappa + \lg m + k \lg(\kappa/k) + \frac{1}{8}m \lg \kappa + O(k) \end{aligned}$$

$$\begin{aligned}
&\leq \lg \binom{V}{m} + m \lg m - 2c\kappa + \lg m + m \lg(\kappa/m) + \frac{1}{8}m \lg \kappa + O(m) \\
&\leq \lg \binom{V}{m} - 2c\kappa + \frac{9}{8}m \lg \kappa + O(m).
\end{aligned}$$

By the fact that  $m \leq c\kappa / \lg \kappa$  and  $c \geq 2$ , it is at most

$$\lg \binom{V}{m} - \kappa.$$

The lookup table includes  $\text{tableRankS}_{V,k,\text{few}}$  for all  $k \leq m$ , and has  $2^{(\epsilon+o(1))\kappa}$  size.

To answer query  $x$ , the query algorithm reads  $k$  and  $i$ . Then it runs the query algorithm for  $\mathcal{T}_i$  for query  $x$  on the data structure for  $S_{\leq k}$ , as well as  $\text{qAlgRankL}$  for  $x$  on the data structure for  $\{y_{k+1}, \dots, y_m\}$ . Both algorithms run in constant time. The answer to the query is simply the sum of the two answers.

**Combining the two cases.** Finally, we combine the two cases using Proposition 10, and construct a data structure that works for all sets  $S$ . Given set  $S$ ,  $\text{prepRankS}$  computes  $T(S)$  and the number of shared edges. If it has no more than  $(2c+1)\kappa$  shared edges, it sets  $b := 1$ , runs the preprocessing algorithm for “many shared edges” and obtains a data structure  $\mathcal{D}_1$ . Otherwise, it sets  $b := 2$ , runs the preprocessing algorithm for “few shared edges” and obtains a data structure  $\mathcal{D}_2$ . At last, it applies Proposition 10 to store the pair  $(b, \mathcal{D}_b)$ . The space usage is

$$\begin{aligned}
&\lg \left( \binom{V}{m} \cdot 2^{2^{-\frac{3}{4}\kappa}} + \binom{V}{m} \cdot 2^{-\kappa} \right) + 2^{-\kappa+2} \\
&\leq \lg \binom{V}{m} + 2^{-\frac{3}{4}\kappa} + \lg(1 + 2^{-\kappa-2^{-\frac{3}{4}\kappa}}) + 2^{-\kappa+2} \\
&\leq \lg \binom{V}{m} + 2^{-\frac{1}{2}\kappa}.
\end{aligned}$$

To answer query  $x$ , we simply decode  $b$  and  $\mathcal{D}_b$  using Proposition 10, and use the corresponding query algorithm based on  $b$ .

The lookup table  $\text{tableRankS}_{V,m}$  also includes all  $\text{tableRankS}_{V,k}$  for  $k \leq m$ , which has size  $2^{O(\epsilon\kappa)}$ . This proves the lemma.  $\square$

Finally, we prove Lemma 28, which constructs a rank data structure for  $m \leq \kappa^c$ .

**Lemma 28 (restated).** *Let  $c$  be any constant positive integer and  $\epsilon$  be any positive constant. There is a preprocessing algorithm  $\text{prepRank}$ , query algorithm  $\text{qAlgRank}$  and lookup tables  $\text{tableRank}_{V,m}$  of sizes  $\tilde{O}(2^{\epsilon\kappa})$ , such that for any integers  $V \leq 2^{\kappa/2}$ ,  $m \leq \kappa^c$ , given a set  $S \subset [V]$  of size  $m$ ,  $\text{prepRank}(V, m, S)$  outputs a data structure  $\mathcal{D}$  of length*

$$\lg \binom{V}{m} + (m-1) \cdot 2^{-\kappa/2}.$$

*Given  $x \in [V]$ ,  $\text{qAlgRank}(V, m, x)$  computes  $\text{rank}_S(x)$  in constant time, by accessing  $\mathcal{D}$  and  $\text{tableRank}_{V,m}$ . In particular, by computing both  $\text{rank}_S(x)$  and  $\text{rank}_S(x-1)$ , one can decide if  $x \in S$  in constant time. The algorithms run on a random access machine with word-size  $w = \Theta(\kappa)$ .*

*Proof.* The data structure construction is based on recursion. As the base case, if  $m \leq 16\kappa/\lg \kappa$ , we simply use the data structure from Lemma 34, and the statement holds. Otherwise for  $m > 16\kappa/\lg \kappa$ , we divide  $V$  into  $B$  blocks of equal size, for  $B = \lceil \kappa^{1/2} \rceil$ . For a typical set  $S$ , we would expect each block to contain roughly  $m/B$  elements. If it indeed happens, the size of  $S$  would be reduced by a factor of  $B$ . Hence, we will reach the base case in constant rounds. On the other hand, input sets  $S$  which have at least one block with significantly more than  $m/B$  elements are very rare. If such blocks occur, we are going to apply Lemma 33 on them. Although Lemma 33 introduces a large redundancy, such cases occur sufficiently rarely, so that the overall redundancy is still small.

We partition the input set  $S$  into  $B$  subsets  $S_1, \dots, S_B$  such that  $S_i$  contains all elements of  $S$  between  $\lceil (i-1)V/B \rceil$  and  $\lceil iV/B \rceil - 1$ . Let  $V_i := \lceil iV/B \rceil - \lceil (i-1)V/B \rceil$  be the size of the  $i$ -th block. By definition,  $|S_1| + \dots + |S_B| = m$  and  $V_1 + \dots + V_B = V$ . We construct a data structure for each  $S_i$ , over a universe of size  $V_i$ . Then we apply Proposition 8 to concatenate the  $B$  data structures *given* the sizes of  $S_1, \dots, S_B$ . Finally, we apply Proposition 10 to union all combinations of sizes. We present the details below.

**Preprocessing algorithm.** Given a set  $S$  of size  $m$ , if  $2m \geq V$ , we take the complement. Note that the space bound stated in the lemma becomes smaller after taking the complement. It is also easy to derive the answer from the data structure for the complement. Then if  $m = 1$ , we simply write down the element; if  $m \leq 16\kappa/\lg \kappa$ , we apply Lemma 34.

**preprocessing algorithm**  $\text{prepRank}(V, m, S)$ :

1. if  $V \leq 2m$
2.    $m := V - m$  and  $S := [V] \setminus S$
3. if  $m = 1$
4.   return the only element in  $S$
5. if  $m \leq 16\kappa/\lg \kappa$
6.   return  $\mathcal{D} := \text{prepRankS}(V, m, S)$  using Lemma 34 (to be cont'd)

If  $m > 16\kappa/\lg \kappa$ , we divide  $[V]$  into  $\kappa^{1/4}$  chunks, and construct a data structure for each chunk.

7.  $B := \lfloor \kappa^{1/4} \rfloor$
8. compute  $S_i := S \cap [(i-1)V/B, iV/B)$  and  $m_i := |S_i|$
9. let  $V_i := \lceil iV/B \rceil - \lceil (i-1)V/B \rceil$
10. for  $i = 1, \dots, B$
11.   if  $m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}$
12.     compute  $\mathcal{D}_i := \text{prepRankL}(V_i, m_i, S_i)$  using Lemma 33
13.   else
14.     compute  $\mathcal{D}_i := \text{prepRank}(V_i, m_i, S_i)$  recursively

If the chunk has too many elements, we apply Lemma 33 to construct a data structure with larger redundancy. Otherwise, the size of the set at least decreases by a factor of  $\kappa^{1/4}$ , and we recurse.

Next, we concatenate the data structures for all chunks, and fuse the tuple  $(m_1, \dots, m_B)$  into the data structure.

15. apply Proposition 8 to concatenate  $\mathcal{D}_1, \dots, \mathcal{D}_B$ , and obtain  $\mathcal{D}_{\text{cat}}$
16. let  $C := \binom{m+B-1}{B-1}$  be the number of different tuples  $(m_1, \dots, m_B)$  such that  $m_i \geq 0$  and  $m_1 + \dots + m_B = m$
17. let  $1 \leq j \leq C$  be the index such that the current  $(m_1, \dots, m_B)$  is the  $j$ -th in the lexicographic order
18. apply Proposition 10 to fuse  $j$  into  $\mathcal{D}_{\text{cat}}$ , and obtain  $\mathcal{D}$
19. return  $\mathcal{D}$

**Space analysis.** In the following, we analyze the size of the data structure. We will prove by induction that  $\text{prepRank}(V, m, S)$  outputs a data structure of size at most

$$\lg \binom{V}{m} + (m-1)2^{-\kappa/2}.$$

The base case when  $m \leq 16\kappa/\lg \kappa$  is a direct implication of Lemma 34 (or if  $m = 1$ , the space usage if  $\lg V = \lg \binom{V}{1}$ ). Now, let us consider larger  $m$ .

To prove the inductive step, let us fix a  $B$ -tuple  $(m_1, \dots, m_B)$ , and consider the size of  $\mathcal{D}_{\text{cat}}$  from line 15. By Proposition 8, when all  $m_i \leq \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}$ , its size is at most

$$s(m_1, \dots, m_B) := \lg \prod_{i=1}^B \binom{V_i}{m_i} + (m-B) \cdot 2^{-\kappa/2} + (B-1)2^{-\kappa+4};$$

otherwise, its size is at most

$$s(m_1, \dots, m_B) := \lg \prod_{i=1}^B \binom{V_i}{m_i} + \sum_{i: m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}} \frac{1}{8} m_i \lg \kappa + B. \quad (7)$$

It turns out that in the latter case, (7) is *significantly* smaller than  $\lg \binom{V}{m}$ .

**Claim 35.** *If there is at least one  $m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}$ , then (7) is at most  $\lg \binom{V}{m} - \kappa$ .*

We defer its proof to the end. Assuming the claim, by Proposition 10, the size of  $\mathcal{D}$  from line 18 is at most

$$\lg \left( \sum_{\substack{m_1, \dots, m_B: \\ \sum_i m_i = m}} 2^{s(m_1, \dots, m_B)} \right) + C \cdot 2^{-\kappa+2}. \quad (8)$$

To bound the sum in the logarithm, we first take the sum only over all tuples such that  $m_i \leq \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}$ , the sum is at most

$$\begin{aligned} \sum 2^{s(m_1, \dots, m_B)} &\leq \sum \prod_{i=1}^B \binom{V_i}{m_i} \cdot 2^{(m-B) \cdot 2^{-\kappa/2} + (B-1)2^{-\kappa+4}} \\ &\leq \binom{V}{m} \cdot 2^{(m-B) \cdot 2^{-\kappa/2} + (B-1)2^{-\kappa+4}}, \end{aligned}$$

where the second inequality uses the fact that  $\sum_{m_1, \dots, m_B: \sum m_i = m} \prod_{i=1}^B \binom{V_i}{m_i} \leq \binom{\sum_{i=1}^B V_i}{m}$ , and we are taking this sum over a subset of all such  $B$ -tuples. By Claim 35,  $s(m_1, \dots, m_B) \leq \lg \binom{V}{m} - \kappa$  for all other tuples. Thus, the sum in the logarithm is at most

$$\binom{V}{m} \cdot 2^{(m-B) \cdot 2^{-\kappa/2} + (B-1)2^{-\kappa+4}} + \binom{V}{m} \cdot C \cdot 2^{-\kappa}.$$

Finally, since  $C \leq m^B$  and  $m \leq \kappa^c$ , (8) is at most

$$(8) \leq \lg \left( \binom{V}{m} \cdot 2^{(m-B) \cdot 2^{-\kappa/2} + (B-1)2^{-\kappa+4}} + \binom{V}{m} \cdot m^B \cdot 2^{-\kappa} \right) + m^B \cdot 2^{-\kappa+2}$$

$$\begin{aligned}
&\leq \lg \binom{V}{m} + (m - B)2^{-\kappa/2} + (B - 1)2^{-\kappa+4} + \lg(1 + 2^{-\kappa+B \lg m}) + 2^{-\kappa+B \lg m+2} \\
&\leq \lg \binom{V}{m} + (m - B)2^{-\kappa/2} + (B - 1)2^{-\kappa+4} + 2^{-\kappa+c\kappa^{1/4} \lg \kappa+3} \\
&\leq \lg \binom{V}{m} + (m - 1)2^{-\kappa/2}.
\end{aligned}$$

By induction, it proves the data structure uses space as claimed.

**Lookup table.** We store the following information in the lookup table.

<b>lookup table</b> <code>tableRank<sub>V,m</sub></code> : 1. if $m \leq 16\kappa/\lg \kappa$ , include <code>tableRankS<sub>V,m</sub></code> from Lemma 34 2. the value of $C = \binom{m+B-1}{B-1}$ 3. for all $1 \leq j \leq C$ 4. the $j$ -th $B$ -tuple $(m_1, \dots, m_B)$ in the lexicographic order 5. for $i = 1, \dots, B$ 6. $m_1 + \dots + m_i$ 7. lookup table for Proposition 8 in line 15, for all possible $B$ -tuples $(m_1, \dots, m_B)$ 8. lookup table for Proposition 10 in line 18 9. include all tables <code>tableRank<sub>V',m'</sub></code> and <code>tableRankL<sub>V',m'</sub></code> for $V' = \lfloor V/B^i \rfloor$ or $\lceil V/B^i \rceil$ for $i \geq 1$ , and $m' \leq m$
--

Since  $C = \binom{m+B-1}{B-1} \leq 2^{o(\kappa)}$ , line 2 to 8 all have size  $2^{o(\kappa)}$ . Finally, we are only including  $\kappa^{O(1)}$  other tables in line 1 and 6, each taking at most  $\tilde{O}(2^{\epsilon\kappa})$  bits by Lemma 33 and 34. The total size of `tableRankV,m` is  $\tilde{O}(2^{\epsilon\kappa})$ .

**Query algorithm.** Given a query  $x$ , if  $V \leq 2m$ , we retreat the data structures as storing the complement of  $S$ , and use the fact that  $\text{rank}_S(x) = x + 1 - \text{rank}_{[V] \setminus S}(x)$ . Then if  $m = 1$ , we simply compare it with  $x$ . If  $m \leq 16\kappa/\lg \kappa$ , we invoke the query algorithm from Lemma 34.

<b>query algorithm</b> <code>qAlgRank(V, m, x)</code> : 1. if $V \leq 2m$ 2. $m := V - m$ 3. in the following, when about to return answer $r$ , return $x + 1 - r$ 4. if $m = 1$ 5. retrieve the element, compare it with $x$ , and return 0 or 1 6. if $m \leq 16\kappa/\lg \kappa$ , 7. return <code>qAlgRankS(V, m, x)</code> (from Lemma 34) (to be cont'd)
---

If  $m > 16\kappa/\lg \kappa$ , we decode  $j$ , which encodes the tuple  $(m_1, \dots, m_B)$  and  $\mathcal{D}_{\text{cat}}$ . Then if  $x$  is in the  $i$ -th chunk, we decode  $m_i$  and the corresponding  $\mathcal{D}_i$ .

8. apply Proposition 10 to decode $j$ and $\mathcal{D}_{\text{cat}}$ 9. let $i$ be the chunk that contains $x$ 10. apply Proposition 8 to decode $\mathcal{D}_i$ 11. retrieve $m_1 + \dots + m_{i-1}$ and $m_i$ for $j$ -th tuple from the lookup table (to be cont'd)
---

Then depending on the value of  $m_i$ , we invoke the query algorithm from Lemma 33 or recurse.

12. if  $m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}$
13.     return  $(m_1 + \dots + m_{i-1}) + \mathcal{D}_i.\text{qAlgRankL}(V_i, m_i, x - \lceil (i-1)V/B \rceil)$  (from Lemma 33)
14. else
15.     return  $(m_1 + \dots + m_{i-1}) + \mathcal{D}_i.\text{qAlgRank}(V_i, m_i, x - \lceil (i-1)V/B \rceil)$

The query algorithm recurses only when  $m_i \leq m \cdot \kappa^{-1/4}$ . In all other cases, the query is answered in constant time. On the other hand,  $m \leq \kappa^c$ . The level of recursion must be bounded by a constant. Thus, the data structure has constant query time, proving the lemma.  $\square$

Next, we prove the remaining claim.

*Proof of Claim 35.* To prove the claim, let us first compare the first term with  $\lg \binom{V}{m}$ . We have

$$\begin{aligned} & \lg \binom{V}{m} - \lg \prod_{i=1}^B \binom{V_i}{m_i} \\ &= \lg \frac{V! \cdot m_1! \cdots m_B! \cdot (V_1 - m_1)! \cdots (V_B - m_B)!}{V_1! \cdots V_B! \cdot m!(V - m)!}, \end{aligned}$$

which, by Stirling's formula, is at least

$$\geq \sum_{i=1}^B \left( V_i \lg \frac{V}{V_i} - m_i \lg \frac{m}{m_i} - (V_i - m_i) \lg \frac{V - m}{V_i - m_i} \right) - O(B) - \lg V,$$

which by the fact that  $f(\varepsilon) = \varepsilon \log 1/\varepsilon$  is concave and hence  $V \cdot f(\frac{V_i}{V}) \geq m \cdot f(\frac{m_i}{m}) + (V - m) \cdot f(\frac{V_i - m_i}{V - m})$ , is at least

$$\geq \sum_{i: m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}} \left( V_i \lg \frac{V}{V_i} - m_i \lg \frac{m}{m_i} - (V_i - m_i) \lg \frac{V - m}{V_i - m_i} \right) - O(B) - \lg V. \quad (9)$$

For each term in this sum, we have

$$V_i \lg \frac{V}{V_i} = V_i \lg B - V_i \lg \left( 1 + \frac{V_i - V/B}{V/B} \right) \geq V_i \lg B - O(1),$$

since  $|V_i - V/B| \leq 1$ ; and

$$\begin{aligned} (V_i - m_i) \lg \frac{V - m}{V_i - m_i} &= (V_i - m_i) \left( \lg B + \lg \left( 1 + \frac{m_i - m/B + (V/B - V_i)}{V_i - m_i} \right) \right) \\ &\leq (V_i - m_i) \lg B + (V_i - m_i) \cdot \frac{m_i - m/B + 1}{V_i - m_i} \cdot \lg e \\ &\leq (V_i - m_i) \lg B + 2m_i. \end{aligned}$$

Plugging into (9), we have

$$\lg \binom{V}{m} - \lg \prod_{i=1}^B \binom{V_i}{m_i}$$

$$\begin{aligned}
&\geq \sum_{i:m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}} \left( V_i \lg B - m_i \lg \frac{m}{m_i} - (V_i - m_i) \lg B - 2m_i \right) - O(B) - \lg V \\
&= \sum_{i:m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}} m_i \left( \lg \frac{Bm_i}{m} - 2 \right) - O(B) - \lg V \\
&\geq \sum_{i:m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}} m_i \left( \frac{1}{4} \lg \kappa - 2 \right) - O(B) - \lg V.
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
(7) &\leq \lg \binom{V}{m} - \sum_{i:m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}} m_i \left( \frac{1}{4} \lg \kappa - 2 \right) + O(B) + \lg V \\
&\quad + \sum_{i:m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}} \frac{1}{8} m_i \lg \kappa \\
&\leq \lg \binom{V}{m} - \sum_{i:m_i > \max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}} m_i \left( \frac{1}{8} \lg \kappa - 2 \right) + O(B) + \lg V \\
&\leq \lg \binom{V}{m} - \kappa.
\end{aligned}$$

The last inequality is due to the fact that there is at least one  $m_i$  that is larger than  $\max\{m \cdot \kappa^{-1/4}, 16\kappa/\lg \kappa\}$  (in particular,  $m_i > 16\kappa/\lg \kappa$ ),  $B = \Theta(\kappa^{1/2})$  and  $\lg V \leq \kappa/2$ .  $\square$

rank queries can be viewed as mapping that maps  $S \rightarrow [m]$ , and  $[V] \setminus S \rightarrow [V - m]$ . Thus, Lemma 25 is an immediate corollary.

**Lemma 25 (restated).** *Let  $c$  be any constant positive integer and  $\epsilon$  be any positive constant. There is a preprocessing algorithm `perfHashS`, query algorithm `qalgS` and lookup tables `tableSV,m` of sizes  $\tilde{O}(2^{\epsilon\kappa})$ , such that for any  $V \leq 2^{\kappa/2}$  and  $m \leq \kappa^c$ , given a set  $S \subset [V]$  of  $m$  keys, `perfHashS` preprocesses  $S$  into a data structure of size at most*

$$\text{OPT}_{V,m} + (m - 1) \cdot 2^{-\kappa/2+1},$$

*such that it defines a bijection  $h$  between  $S$  and  $[m]$  and a bijection  $\bar{h}$  between  $[V] \setminus S$  and  $[V - m]$ . Given any  $x \in [V]$ , `qalgS` answers `hash(x)` in constant time, by accessing the data structure and `tableSV,m`.*

## 8 Perfect Hashing for Sets of Any Size

In this section, we generalize the data structure from Section 6 to sets of all sizes, proving our main theorem.

**Theorem 36 (main theorem).** *For any constant  $\epsilon > 0$ , there is a preprocessing algorithm `perfHash`, a query algorithm `qAlg` and lookup tables `tableU,n` of size  $n^\epsilon$ , such that given*

- *a set  $S$  of  $n$  keys over the key space  $[U]$ ,*
- *a uniformly random string  $\mathcal{R}$  of length  $O(\lg^{12} n)$ ,*

*`perfHash` preprocesses  $S$  into a data structure  $\mathcal{D}$  of (worst-case) length*

$$\text{OPT}_{U,n} + O(\lg \lg U),$$

such that  $\mathcal{D}$  defines a bijection  $h$  between  $S$  and  $[n]$  and a bijection  $\bar{h}$  between  $[U] \setminus S$  and  $[U - n]$ . Given access to  $\mathcal{D}$ ,  $\mathcal{R}$  and  $\text{table}_{U,n}$ , for any key  $x \in [U]$ ,  $\text{qAlg}(U, n, x)$  outputs  $\text{hash}(x)$  on a RAM with word-size  $w \geq \Omega(\lg U)$ , in time

- $O(1)$  with probability  $1 - O(\lg^{-7} U)$  and
- $O(\lg^7 U)$  in worst-case,

where the probability is taken over the random  $\mathcal{R}$ . In particular, the query time is constant in expectation and with high probability.

*Proof.* Similar to the proof of Theorem 18, we first assume  $2n \leq U$  (otherwise, we take the complement of  $S$ ). Then if  $n \geq U^{1/12}$ , Theorem 18 already gives the desired result. From now on, we assume  $n < U^{1/12}$ .

We partition  $[U]$  into  $n^{12}$  blocks. A typical set  $S$  has all the keys in different blocks. In this case, we may view the universe size being only  $n^{12}$ , and apply Theorem 18. On the other hand, only roughly  $1/n^{11}$ -fraction of the inputs have at least one pair of keys in the same block, which suggests that we may use at least  $10 \lg n$  extra bits.

**No collision in blocks and last block empty.** More specifically, given  $S$  such that  $n = |S| < U^{1/12}$ , let  $V := \lceil U \cdot n^{-12} \rceil$  be the block size. We partition  $U$  into blocks:  $(U \div V)$  blocks of size  $V$  and one last block of size  $(U \bmod V)$ . Let us first only consider inputs that have at most one key in every block and no key in the last block. We apply Theorem 18 on the universe of all blocks. That is, let the universe size  $U_{\text{new}} = U \div V$ , number of keys  $n_{\text{new}} = n$ . We construct the new set  $S_{\text{new}}$  such that  $i \in S_{\text{new}}$  if and only if block  $i$  contains a key  $x \in S$ . By Theorem 18, we construct a data structure of size

$$\lg \binom{U_{\text{new}}}{n_{\text{new}}} + 1/U_{\text{new}} = \lg \binom{U \div V}{n} + 1/(U \div V),$$

which defines hash functions  $h_{\text{new}}$  and  $\bar{h}_{\text{new}}$ . Besides this data structure, we also apply Lemma 16 to store for each  $i \in S_{\text{new}}$ , the key  $x$  within block  $i$ , according to  $h_{\text{new}}(i)$ . That is, we store  $x - (i - 1)V$  in coordinate  $h_{\text{new}}(i)$ . Hence, this part takes

$$n \lg V + (n - 1)2^{-\kappa+5}$$

bits. Then we apply Proposition 8 to concatenate the two data structures. The total space is at most

$$\begin{aligned} & \lg \binom{U \div V}{n} + 1/(U \div V) + n \lg V + n \cdot 2^{-\kappa+5} \\ & \leq \lg \frac{V^n \prod_{i=0}^{n-1} (U \div V - i)}{n!} + O(1/n) \\ & \leq \lg \frac{\prod_{i=0}^{n-1} (U - i)}{n!} + O(1/n) \\ & = \lg \binom{U}{n} + O(1/n). \end{aligned}$$

To define the hash functions in this case, for each  $x \in S$ , which is in block  $i$ , we simply let  $h(x) := h_{\text{new}}(i)$ , the hash value of the block. For  $x \notin S$  in block  $i$ ,

- if  $i \in S_{\text{new}}$ , let  $x^*$  be the key in block  $i$ ,
  - if  $x < x^*$ , we let  $\bar{h}(x) := (V - 1) \cdot h_{\text{new}}(i) + (x - (i - 1)V)$ ,
  - if  $x > x^*$ , we let  $\bar{h}(x) := (V - 1) \cdot h_{\text{new}}(i) + (x - (i - 1)V - 1)$ ,
- if  $i \notin S_{\text{new}}$ , we let  $\bar{h}(x) := (V - 1)n + V \cdot \bar{h}_{\text{new}}(i) + (x - (i - 1)V)$ .

- if  $x$  is in the last block, we let  $\bar{h}(x) := (U \operatorname{div} V) \cdot V - n + (x - (U \operatorname{div} V) \cdot V)$

That is, we order all non-keys in block  $i$  for  $i \in S_{\text{new}}$  first, in the increase order of  $(h_{\text{new}}(i), x)$ ; then we order all non-keys not in the last block, in the increasing order of  $(\bar{h}_{\text{new}}(i), x)$ ; finally we order all non-keys in the last block.

To answer a query  $x$  in block  $i$ , we first query if  $i \in S_{\text{new}}$ . If  $i \notin S_{\text{new}}$ , then we know  $x$  is not a key, calculate  $\bar{h}(x)$  by its definition, and return. Otherwise, we query the  $(h_{\text{new}}(i) + 1)$ -th value in the second data structure, using Lemma 16, to retrieve the key in block  $i$ . If  $x$  happens to be this key, we return  $(1, h_{\text{new}}(i))$ . Otherwise,  $x \notin S$ , and  $\bar{h}(x)$  can be calculated by its definition. Finally, for queries  $x$  in the last block,  $x$  is not a key, and we calculate  $\bar{h}(x)$  according to its definition.

**Exist collision in blocks or keys in last block.** Next, we consider the case where at least one block contains more than one key, or the last block contains at least one key. We spend the first  $3\lceil \lg n \rceil$  bits to store

- $N$ , the number of blocks with at least two keys (blocks with collisions, or simply *collision blocks*),
- $n_{\text{cl}}$ , the total number of keys in all collision blocks,
- $n_{\text{last}}$ , the number of keys in the last block.

Next, we apply Lemma 26, and construct a membership data structure for  $N$  collision blocks using

$$\lg \binom{U \operatorname{div} V}{N} + O(N)$$

bits, which defines a bijection  $h_{\text{cl}}$  between all collision blocks and  $[N]$ , and a bijection  $\bar{h}_{\text{cl}}$  between all other blocks (except for the last block) and  $[(U \operatorname{div} V) - N]$ .

The final data structure has three *more* components:

1. store all keys in  $N$  collision blocks using Lemma 26, where each element  $x$  in block  $i$  is stored as  $V \cdot h_{\text{cl}}(i) + (x - (i - 1)V)$ , which uses at most

$$\lg \binom{NV}{n_{\text{cl}}} + O(n_{\text{cl}} + \lg \lg V)$$

bits;

2. store all other  $(U \operatorname{div} V) - N$  blocks using the data structure for no collisions, where each element  $x$  in block  $i$  is stored as  $V \cdot \bar{h}_{\text{cl}}(i) + (x - (i - 1)V)$ , which uses at most

$$\lg \binom{(U \operatorname{div} V)V - NV}{n - n_{\text{cl}} - n_{\text{last}}} + 1$$

bits;

3. store the last block using Lemma 26, which uses

$$\lg \binom{U \operatorname{mod} V}{n_{\text{last}}} + O(n_{\text{last}} + \lg \lg V)$$

bits.

Summing up the sizes of these three data structures, we get

$$\lg \binom{NV}{n_{\text{cl}}} \binom{(U \operatorname{div} V)V - NV}{n - n_{\text{cl}} - n_{\text{last}}} \binom{U \operatorname{mod} V}{n_{\text{last}}} + O(n_{\text{cl}} + n_{\text{last}} + \lg \lg V).$$

By the fact that  $\binom{n}{k} \leq (en/k)^k$  and  $n_{\text{cl}} \geq 2N$ , the first term is at most

$$\begin{aligned} & n_{\text{cl}} \lg \frac{NV}{n_{\text{cl}}} + (n - n_{\text{cl}} - n_{\text{last}}) \lg \frac{(U \operatorname{div} V)V - NV}{n - n_{\text{cl}} - n_{\text{last}}} + n_{\text{last}} \lg \frac{U \operatorname{mod} V}{n_{\text{last}}} + n \lg e \\ & \leq n_{\text{cl}} \lg \frac{V}{2} + (n - n_{\text{cl}} - n_{\text{last}}) \lg \frac{U}{n - n_{\text{cl}} - n_{\text{last}}} + n_{\text{last}} \lg V + n \lg e \\ & \leq n \lg \frac{eU}{n} + n_{\text{cl}} \lg \frac{nV}{2U} + (n - n_{\text{cl}} - n_{\text{last}}) \lg \frac{n}{n - n_{\text{cl}} - n_{\text{last}}} + n_{\text{last}} \lg \frac{nV}{U} \end{aligned}$$

which by the fact that  $V \leq 2U \cdot n^{-12}$ , is at most

$$\begin{aligned} & \leq n \lg \frac{eU}{n} + n_{\text{cl}} \lg n^{-11} + (n - n_{\text{cl}} - n_{\text{last}}) \lg \left( 1 + \frac{n_{\text{cl}} + n_{\text{last}}}{n - n_{\text{cl}} - n_{\text{last}}} \right) + n_{\text{last}} \lg(2n^{-11}) \\ & \leq n \lg \frac{eU}{n} + n_{\text{cl}} \lg n^{-11} + (n_{\text{cl}} + n_{\text{last}}) \lg e + n_{\text{last}} \lg(2n^{-11}) \\ & \leq n \lg \frac{eU}{n} - (n_{\text{cl}} + n_{\text{last}})(11 \lg n - O(1)). \end{aligned}$$

On the other hand, by Stirling's formula,

$$\begin{aligned} \lg \binom{U}{n} &= \lg \frac{U!}{n!(U-n)!} \\ &\geq \lg \frac{\sqrt{U} U^U}{\sqrt{n} n^n \cdot \sqrt{U-n} (U-n)^{U-n}} - O(1) \\ &\geq n \lg \frac{U}{n} + (U-n) \lg \frac{U}{U-n} - \frac{1}{2} \lg n - O(1) \end{aligned}$$

which by the fact that  $\ln(1+x) \geq x - x^2/2$  for  $x \geq 0$ , is at most

$$\begin{aligned} & \geq n \lg \frac{U}{n} + (U-n) \left( \frac{n}{U-n} - \frac{n^2}{2(U-n)^2} \right) \lg e - \frac{1}{2} \lg n - O(1) \\ &= n \lg \frac{eU}{n} - \frac{n^2 \lg e}{2(U-n)} - \frac{1}{2} \lg n - O(1) \\ &\geq n \lg \frac{eU}{n} - \frac{1}{2} \lg n - O(1). \end{aligned}$$

Thus, the total size of the data structure when  $N \geq 1$  is at most

$$\begin{aligned} & \lg \binom{U}{n} + \frac{1}{2} \lg n - (n_{\text{cl}} + n_{\text{last}})(11 \lg n - O(1)) + 3 \lg n + \lg \binom{U \operatorname{div} V}{N} + O(N + \lg \lg U) \\ & \leq \lg \binom{U}{n} + \frac{1}{2} \lg n - (n_{\text{cl}} + n_{\text{last}})(11 \lg n - O(1)) + 3 \lg n + 12N \lg n + O(N + \lg \lg U) \end{aligned}$$

which by the fact that  $n_{\text{cl}} \geq 2N$ , is at most

$$\begin{aligned} &\leq \lg \binom{U}{n} + \frac{1}{2} \lg n - (n_{\text{cl}} + n_{\text{last}})(5 \lg n - O(1)) + 3 \lg n + O(\lg \lg U) \\ &= \text{OPT}_{U,n} - \lg n + O(\lg \lg U). \end{aligned}$$

In this case, the hash functions are defined as follows. For both  $h$  and  $\bar{h}$ , we first order all elements in the  $N$  collision blocks according to their hash values from component 1, which are mapped to  $[n_{\text{cl}}]$  and  $[N \cdot V - n_{\text{cl}}]$  respectively. Then we order all elements in the  $(U \div V) - N$  non-collision blocks according to their hash values from component 2, which are mapped to  $\{n_{\text{cl}}, \dots, n - n_{\text{last}} - 1\}$  and  $\{N \cdot V - n_{\text{cl}}, \dots, (U \div V) \cdot V - (n - n_{\text{last}}) - 1\}$  respectively. Finally, we order all elements in the last block according to their hash values from component 3, which are mapped to  $\{n - n_{\text{last}}, n - 1\}$  and  $\{(U \div V) \cdot V - (n - n_{\text{last}}), U - n - 1\}$  respectively.

To answer a query  $x$  in block  $i$ , we retrieve  $N$ ,  $n_{\text{cl}}$  and  $n_{\text{last}}$ , and query if  $i$  is a collision block and  $h(i)$  (or  $\bar{h}(i)$ ). If  $i$  is a collision block, we query component 1; if  $i$  is not a collision block, we query component 2; if  $i$  is the last block, we query component 3. In any case, the hash value of  $x$  can be computed according to its definition in constant time.

Finally, we apply Proposition 10 to combine the two cases, by fusing a bit indicating whether there is any collision block. The final data structure has space bounded by

$$\begin{aligned} &\lg \left( 2^{\text{OPT}_{U,n} + O(1/n)} + 2^{\text{OPT}_{U,n} - \lg n + O(\lg \lg U)} \right) + 2^{-\kappa+2} \\ &= \text{OPT}_{U,n} + \lg(2^{O(1/n)} + (\lg^{O(1)} U)/n) + 2^{-\kappa+2} \\ &\leq \text{OPT}_{U,n} + O(\lg \lg U). \end{aligned}$$

The query algorithm is straightforward. To answer a query  $x$ , we apply Proposition 10 to decode the data structure, and the bit indicating whether there is any collision block or any element in the last block. Then we apply the corresponding query algorithm as described above. This proves the theorem.  $\square$

*Remark.* When the  $O(\lg \lg U)$  term is at most  $0.5 \lg n$ , the above data structure uses  $\text{OPT} + o(1)$  bits. To improve the  $O(\lg \lg U)$  term when  $U$  is large, we partition the universe into  $\lg^{10} U$  blocks, and check if any block has at least two keys. In this case, the fraction of inputs with some block with at least two keys is only  $1/\lg^{10} U$  fraction. Therefore, we can afford to “waste” about  $10 \lg \lg U$  bits, which dominates the  $O(\lg \lg U)$  term. This strategy reduces the problem to storing  $n$  non-empty blocks among a total of  $\lg^{O(1)} U$  blocks, i.e., the universe size is reduced from  $U$  to  $\lg^{O(1)} U$ . Thus, repeatedly applying it improves the  $O(\lg \lg U)$  term to  $O(\lg \lg \dots \lg U)$  for logarithm iterated for any constant number of times.

## 9 Discussions and Open Problems

In this paper, we assumed that the word-size  $w$  is at least  $\Omega(\lg U + \lg \sigma)$ , i.e., each (key, value) pair fits in  $O(1)$  words. When either the key or the value is larger than  $\Theta(w)$  bits, it would take super-constant time to just read the query or write the output on a RAM. The best query time one can hope for is  $O((\lg U + \lg \sigma)/w)$ .

When  $\lg \sigma \gg w$ , the only place being affected is Lemma 16, where we need to retrieve values longer than one word. Our data structure naturally supports such long answers in optimal time. When  $\lg U \gg w$ , a similar strategy to Section 8 applies. We view the first  $O(w)$  bits of an element in  $[U]$  as its “hash value”.

If it turns out that all keys have different “hash values”, it suffices to add the remaining bits of the key into its value. Otherwise, if multiple keys share the same prefix, then we will be able to save  $O(w)$  bits for every extra key with the same prefix.

Our dictionary data structure supports each query in constant expected time. A major open question is to design *deterministic* succinct dictionary with similar bounds, or to prove this is impossible. Our approach crucially relies on sampling a small set of keys to be the “hard queries”. There is always a small portion of the data stored using the rank data structure of Pătraşcu, which takes  $O(\lg n)$  time to decode. “Derandomizing” this data structure seems to require a completely different strategy. On the other hand, proving lower bounds may also be challenging, as the common strategy of “designing a hard distribution and proving average-case lower bound” is doomed to fail. For any fixed input distribution, we could always fix and hardwire the random bits in the data structure, thus, our data structure uses only  $\text{OPT} + 1$  bits of space.

Our data structure only supports value-retrieval queries on a *fixed* set of (key, value) pairs, i.e., it solves the *static* dictionary problem. The *dynamic* dictionary problem further requires the data structure to support (key, value) insertions and deletions. It seems non-trivial to extend our data structure to such updates to the data, even with good amortized expected time, although our data structure has  $\tilde{O}(n)$  preprocessing time, thus using a hash table to store a “buffer” of size  $n^{1-\epsilon}$  and using the technique of *global-rebuilding* [Ove83], one can get  $\tilde{O}(n^{1-\epsilon})$  redundancy,  $n^\epsilon$  update time and expected constant query time. On the other hand, it is also possible to update a (key, value) to a new value in our data structure, in  $O(1)$  expected time.

Finally, the dependence on  $U$  in the extra bits is intriguing. In the RAM model, the dependence is very slow-growing, but still super constant. We believe it is not necessary, but it is unclear how to remove this extra small term. On the other hand, note that in the cell-probe model, it can actually be entirely removed (even for very large  $U$ ). This is because when  $U$  is large enough so that  $\lg \lg U$  becomes unignorable, we could simply apply Lemma 25. This strategy does not work on RAM, since it requires a large lookup table, which can only be hardwired in a cell-probe data structure.

## Acknowledgment

The author would like to thank anonymous reviewers for helpful comments.

## References

- [BL13] Karl Bringmann and Kasper Green Larsen. Succinct sampling from discrete distributions. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 775–782, 2013.
- [BM99] Andrej Brodnik and J. Ian Munro. Membership in constant time and almost-minimum space. *SIAM J. Comput.*, 28(5):1627–1640, 1999.
- [BMRV02] Harry Buhrman, Peter Bro Miltersen, Jaikumar Radhakrishnan, and Srinivasan Venkatesh. Are bitvectors optimal? *SIAM J. Comput.*, 31(6):1723–1744, 2002.
- [CW79] Larry Carter and Mark N. Wegman. Universal classes of hash functions. *J. Comput. Syst. Sci.*, 18(2):143–154, 1979.
- [DPT10] Yevgeniy Dodis, Mihai Pătraşcu, and Mikkel Thorup. Changing base without losing space. In *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, pages 593–602, 2010.

- [FKS84] Michael L. Fredman, János Komlós, and Endre Szemerédi. Storing a sparse table with  $O(1)$  worst case access time. *J. ACM*, 31(3):538–544, 1984.
- [FM95] Faith E. Fich and Peter Bro Miltersen. Tables should be sorted (on random access machines). In *Algorithms and Data Structures, 4th International Workshop, WADS '95, Kingston, Ontario, Canada, August 16-18, 1995, Proceedings*, pages 482–493, 1995.
- [FN93] Amos Fiat and Moni Naor. Implicit  $O(1)$  probe search. *SIAM J. Comput.*, 22(1):1–10, 1993.
- [FNSS92] Amos Fiat, Moni Naor, Jeanette P. Schmidt, and Alan Siegel. Nonoblivious hashing. *J. ACM*, 39(4):764–782, 1992.
- [FW93] Michael L. Fredman and Dan E. Willard. Surpassing the information theoretic bound with fusion trees. *J. Comput. Syst. Sci.*, 47(3):424–436, 1993.
- [GORR09] Roberto Grossi, Alessio Orlandi, Rajeev Raman, and S. Srinivasa Rao. More haste, less waste: Lowering the redundancy in fully indexable dictionaries. In *26th International Symposium on Theoretical Aspects of Computer Science, STACS 2009, February 26-28, 2009, Freiburg, Germany, Proceedings*, pages 517–528, 2009.
- [Jac89] Guy Jacobson. Space-efficient static trees and graphs. In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 549–554, 1989.
- [MHMP15] A. Makhdoumi, S. Huang, M. Médard, and Y. Polyanskiy. On locally decodable source coding. In *2015 IEEE International Conference on Communications (ICC)*, pages 4394–4399, 2015.
- [Mil96] Peter Bro Miltersen. Lower bounds for static dictionaries on rams with bit operations but no multiplication. In *Automata, Languages and Programming, 23rd International Colloquium, ICALP96, Paderborn, Germany, 8-12 July 1996, Proceedings*, pages 442–453, 1996.
- [MNSW98] Peter Bro Miltersen, Noam Nisan, Shmuel Safra, and Avi Wigderson. On data structures and asymmetric communication complexity. *J. Comput. Syst. Sci.*, 57(1):37–49, 1998.
- [Ove83] Mark H. Overmars. *The Design of Dynamic Data Structures*. Lecture Notes in Economic and Mathematical Systems. Springer-Verlag, 1983.
- [Pag01a] Rasmus Pagh. Low redundancy in static dictionaries with constant query time. *SIAM J. Comput.*, 31(2):353–363, 2001.
- [Pag01b] Rasmus Pagh. On the cell probe complexity of membership and perfect hashing. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, pages 425–432, 2001.
- [Păt08] Mihai Pătraşcu. Succincter. In *Proc. 49th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 305–313, 2008.
- [PT06] Mihai Pătraşcu and Mikkel Thorup. Time-space trade-offs for predecessor search. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 232–240, 2006.

- [PT07] Mihai Pătraşcu and Mikkel Thorup. Randomization does not help searching predecessors. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, pages 555–564, 2007.
- [PV10] Mihai Pătraşcu and Emanuele Viola. Cell-probe lower bounds for succinct partial sums. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 117–122, 2010.
- [RRR07] Rajeev Raman, Venkatesh Raman, and Srinivasa Rao Satti. Succinct indexable dictionaries with applications to encoding  $k$ -ary trees, prefix sums and multisets. *ACM Trans. Algorithms*, 3(4):43, 2007.
- [SS90] Jeanette P. Schmidt and Alan Siegel. The spatial complexity of oblivious  $k$ -probe hash functions. *SIAM J. Comput.*, 19(5):775–786, 1990.
- [Tho13] Mikkel Thorup. Mihai Pătraşcu: Obituary and open problems. *Bulletin of the EATCS*, 109:7–13, 2013.
- [TY79] Robert Endre Tarjan and Andrew Chi-Chih Yao. Storing a sparse table. *Commun. ACM*, 22(11):606–611, 1979.
- [Vio12] Emanuele Viola. Bit-probe lower bounds for succinct data structures. *SIAM J. Comput.*, 41(6):1593–1604, 2012.
- [VWY19] Emanuele Viola, Omri Weinstein, and Huacheng Yu. How to store a random walk. *CoRR*, abs/1907.10874, 2019.
- [Yao81] Andrew Chi-Chih Yao. Should tables be sorted? *J. ACM*, 28(3):615–628, 1981.
- [Yu19] Huacheng Yu. Optimal succinct rank data structure via approximate nonnegative tensor decomposition. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019.*, pages 955–966, 2019.

## A Proofs for Fractional-length Strings

In the section, we prove the propositions from Section 4. We first show that two strings can be concatenated.

**Proposition 37.** *Let  $s_1, s_2 \geq 0$ . Given two binary strings  $S_1 = (M_1, K_1)$  and  $S_2 = (M_2, K_2)$  of  $s_1$  and  $s_2$  bits respectively, they can be concatenated into one string  $S = (M, K)$  of length at most  $s_1 + s_2 + 2^{-\kappa+2}$ , and both  $M_1$  and  $M_2$  are (consecutive) substrings of  $M$ . Moreover, given the values of  $s_1$  and  $s_2$ , both  $S_1$  and  $S_2$  can be decoded using constant time and one access to  $S$ , i.e., a decoding algorithm recovers  $K_1$  and  $K_2$ , and finds the starting locations of  $M_1$  and  $M_2$  within  $M$  using constant time and one access to  $S$ .*

After decoding  $S_1$  and  $S_2$ , any further access to the two strings can be performed as if they were stored explicitly.

*Proof.* To concatenate two strings, let us first combine  $K_1$  and  $K_2$  into a single integer  $K' \in [\text{range}(K_1) \cdot \text{range}(K_2)]$ :

$$K' := K_1 \cdot \text{range}(K_2) + K_2.$$

If  $s_1 + s_2 < \kappa + 1$ , we simply let  $K = K'$ , and let  $M$  be the empty string. Then  $\lg(\text{range}(K)) = s_1 + s_2 < \kappa + 1$ ,  $(M, K)$  is the concatenation.

Next, we assume  $s_1 + s_2 \geq \kappa + 1$ . In this case, in the final string  $\mathcal{S} = (M, K)$ ,  $M$  will be the concatenation of  $M_1$ ,  $M_2$  and the lowest bits of  $K'$ . More specifically, let  $|M|$  be  $\lfloor s_1 + s_2 + 2^{-\kappa+2} \rfloor - \kappa$ , and let  $\text{range}(K)$  be  $\lfloor 2^{\kappa + \text{frac}(s_1 + s_2 + 2^{-\kappa+2})} \rfloor$ .<sup>6</sup> It is easy to verify that  $|M| + \lg(\text{range}(K)) \leq s_1 + s_2 + 2^{-\kappa+2}$  and  $\text{range}(K) \in [2^\kappa, 2^{\kappa+1})$ .

We set

$$M := M_1 \circ M_2 \circ (K' \bmod 2^{|M| - |M_1| - |M_2|})_2,$$

where  $(x)_2$  is the binary representation of  $x$ , and

$$K := K' \text{div } 2^{|M| - |M_1| - |M_2|}.$$

To see why  $K$  is at most  $\text{range}(K) - 1$ , we have

$$\begin{aligned} K &\leq (\text{range}(K_1) \cdot \text{range}(K_2) - 1) \text{div } 2^{|M| - |M_1| - |M_2|} \\ &= \left( 2^{s_1 - |M_1| + s_2 - |M_2|} - 1 \right) \text{div } 2^{|M| - |M_1| - |M_2|} \\ &\leq 2^{s_1 + s_2 - |M|} \\ &\leq (\text{range}(K) + 1) \cdot 2^{-2^{-\kappa+2}} \\ &\leq \text{range}(K) - 1, \end{aligned}$$

where the last inequality uses the fact that  $2^{-2^{-\kappa+2}} \leq 1 - 2^{-\kappa+1}$  and  $\text{range}(K) \geq 2^\kappa$ .

To decode  $\mathcal{S}_1$  and  $\mathcal{S}_2$ , observe that  $\mathcal{S}[|M_1| + |M_2|, |M|]$  encodes exactly  $K'$ . We access  $\mathcal{S}$  to retrieve its value, and compute  $K_1$  and  $K_2$  using  $K_1 := K' \text{div } \text{range}(K_2)$  and  $K_2 := K' \bmod \text{range}(K_2)$ . By our construction,  $M_1$  is  $M[0, |M_1| - 1]$  and  $M_2$  is  $M[|M_1|, |M_1| + |M_2| - 1]$ . Hence, we decode  $\mathcal{S}_1$  and  $\mathcal{S}_2$  in constant time and one access to  $\mathcal{S}$ .  $\square$

Using ideas similar to [DPT10], we show that multiple strings can be concatenated, allowing fast decoding of any given string.

**Proposition 38.** *Let  $s_1, \dots, s_B \geq \kappa$ . Suppose there are numbers  $\tilde{T}_0, \dots, \tilde{T}_B$  such that*

- $\tilde{T}_0 = 0$  and  $s_i + 1 \geq \tilde{T}_i - \tilde{T}_{i-1} \geq s_i + 2^{-\kappa+2}$ ;
- each  $\tilde{T}_i$  is of the form  $\tilde{T}_i = \tilde{m}_i + \lg \tilde{R}_i$ , where  $\tilde{R}_i \in [2^\kappa, 2^{\kappa+1})$  and  $\tilde{m}_i \geq 0$  are integers;
- for any given  $i$ ,  $\tilde{T}_i$  can be computed in  $O(t)$  time.

*Then given  $B$  strings  $\mathcal{S}_1, \dots, \mathcal{S}_B$ , where  $\mathcal{S}_i = (M_i, K_i)$  has length  $s_i$ , they can be concatenated into one string  $\mathcal{S} = (M, K)$  of length  $\tilde{T}_B$ , and each  $M_i$  is a (consecutive) substring of  $M$ . Moreover, for any given  $i$ ,  $\mathcal{S}_i$  can be decoded using  $O(t)$  time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $K_i$ , and finds the starting location of  $M_i$  using  $O(t)$  time and two accesses to  $\mathcal{S}$ .*

*Proof.* Without loss of generality, we assume  $B$  is odd, since otherwise, we could first apply the following argument to the first  $B - 1$  strings, then apply Proposition 37 to concatenate the outcome with the last string  $\mathcal{S}_B$ .

To concatenate all strings, we break  $\mathcal{S}_2, \dots, \mathcal{S}_B$  into  $(B - 1)/2$  pairs, where the  $j$ -th pair consists of  $\mathcal{S}_{2j}$  and  $\mathcal{S}_{2j+1}$ . We start with the first string  $\mathcal{S}_1$ , and repeatedly “append” the pairs to it. More specifically,

---

<sup>6</sup>Recall that  $\text{frac}(x) = x - \lfloor x \rfloor$ .

let  $\mathcal{S}^{(0)} := \mathcal{S}_1$ . Suppose we have concatenated  $\mathcal{S}_1$  and the first  $j-1$  pairs into  $\mathcal{S}^{(j-1)} = (M^{(j-1)}, K^{(j-1)})$ , such that  $|M^{(j-1)}| = \tilde{m}_{2j-1}$  and  $\text{range}(K^{(j-1)}) = \tilde{R}_{2j-1}$ . In particular, it has length  $\tilde{T}_{2j-1}$ . Now, we show how to “append”  $\mathcal{S}_{2j}$  and  $\mathcal{S}_{2j+1}$  to it.

To this end, we combine  $K_{2j}$  and  $K_{2j+1}$  into a single integer  $L_j$ ,

$$L_j := K_{2j} \cdot \text{range}(K_{2j+1}) + K_{2j+1}.$$

Thus,  $\text{range}(L_j) = \text{range}(K_{2j}) \cdot \text{range}(K_{2j+1})$ , and we have  $\text{range}(L_j) \in [2^{2\kappa}, 2^{2\kappa+2})$ . Then, we re-break  $L_j$  into a pair  $(X_j, Y_j)$ , such that the product of  $\text{range}(X_j)$  and  $\text{range}(K^{(j-1)})$  is close to a power of two: we set

$$\text{range}(X_j) := \left\lfloor \frac{2^{\lfloor \tilde{T}_{2j+1} \rfloor - \lfloor \tilde{T}_{2j-1} \rfloor - |M_{2j}| - |M_{2j+1}|}}{\text{range}(K^{(j-1)})} \right\rfloor,$$

and

$$\text{range}(Y_j) := \left\lceil \frac{\text{range}(L_j)}{\text{range}(X_j)} \right\rceil.$$

Note that

$$2\kappa \leq \lfloor \tilde{T}_{2j+1} \rfloor - \lfloor \tilde{T}_{2j-1} \rfloor - |M_{2j}| - |M_{2j+1}| \leq 2\kappa + 4.$$

To break  $L_j$  into such a pair, we let  $Y_j := L_j \text{ div } \text{range}(X_j)$  and  $X_j := L_j \text{ mod } \text{range}(X_j)$ . Next, we combine  $K^{(j-1)}$  and  $X_j$  into an integer  $Z_j$  smaller than  $2^{\lfloor \tilde{T}_{2j+1} \rfloor - \lfloor \tilde{T}_{2j-1} \rfloor - |M_{2j}| - |M_{2j+1}|}$ : let  $Z_j := K^{(j-1)} \cdot \text{range}(X_j) + X_j$ .

Finally, we let  $\mathcal{S}^{(j)} := (M^{(j)}, K^{(j)})$ , where

$$M^{(j)} := M^{(j-1)} \circ (Z_j)_2 \circ M_{2j} \circ M_{2j+1},$$

and

$$K^{(j)} := Y_j.$$

The length of  $M^{(j)}$

$$\begin{aligned} |M^{(j)}| &= |M^{(j-1)}| + (\lfloor \tilde{T}_{2j+1} \rfloor - \lfloor \tilde{T}_{2j-1} \rfloor - |M_{2j}| - |M_{2j+1}|) + |M_{2j}| + |M_{2j+1}| \\ &= \lfloor \tilde{T}_{2j+1} \rfloor - \kappa \\ &= \tilde{m}_{2j+1}. \end{aligned}$$

The range of  $K^{(j)}$  has size

$$\begin{aligned} \text{range}(K^{(j)}) &< \frac{\text{range}(L_j)}{\text{range}(X_j)} + 1 \\ &\leq \frac{\text{range}(K_{2j}) \cdot \text{range}(K_{2j+1})}{\frac{2^{\lfloor \tilde{T}_{2j+1} \rfloor - \lfloor \tilde{T}_{2j-1} \rfloor - |M_{2j}| - |M_{2j+1}|}}{\text{range}(K^{(j-1)})} - 1} + 1 \\ &\leq \frac{\text{range}(K_{2j}) \cdot \text{range}(K_{2j+1}) \cdot \text{range}(K^{(j-1)})}{2^{\lfloor \tilde{T}_{2j+1} \rfloor - \lfloor \tilde{T}_{2j-1} \rfloor - |M_{2j}| - |M_{2j+1}|}} \cdot (1 - 2^{-\kappa+1})^{-1} + 1 \\ &= 2^{s_{2j} + s_{2j+1} + \tilde{T}_{2j-1} - |M^{(j)}|} \cdot (1 - 2^{-\kappa+1})^{-1} + 1 \\ &\leq 2^{\kappa + \text{frac}(\tilde{T}_{2j+1}) - 2^{\kappa+3}} \cdot (1 - 2^{-\kappa+1})^{-1} + 1 \\ &\leq 2^{\kappa + \text{frac}(\tilde{T}_{2j+1})} \end{aligned}$$

$$= \tilde{R}_{2j+1}.$$

Thus,  $\mathcal{S}^{(j)}$  has length  $\tilde{T}_{2j+1}$ , and hence, the final string  $\mathcal{S} := \mathcal{S}^{((B-1)/2)}$  has length  $\tilde{T}_B$ .

Next, we show that each  $\mathcal{S}_i$  can be decoded in  $O(t)$  time and two accesses to  $\mathcal{S}$ . If  $i = 1$ , we compute  $\tilde{T}_3$  in  $O(t)$  time, and then compute  $\text{range}(Z_1)$ ,  $\text{range}(X_1)$  and  $|M_1|$ . Thus,  $M_1 = M[0, |M_1| - 1]$  and  $Z_1$  is stored in  $M$  immediately after  $M_1$ . By making one access to  $\mathcal{S}$ , we recover the value of  $Z_1$ , and hence,  $K_1$  can be computed using  $K_1 = Z_1 \text{ div } \text{range}(X_1)$ . This decodes  $\mathcal{S}_1$  in  $O(t)$  time and one access to  $\mathcal{S}$ .

If  $i > 1$ , let  $j = \lfloor i/2 \rfloor$ , i.e.,  $\mathcal{S}_i$  is in the  $j$ -th pair. We first compute  $\tilde{T}_{2j-1}$  and  $\tilde{T}_{2j+1}$  in  $O(t)$  time. They determine  $\text{range}(Z_j)$ ,  $\text{range}(X_j)$  and  $|M^{(j-1)}|$ , as well as the starting location of  $M_i$ . Thus,  $Z_j$  can be recovered with one access to  $\mathcal{S}$ .  $X_j$  can be computed using  $X_j = Z_j \text{ mod } \text{range}(X_j)$ . Similarly, we then recover  $Z_{j+1}$ , and  $Y_j$  can be computed using  $Y_j = K^{(j)} = Z_{j+1} \text{ div } \text{range}(X_{j+1})$  (if  $\mathcal{S}_i$  is in the last pair,  $Y_j$  is simply  $K$  in the final string). This recovers both  $X_j$  and  $Y_j$ . Next, we recover  $L_j$  using  $L_j = Y_j \cdot \text{range}(X_j) + X_j$ , and compute  $K_{2j}$  and  $K_{2j+1}$  using  $K_{2j} = L_j \text{ div } \text{range}(K_{2j+1})$  and  $K_{2j+1} = L_j \text{ mod } \text{range}(K_{2j+1})$ . In particular, it recovers the value of  $K_i$ , and hence, it decodes  $\mathcal{S}_i$ .  $\square$

**Proposition 7 (restated).** *Let  $s_1, \dots, s_B \geq \kappa$ . Suppose for any given  $i$ ,  $s_1 + \dots + s_i$  can be approximated (deterministically) in  $O(t)$  time with an additive error of at most  $2^{-\kappa}$ . Then given  $B$  strings  $\mathcal{S}_1, \dots, \mathcal{S}_B$ , where  $\mathcal{S}_i = (M_i, K_i)$  has length  $s_i$ , they can be concatenated into one string  $\mathcal{S} = (M, K)$  of length at most*

$$s_1 + \dots + s_B + (B - 1) \cdot 2^{-\kappa+4},$$

*so that each  $M_i$  is a (consecutive) substring of  $M$ . Moreover, for any given  $i$ ,  $\mathcal{S}_i$  can be decoded using  $O(t)$  time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $K_i$ , and finds the starting location of  $M_i$  using  $O(t)$  time and two accesses to  $\mathcal{S}$ .*

*Proof.* Suppose we can compute  $\tilde{S}_i = s_1 + \dots + s_i \pm 2^{-\kappa}$ . We set  $\tilde{m}_i = \lfloor \tilde{S}_i + (i - 1) \cdot 2^{-\kappa+3} \rfloor - \kappa$ ,  $\tilde{R}_i = \lfloor 2^{\tilde{S}_i + (i-1) \cdot 2^{-\kappa+3} - \tilde{m}_i} \rfloor$  and  $\tilde{T}_i = \tilde{m}_i + \lg \tilde{R}_i$ . Then  $\tilde{T}_i \leq \tilde{S}_i + (i - 1) \cdot 2^{-\kappa+3}$  and  $\tilde{T}_i > \tilde{S}_i + (i - 1) \cdot 2^{-\kappa+3} - 2^{-\kappa+1}$ . Therefore,

$$\begin{aligned} \tilde{T}_i - \tilde{T}_{i-1} &\geq \tilde{S}_i - \tilde{S}_{i-1} + 2^{-\kappa+3} - 2^{-\kappa+1} \\ &\geq s_i + 2^{-\kappa+2}. \end{aligned}$$

Also,  $\tilde{T}_i - \tilde{T}_{i-1} \leq s_i + 2^{-\kappa+4}$ . Finally, by Proposition 38, the size of the data structure is at most  $\tilde{T}_B \leq s_B + (B - 1) \cdot 2^{-\kappa+3} \leq s_1 + \dots + s_B + (B - 1) \cdot 2^{-\kappa+4}$ .  $\square$

In particular, by storing approximations of all  $B$  prefix sums in a lookup table of size  $O(B)$ , the length of  $\mathcal{S}$  is at most  $s_1 + \dots + s_B + (B - 1)2^{-\kappa+4}$  and each  $\mathcal{S}_i$  can be decoded in  $O(1)$  time.

**Proposition 8 (restated).** *Let  $s_1, \dots, s_B \geq 0$ . There is a lookup table of size  $O(B)$ . Given  $B$  strings  $\mathcal{S}_1, \dots, \mathcal{S}_B$ , where  $\mathcal{S}_i = (M_i, K_i)$  has length  $s_i$ , they can be concatenated into one string  $\mathcal{S} = (M, K)$  of length at most*

$$s_1 + \dots + s_B + (B - 1)2^{-\kappa+4},$$

*so that each  $M_i$  is a (consecutive) substring of  $M$ . Moreover, assuming we can make random accesses to the lookup table,  $\mathcal{S}_i$  can be decoded using constant time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $K_i$ , and finds the starting location of  $M_i$  using constant time and two accesses to  $\mathcal{S}$ .*

*Proof.* If all  $s_i \geq \kappa$ , the proposition is an immediate corollary of Proposition 7, as we could simply store the approximations of all  $B$  prefix sums. For general  $s_i \geq 0$ , we group the strings so that each group has length at least  $\kappa$ .

We greedily divide all strings into groups: Pick the first  $i_1$  such that  $s_1 + \dots + s_{i_1} \geq \kappa$ , then pick the first  $i_2$  such that  $s_{i_1+1} + \dots + s_{i_2} \geq \kappa$ , etc. Then each group has total length at least  $\kappa$ , possibly except for the last group. We store in the lookup table, which group each string belongs to, and the values of  $i_1, i_2, \dots$ . Then consider a group consisting of  $\mathcal{S}_a, \dots, \mathcal{S}_b$ , we must have  $s_a + \dots + s_{b-1} < \kappa$ , which means that they can be combined into one single integer smaller than  $\prod_{i=a}^{b-1} \text{range}(K_i) < 2^\kappa$ , e.g.,

$$K := \sum_{i=a}^{b-1} K_i \cdot \prod_{j=a}^{i-1} \text{range}(K_j).$$

If we store  $\prod_{j=a}^{i-1} \text{range}(K_j)$  and  $\text{range}(K_i)$  in the lookup table for each  $i$  in the group, then  $K_i$  can be recovered from  $K$  using

$$K_i = (K \text{ div } \prod_{j=a}^{i-1} \text{range}(K_j)) \bmod \text{range}(K_i).$$

This concatenates all strings in the group except the last one. We then apply Proposition 37 to concatenate the last string in the group to it. Then we apply Proposition 7 to concatenate the strings obtained from each group (except for the last group), using the lookup table. Finally, we concatenate the string obtained from the last group to it.

Concatenating strings in each group loses at most  $2^{-\kappa+2}$  due to Proposition 37. The length of the final string is at most  $s_1 + \dots + s_B + (B-1)2^{-\kappa+4}$ . The lookup table has size  $O(B)$ .  $\square$

Next, we show that an integer  $i \in [C]$  can be *fused* into a string.

**Proposition 39.** *Let  $s_1, \dots, s_C \geq 0$ . Suppose there are numbers  $\tilde{T}_1, \dots, \tilde{T}_C$  such that*

- $2^{\tilde{T}_i} - 2^{\tilde{T}_{i-1}} \geq 2^{s_i}$ ;
- *each  $\tilde{T}_i$  is of the form  $\tilde{T}_i = \tilde{m} + \lg \tilde{R}_i$ , where  $\tilde{m}, \tilde{R}_i$  are integers;*
- *$\tilde{T}_C$  is a valid length, i.e.,  $\tilde{m} = 0$  and  $\tilde{R}_C \in [1, 2^\kappa)$ , or  $\tilde{m} \geq 0$  and  $\tilde{R}_C \in [2^\kappa, 2^{\kappa+1})$ ;*
- *for any given  $K$ , the largest  $i \leq C$  such that  $\tilde{R}_i \leq K$  can be computed in  $O(t)$  time.*

*Then given  $i \in \{1, \dots, C\}$  and string  $\mathcal{S}_i = (M_i, K_i)$  of length  $s_i$ , the pair  $(i, \mathcal{S}_i)$  can be stored in  $\mathcal{S} = (M, K)$  of length  $\tilde{T}_C$ , and  $M_i$  is a (consecutive) substring of  $M$ . Moreover, we can recover the value of  $i$  and decode  $\mathcal{S}_i$  using  $O(t)$  time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $i$ ,  $K_i$ , and finds the starting location of  $M_i$  using  $O(t)$  time and two accesses to  $\mathcal{S}$ .*

*Proof.* Clearly, we have  $s_i \leq \tilde{T}_C$  for all  $i$ , and hence,  $|M_i| \leq \tilde{m}$ . We first increase the length of  $|M_i|$  to  $\tilde{m}$  by appending the least significant bits of  $K_i$  to it. That is, let

$$M := M_i \circ (K_i \bmod 2^{\tilde{m}-|M_i|})_2.$$

Next, we encode the remaining information of  $(i, \mathcal{S}_i)$  in  $K$ , i.e., encode  $i$  and the top bits of  $K_i$ :

$$K := \tilde{R}_{i-1} + (K_i \text{ div } 2^{\tilde{m}-|M_i|}),$$

where  $\tilde{R}_0$  is assumed to be 0. Note that we have

$$\tilde{R}_{i-1} + (\text{range}(K_i) - 1) \text{ div } 2^{\tilde{m}-|M_i|} < \tilde{R}_{i-1} + \text{range}(K_i) \cdot 2^{|M_i|-\tilde{m}}$$

$$\begin{aligned}
&= \tilde{R}_{i-1} + 2^{s_i - \tilde{m}} \\
&= 2^{-\tilde{m}}(2^{\tilde{T}_{i-1}} + 2^{s_i}) \\
&\leq 2^{\tilde{T}_i - \tilde{m}} \\
&= \tilde{R}_i.
\end{aligned}$$

That is, the value of  $K$  determines both  $i$  and  $K_i \div 2^{\tilde{m} - |M_i|}$ , and  $\text{range}(K)$  is at most  $\tilde{R}_C$ . Thus,  $\mathcal{S}$  is a string of length  $\tilde{T}_C$ .

To decode  $i$  and  $\mathcal{S}_i$ , we first access  $\mathcal{S}$  to retrieve  $K$ . Then we compute the largest  $i \leq C$  such that  $\tilde{R}_i \leq K$  in  $O(t)$  time. By the argument above, it recovers the value of  $i$  and determines

$$(K_i \div 2^{\tilde{m} - |M_i|}) = K - \tilde{R}_i.$$

To decode  $\mathcal{S}_i$ , observe that  $M_i = M[0, |M_i| - 1]$ , and  $M[|M_i|, \tilde{m} - 1]$  stores the value of  $K_i \bmod 2^{\tilde{m} - |M_i|}$ . If  $\tilde{m} - |M_i| \leq \kappa + 1$ , we retrieve its value using one access, and together with  $K_i \div 2^{\tilde{m} - |M_i|}$ , it determines  $K_i$ . Otherwise, since  $K_i < 2^{\kappa+1}$ , its value is entirely stored in  $M$  (in its binary representation). We simply make one access to retrieve it. In both cases, we recover the value of  $i$  and decode  $\mathcal{S}_i$  in  $O(t)$  time and two accesses to  $\mathcal{S}$ .  $\square$

**Proposition 9 (restated).** *Let  $s_1, \dots, s_C \geq 0$ . Suppose for any given  $i$ ,  $2^{s_1} + \dots + 2^{s_i}$  can be approximated (deterministically) in  $O(t)$  time with an additive error of at most  $(2^{s_1} + \dots + 2^{s_C}) \cdot 2^{-\kappa-3}$ . Then given  $i \in \{1, \dots, C\}$  and string  $\mathcal{S}_i = (M_i, K_i)$  of length  $s_i$ , the pair  $(i, \mathcal{S}_i)$  can be stored in  $\mathcal{S} = (M, K)$  of length at most*

$$\lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+4},$$

*so that  $M_i$  is a (consecutive) substring of  $M$ . Moreover, we can recover the value of  $i$  and decode  $\mathcal{S}_i$  using  $O(t \lg C)$  time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $i$ ,  $K_i$ , and finds the starting location of  $M_i$  using  $O(t \lg C)$  time and two accesses to  $\mathcal{S}$ .*

*Proof.* We compute  $\tilde{S}_i = (2^{s_1} + \dots + 2^{s_i}) \pm (2^{s_1} + \dots + 2^{s_C}) \cdot 2^{-\kappa-3}$ . If  $2^{s_1} + \dots + 2^{s_C} < \kappa$ , then the error term  $(2^{s_1} + \dots + 2^{s_C}) \cdot 2^{-\kappa-3} < 1/8$ . However, each  $2^{s_1} + \dots + 2^{s_i}$  must be an integer by definition.  $\tilde{S}_i$  rounded to the nearest integer is the accurate value of  $2^{s_1} + \dots + 2^{s_i}$ . To apply Proposition 39, we simply set  $\tilde{m} := 0$ ,  $\tilde{R}_i := \lfloor \tilde{S}_i + 1/2 \rfloor$  for  $i = 1, \dots, C$  and  $\tilde{T}_i = \tilde{m} + \lg \tilde{R}_i$ . It is easy to verify that  $\tilde{R}_i - \tilde{R}_{i-1} \geq 2^{s_i}$ ;  $\tilde{T}_C$  is a valid length. For any given  $K$ , by doing a binary search, the largest  $i$  such that  $\tilde{R}_i \leq K$  can be found in  $O(t \lg C)$  time. Thus, by Proposition 39, the pair  $(i, \mathcal{S}_i)$  can be stored using space

$$\tilde{T}_C = \lg(2^{s_1} + \dots + 2^{s_C}),$$

and allowing  $O(t \lg C)$  time for decoding.

Next, we consider the case where  $2^{s_1} + \dots + 2^{s_C} \geq \kappa$ . To apply Proposition 39, we let  $\tilde{T}_C$  be the largest valid length smaller than  $\lg \tilde{S}_C + C \cdot 2^{-\kappa+3}$ . That is, we set

$$\tilde{m} := \lfloor \lg \tilde{S}_C + C \cdot 2^{-\kappa+3} \rfloor - \kappa.$$

Then

$$\tilde{R}_C := \lfloor \tilde{S}_C \cdot 2^{C \cdot 2^{-\kappa+3}} \cdot 2^{-\tilde{m}} \rfloor,$$

and  $\tilde{T}_C = \tilde{m} + \lg \tilde{R}_C$ . Then for  $i < C$ , we let

$$\tilde{R}_i := \lfloor \tilde{S}_i \cdot 2^{-\tilde{m}} \rfloor + 2(i-1),$$

and  $\tilde{T}_i = \tilde{m} + \lg \tilde{R}_i$ .

To apply Proposition 39, we verify that  $2^{\tilde{T}_i} - 2^{\tilde{T}_{i-1}} \geq 2^{s_i}$ . To see this, for  $i < C$ , we have

$$\begin{aligned} 2^{\tilde{T}_i} - 2^{\tilde{T}_{i-1}} &= 2^{\tilde{m}} \cdot (\tilde{R}_i - \tilde{R}_{i-1}) \\ &\geq 2^{\tilde{m}} \cdot (\tilde{S}_i \cdot 2^{-\tilde{m}} - \tilde{S}_{i-1} \cdot 2^{-\tilde{m}} + 1) \\ &\geq 2^{s_i} + 2^{\tilde{m}} - (2^{s_1} + \dots + 2^{s_C}) \cdot 2^{-\kappa-2}. \end{aligned}$$

On the other hand,  $\tilde{S}_C = (2^{s_1} + \dots + 2^{s_C}) \cdot (1 \pm 2^{-\kappa-3})$ , i.e.,  $2^{s_1} + \dots + 2^{s_C} = \tilde{S}_C \cdot (1 \pm 2^{-\kappa-3})^{-1}$ .

$$2^{\tilde{m}} - (2^{s_1} + \dots + 2^{s_C}) \cdot 2^{-\kappa-2} \geq 2^{\tilde{m}} - \tilde{S}_C \cdot 2^{-\kappa-1} \geq 0.$$

Thus,  $2^{\tilde{T}_i} - 2^{\tilde{T}_{i-1}} \geq 2^{s_i}$  for  $i < C$ . For  $i = C$ , it suffices to show  $\lfloor \tilde{S}_C \cdot 2^{-\tilde{m}} \rfloor + 2(C-1) \leq \tilde{R}_C$ . Indeed, we have

$$\begin{aligned} \tilde{R}_C - (\lfloor \tilde{S}_C \cdot 2^{-\tilde{m}} \rfloor + 2(C-1)) &\geq \tilde{S}_C \cdot 2^{C \cdot 2^{-\kappa+3}} \cdot 2^{-\tilde{m}} - 1 - \tilde{S}_C \cdot 2^{-\tilde{m}} - 2(C-1) \\ &\geq \tilde{S}_C \cdot 2^{-\tilde{m}} \cdot (2^{C \cdot 2^{-\kappa+3}} - 1) - 2C \\ &\geq \tilde{S}_C \cdot 2^{-\tilde{m}} \cdot C \cdot 2^{-\kappa+2} - 2C. \end{aligned}$$

Since  $\tilde{m} + \kappa \leq \lg \tilde{S}_C + 1$ , it is at least 0.

Since each  $\tilde{R}_i$  can be computed in  $O(t)$  time, by doing a binary search, for any given  $K$ , we can find the largest  $i$  such that  $\tilde{R}_i \leq K$  in  $O(t \lg C)$  time. By Proposition 39, we obtain a data structure of size

$$\tilde{T}_C \leq \lg \tilde{S}_C + C \cdot 2^{-\kappa+3} \leq \lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+4}.$$

This proves the proposition.  $\square$

Similar to the concatenation, the decoding algorithm takes constant time if we use a lookup table of size  $O(C)$ .

**Proposition 10 (restated).** *Let  $s_1, \dots, s_C \geq 0$ . There is a lookup table of size  $O(C)$ . Given  $i \in \{1, \dots, C\}$  and string  $\mathcal{S}_i = (M_i, K_i)$  of length  $s_i$ , the pair  $(i, \mathcal{S}_i)$  can be stored in  $\mathcal{S} = (M, K)$  of length*

$$\lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2},$$

*so that  $M_i$  is a (consecutive) substring of  $M$ . Moreover, assuming we can make random accesses to the lookup table, the value of  $i$  can be recovered and  $\mathcal{S}_i$  can be decoded using constant time and two accesses to  $\mathcal{S}$ , i.e., a decoding algorithm recovers  $i$ ,  $K_i$ , and finds the starting location of  $M_i$  using constant time and two accesses to  $\mathcal{S}$ .*

*Proof.* Without loss of generality, assume  $s_1 \leq \dots \leq s_C$ , since otherwise, we simply sort  $s_1, \dots, s_C$  and store the permutation in the lookup table.

To apply Proposition 39, if  $2^{s_1} + \dots + 2^{s_C} \leq 2^\kappa$ , we set

$$\tilde{m} := 0,$$

$$\tilde{R}_i = 2^{s_1} + \dots + 2^{s_i}$$

and  $\tilde{T}_i = \tilde{m} + \lg \tilde{R}_i$ . Otherwise, if  $2^{s_1} + \dots + 2^{s_C} > 2^\kappa$ , we set

$$\tilde{m} := \lfloor \lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2} \rfloor - \kappa,$$

for  $i < C$ , let

$$\tilde{R}_i := \lceil 2^{s_1 - \tilde{m}} \rceil + \dots + \lceil 2^{s_i - \tilde{m}} \rceil,$$

and

$$\tilde{R}_C := \max \{ \lceil 2^{s_1 - \tilde{m}} \rceil + \dots + \lceil 2^{s_C - \tilde{m}} \rceil, 2^\kappa \}.$$

Finally, let  $\tilde{T}_i = \tilde{m} + \lg \tilde{R}_i$ . Clearly, in both cases, we have  $2^{\tilde{T}_i} - 2^{\tilde{T}_{i-1}} \geq 2^{\tilde{m}} \cdot 2^{s_i - \tilde{m}} \geq 2^{s_i}$ . Also, we have  $\tilde{T}_C \leq \lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2}$ . This is because

$$\begin{aligned} \tilde{R}_C &< \max \{ (2^{s_1} + \dots + 2^{s_C}) \cdot 2^{-\tilde{m}} + C, 2^\kappa \} \\ &= \max \{ 2^{\kappa + \text{frac}(\lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2}) - C \cdot 2^{-\kappa+2}} + C, 2^\kappa \} \\ &\leq \max \{ 2^{\kappa + \text{frac}(\lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2})} \cdot (1 - C \cdot 2^{-\kappa+1}) + C, 2^\kappa \} \\ &\leq 2^{\kappa + \text{frac}(\lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2})}. \end{aligned}$$

Thus,  $\tilde{T}_C = \tilde{m} + \lg \tilde{R}_C \leq \lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2}$ .

To apply Proposition 9, we need to show that for any given  $K$ , the largest  $i$  such that  $\tilde{R}_i \leq K$  can be found in constant time. To this end, we store a *predecessor search* data structure for the set  $\{\tilde{R}_1, \dots, \tilde{R}_C\}$ . Note that the set of integers  $\{\tilde{R}_1, \dots, \tilde{R}_C\}$  has *monotone gaps*. That is, the difference between adjacent numbers is non-decreasing. Pătraşcu [Păt08] showed that for such sets, there is a predecessor search data structure using linear space and constant query time, i.e., there is an  $O(C)$ -sized data structure such that given an integer  $K$ , the query algorithm can answer in constant time the largest value in the set that is at most  $K$ . This data structure is stored in the lookup table (it only depends on  $s_1, \dots, s_C$ , but not the input string). To compute the index  $i$  rather than  $\tilde{R}_i$ , we simply store another hash table using perfect hashing in the lookup table. Hence, the lookup table has size  $O(C)$ .

The premises of Proposition 9 are all satisfied. The size of  $\mathcal{S}$  is  $\tilde{T}_C \leq \lg(2^{s_1} + \dots + 2^{s_C}) + C \cdot 2^{-\kappa+2}$ , and  $i$  and  $\mathcal{S}_i$  can be decoded in constant time. This proves the proposition.  $\square$

Next, we show that it is possible to divide a binary string into two substrings.

**Proposition 14 (restated).** *Let  $s_1, s_2, s \geq 3\kappa$  and  $s \leq s_1 + s_2 - 2^{-\kappa+2}$ . Then given a double-ended string  $\mathcal{S} = (K_h, M, K_t)$  of length  $s$ , a division algorithm outputs two double-ended strings  $\mathcal{S}_1 = (K_{1,h}, M_1, K_{1,t})$  and  $\mathcal{S}_2 = (K_{2,h}, M_2, K_{2,t})$  of lengths at most  $s_1$  and  $s_2$  respectively. Moreover,  $(K_{1,h}, M_1)$  is a prefix of  $\mathcal{S}$ ,  $(M_2, K_{2,t})$  is a suffix of  $\mathcal{S}$ , and  $K_{1,t}$  and  $K_{2,h}$  together determine  $M[|M_1|, |M| - |M_2| - 1]$ , i.e., the remaining bits of  $M$ .  $\text{range}(K_{i,h})$ ,  $\text{range}(K_{i,t})$  and  $|M_i|$  can be computed in  $O(1)$  time given  $\text{range}(K_h)$ ,  $\text{range}(K_t)$ ,  $|M|$  and  $s_1, s_2$ , for  $i = 1, 2$ .*

*Proof.* We first calculate the length of  $M_1$  and  $M_2$ , let  $|M_1| := \lfloor s_1 - \lg(\text{range}(K_h)) \rfloor - \kappa$  and  $M_2 := \lfloor s_2 - \lg(\text{range}(K_t)) \rfloor - \kappa$ . Then let

$$(K_{h,1}, M_1) := \mathcal{S}[-1, |M_1| - 1]$$

be a prefix, and

$$(M_2, K_{t,2}) := \mathcal{S}[|M| - |M_2|, |M|]$$

be a suffix. The remaining task is to divide the middle  $|M| - |M_1| - |M_2|$  bits of  $M$  into  $K_{t,1}$  and  $K_{h,2}$ .

To this end, we represent the middle bits as an integer  $L$  in the range  $[2^{|M|-|M_1|-|M_2|}]$ . The sizes of ranges of  $K_{t,1}$  and  $K_{h,2}$  can be calculated using

$$\text{range}(K_{t,1}) = \lfloor 2^{s_1 - \lg(\text{range}(K_h)) - |M_1|} \rfloor$$

and

$$\text{range}(K_{h,2}) = \lfloor 2^{s_2 - \lg(\text{range}(K_t)) - |M_2|} \rfloor.$$

Then let  $K_{t,1} := L \bmod \text{range}(K_{t,1})$  and  $K_{h,2} := L \text{ div } \text{range}(K_{t,1})$ . Clearly,  $K_{t,1} \in [\text{range}(K_{t,1})]$ . It suffices to show that  $K_{h,2}$  is in its range:

$$\begin{aligned} K_{h,2} &< \frac{2^{|M|-|M_1|-|M_2|}}{2^{s_1 - \lg(\text{range}(K_h)) - |M_1|} - 1} \\ &= \frac{2^{|M|-|M_2|-s_1 + \lg(\text{range}(K_h))}}{1 - 2^{-s_1 + \lg(\text{range}(K_h)) + |M_1|}} \\ &= \frac{2^{s - |M_2| - s_1 - \lg(\text{range}(K_t))}}{1 - 2^{-\kappa}} \\ &\leq \frac{2^{s_2 - |M_2| - \lg(\text{range}(K_t)) - 2^{-\kappa+2}}}{1 - 2^{-\kappa}} \\ &< (\text{range}(K_{h,2}) + 1) \cdot \frac{2^{-2^{-\kappa+2}}}{1 - 2^{-\kappa}} \\ &\leq \text{range}(K_{h,2}) \cdot \frac{(1 + 2^{-\kappa})(1 - 2^{-\kappa+1})}{1 - 2^{-\kappa}} \\ &< \text{range}(K_{h,2}). \end{aligned}$$

Thus,  $\mathcal{S}_1$  has at most  $s_1$  bits and  $\mathcal{S}_2$  has at most  $s_2$  bits. This proves the proposition.  $\square$

Finally, we show that the inverse of fusion can be done efficiently.

**Proposition 15 (restated).** *Let  $s_1, \dots, s_C \geq 0$ ,  $R_h, R_t \in [2^\kappa, 2^{\kappa+1}]$  and  $m \geq \kappa$ , let  $s = m + \lg R_h + \lg R_t$ , and  $s \leq \lg(2^{s_1} + \dots + 2^{s_C}) - C \cdot 2^{-\kappa+2}$ , there is a lookup table of size  $O(C)$ . Given a double-ended string  $\mathcal{S} = (K_h, M, K_t)$  such that  $\text{range}(K_h) = R_h$ ,  $\text{range}(K_t) = R_t$  and  $|M| = m$ , there is an extraction algorithm that generates a pair  $(i, \mathcal{S}_i)$  such that  $i \in \{1, \dots, C\}$ , and  $\mathcal{S}_i = (K_{i,h}, M_i, K_{i,t})$  has length at most  $s_i$ . Moreover,  $(M_i, K_{i,t})$  is a suffix of  $\mathcal{S}$ , and given  $i$  and  $K_{i,h}$ , the rest of  $\mathcal{S}$  (i.e.,  $\mathcal{S}[-1, |M| - |M_i| - 1]$ ) can be recovered in constant time, assuming random access to the lookup table.  $\text{range}(K_{i,h})$ ,  $\text{range}(K_{i,t})$  and  $|M_i|$  does not depend on  $\mathcal{S}$ , and can be stored in the lookup table.*

*Proof.* By setting  $K_{i,t} := K_t$ , the task becomes to encode  $(K_h, M)$  using  $(i, (K_{i,h}, M_i))$ . Next, we show how to determine  $i$ . To this end, we divide the range of  $K_h$  into  $C$  disjoint intervals  $\{[l_i, r_i)\}_{i=1, \dots, C}$ , such that the  $i$ -th interval has size at most

$$\lfloor 2^{s_i - |M| - \lg(\text{range}(K_t))} \rfloor.$$

Such division is possible, because

$$\sum_{i=1}^C \lfloor 2^{s_i - |M| - \lg(\text{range}(K_t))} \rfloor > \sum_{i=1}^C 2^{s_i - |M| - \lg(\text{range}(K_t))} - C$$

$$\begin{aligned}
&\geq 2^{-|M|-\lg(\text{range}(K_t))} \cdot 2^{s+(C-1)\cdot 2^{-\kappa+2}} - C \\
&\geq 2^{s-|M|-\lg(\text{range}(K_t))} \cdot (2^{(C-1)\cdot 2^{-\kappa+2}} - C \cdot 2^{-\kappa}) \\
&\geq \text{range}(K_h) \cdot (1 + (C-1)2^{-\kappa+1} - C \cdot 2^{-\kappa}) \\
&\geq \text{range}(K_h).
\end{aligned}$$

Fix one such division, e.g., the  $i$ -th interval is from

$$l_i := \min\{\lfloor 2^{s_1-|M|-\lg(\text{range}(K_t))} \rfloor + \dots + \lfloor 2^{s_{i-1}-|M|-\lg(\text{range}(K_t))} \rfloor, \text{range}(K_h) - 1\}$$

to

$$r_i := \min\{\lfloor 2^{s_1-|M|-\lg(\text{range}(K_t))} \rfloor + \dots + \lfloor 2^{s_i-|M|-\lg(\text{range}(K_t))} \rfloor, \text{range}(K_h) - 1\}$$

excluding the right endpoint. We store all endpoints  $l_i, r_i$  in the lookup table, taking  $O(C)$  space.

Now, find  $i$  such that  $K_h \in [l_i, r_i)$ . Then compute  $|M_i| = \lfloor s_i - \lg(\text{range}(K_t)) \rfloor - \kappa$ , and let

$$M_i := M[|M| - |M_i|, |M| - 1].$$

Finally, we view the first  $|M| - |M_i|$  bits of  $M$  as a nonnegative integer  $Z \in [2^{|M|-|M_i|}]$  and let

$$K_{i,h} := 2^{|M|-|M_i|} \cdot (K_h - l_i) + Z.$$

Observe that  $K_{i,h} < \lfloor 2^{\kappa+\text{frac}(s_i-\lg(\text{range}(K_t)))} \rfloor$ , because

$$\begin{aligned}
K_{i,h} &< 2^{|M|-|M_i|} \cdot (r_i - l_i) \\
&\leq 2^{|M|-(\lfloor s_i-\lg(\text{range}(K_t)) \rfloor - \kappa)} \cdot 2^{s_i-|M|-\lg(\text{range}(K_t))} \\
&\leq 2^{\kappa+\text{frac}(s_i-\lg(\text{range}(K_t)))}.
\end{aligned}$$

Thus, the length of  $\mathcal{S}_i = (K_{i,h}, M_i, K_{i,t})$  is at most

$$\lg(\text{range}(K_{i,h})) + |M_i| + \lg(\text{range}(K_{i,t})) \leq s_i.$$

We also store the sizes of  $\mathcal{S}_i$  for every  $i$  in the lookup table.

It is clear that  $(M_i, K_{i,t})$  is a suffix of  $\mathcal{S}$ . Given  $i$  and  $K_{i,h}$ , we retrieve  $l_i$  and  $M_i$  from the lookup table. Then  $\mathcal{S}[-1] = K_h$  can be recovered using

$$K_h = l_i + K_{i,h} \text{ div } 2^{|M|-|M_i|}.$$

Also,  $Z$  can be recovered using

$$Z = K_{i,h} \bmod 2^{|M|-|M_i|},$$

which determines  $\mathcal{S}[0, |M| - |M_i| - 1]$ . This proves the proposition.  $\square$

## B Approximating Binomial Coefficients

In this section, we prove Claim 22 and Claim 23 from Section 6.1.

**Claim 22 (restated).** *Both  $\text{OPT}_{(k-i+1)V_{\text{bl}},m_1} - (k-i+1)\text{SIZE}_{\text{main}} + (m_1-1)2^{-\kappa/2+2}$  and  $\text{OPT}_{(j-k)V_{\text{bl}},m_2} - (j-k)\text{SIZE}_{\text{main}} + (m_2-1)2^{-\kappa/2+2}$  can be approximated with an additive error of at most  $2^{-\kappa}$  in  $O(1)$  time.*

*Proof.* (sketch) For Claim 22, the goal is essentially to efficiently approximate

$$s_1 = \text{OPT}_{(k-i+1)V_{\text{bl}},m_1} - (k-i+1)\text{SIZE}_{\text{main}} + (m_1-1)2^{-\kappa/2+2}$$

and

$$s_2 = \text{OPT}_{(j-k)V_{\text{bl}},m_2} - (j-k)\text{SIZE}_{\text{main}} + (m_2-1)2^{-\kappa/2+2}.$$

To approximate  $s_1$  and  $s_2$ , we can store an approximation of  $\text{SIZE}_{\text{main}}$  up to  $O(\kappa)$  bits of precision in the lookup table. The task reduces to approximate the two OPTs. Recall that

$$\text{OPT}_{V,m} = \lg \binom{V}{m}.$$

The problem further reduces to approximate  $\lg \binom{(k-i+1)V_{\text{bl}}}{m_1}$  and  $\lg \binom{(j-k)V_{\text{bl}}}{m_2}$ . In the following, we show that for any given  $V, m \leq 2^\kappa$ , it is possible to approximate  $\lg \binom{V}{m}$  in  $O(1)$  time.

$\lg \binom{V}{m}$  can be expanded to  $\lg V! - \lg m! - \lg(V-m)!$ . We approximate each term separately. By Stirling's formula,

$$\ln k! = k \ln \left( \frac{k}{e} \right) + \frac{1}{2} \ln 2\pi n + \sum_{i=2}^d \frac{(-1)^i B_i}{i(i-1)k^{i-1}} + O(k^{-d}),$$

where  $B_i$  is the  $i$ -th Bernoulli number, and  $d \geq 2$ . For any constant  $\epsilon > 0$ , by setting  $d \geq \Omega(1/\epsilon)$ , the above approximation gives an error of  $2^{-\Omega(\kappa)}$  for any  $k \geq 2^{\epsilon\kappa}$ . We store the Bernoulli numbers in the lookup table, and the formula can be evaluated in constant time. On the other hand, for all  $k < 2^{\epsilon\kappa}$ , we simply store an approximation of  $\lg k!$  in a global lookup table, taking  $2^{\epsilon\kappa}$  size. Finally, by approximating  $\lg V!$ ,  $\lg m!$  and  $\lg(V-m)!$  independently with additive error  $2^{-2\kappa-2}$ , we obtain an estimation of  $\lg \binom{V}{m}$  with additive error smaller than  $2^{-2\kappa}$ . Note that each of the three values may be  $2^{\omega(\kappa)}$ , which takes super-constant words to store. However, since the final value is guaranteed to be at most  $2^\kappa$ , we could safely apply mod  $2^\kappa$  over the computation.  $\square$

**Claim 23 (restated).** *For any  $V_1, V_2, m \geq 0$ , and  $0 \leq l \leq m$ ,  $\sum_{i=0}^l 2^{\text{OPT}_{V_1,i} + \text{OPT}_{V_2,m-i}}$  can be approximated up to an additive error of at most  $2^{-\kappa-3} \cdot \sum_{i=0}^m 2^{\text{OPT}_{V_1,i} + \text{OPT}_{V_2,m-i}}$  in  $O(\kappa^5)$  time.*

*Proof.* (sketch) The goal is to approximate

$$\sum_{i=0}^l \binom{V_1}{i} \binom{V_2}{m-i}$$

up to additive error of  $2^{-\kappa-3} \cdot \binom{V_1+V_2}{m}$ , because

$$2^{\text{OPT}_{V_1,i} + \text{OPT}_{V_2,m-i}} = \binom{V_1}{i} \binom{V_2}{m-i}.$$

To this end, we shall use the following lemma from [Yu19] to approximate binomial coefficients.

**Lemma 40** ([Yu19]). *For any large integers  $V, d$  and  $0 < a \leq V/2$ , such that  $d \leq c \cdot a$ , there is a polynomial  $P$  of degree  $d$ , such that*

$$\binom{V}{a+x} \leq \binom{V}{a} \cdot \left(\frac{V-a}{a}\right)^x \cdot P_{V,d}(x) \leq \binom{V}{a+x} \cdot (1 + 2^{-\sqrt{d}+8}),$$

*for all integers  $x \in [0, c \cdot \sqrt{a}]$ , a (small) universal constant  $c > 0$ . Moreover, given  $V$  and  $d$ , the coefficients of  $P_{V,d}$  can be computed in  $O(d^{1.5})$  time.*

This lemma allows us to approximate  $\sum_{l=a}^b \binom{V_1}{l} \binom{V_2}{m-l}$  where  $b - a \leq c \cdot \sqrt{a}$ , up to a *multiplicative* error of  $1 \pm 2^{-2\kappa}$  in  $O(\kappa^4)$  time: it reduces approximating the sum to computing  $\sum_l \alpha^l \cdot P_1(l)P_2(l)$  for two degree- $O(\kappa^2)$  polynomials  $P_1, P_2$ .

Let  $\bar{m} = \frac{V_1}{V_1+V_2} \cdot m$ . For  $l < \bar{m} - 2\sqrt{\bar{m} \cdot \kappa}$ , we return 0 as the approximation; For  $\bar{m} - 2\sqrt{\bar{m} \cdot \kappa} \leq l \leq \bar{m} + 2\sqrt{\bar{m} \cdot \kappa}$ , we divide the range into chunks of size  $O(\sqrt{\bar{m}})$ , apply Lemma 40 to approximate  $\sum_l \binom{V_1}{l} \binom{V_2}{m-l}$  for each chunk in  $O(\kappa^4)$  time, and return the sum; For  $l > \bar{m} + 2\sqrt{\bar{m} \cdot \kappa}$ , we return (an approximation of)  $\binom{V_1+V_2}{m}$  as the estimation. It is not hard to verify that in all cases we return an approximation with desired error. The details are omitted.  $\square$

## C Dictionary with Linear Redundancy

In this section, we show a proof sketch of Lemma 26, and present a dictionary data structure that uses a linear number of extra bits. Recall that  $\text{OPT}_{V,m} := \lg \binom{V}{m}$ . For membership queries only, Pagh [Pag01a] already obtained a better data structure. The data structure in this section is a generalization of Pagh’s static dictionary.

**Lemma 26 (restated).** *Given a set  $S \subset [V]$  of  $m$  keys, there is a data structure of size*

$$\text{OPT}_{V,m} + O(m + \lg \lg V),$$

*such that it defines a bijection  $h$  between  $S$  and  $[m]$  and a bijection  $\bar{h}$  between  $[V] \setminus S$  and  $[V - m]$ . It supports hash queries in constant time.*

We are going to use Pagh’s static dictionary as a subroutine. For this reason, let us first give an overview of this data structure. The data structure uses a minimal perfect hashing of Schmidt and Siegel [SS90]. The hashing has three levels. In the first level, each key  $x$  is mapped to  $h_{k,p}(x) = (kx \bmod p) \bmod m^2$  with *no collisions*, for a prime  $p = \Theta(m^2 \lg V)$  and  $k \in [p]$ . A random pair  $(k, p)$  works with constant probability, and it takes  $O(\lg m + \lg \lg V)$  bits to encode the function. This level effectively reduces the universe size from  $V$  to  $m^2$ . Each key  $x \in S$  is then represented by a pair  $(x^{(1)}, x^{(2)})$  where  $x^{(1)} \in [m^2]$  is the hash value, and  $x^{(2)} = (x \div p) \cdot \lceil p/m^2 \rceil + (kx \bmod p) \div m^2$  (called the quotient function in [Pag01a]). Then  $x^{(2)} \leq O(V/m^2)$  and  $(x^{(1)}, x^{(2)})$  uniquely determines  $x$ .

In the second level, we apply another hash function from the same family on  $x^{(1)}$ ,  $h_{k',p'}(x^{(1)}) = (k'x^{(1)} \bmod p') \bmod m$  to map  $x^{(1)}$  to  $m$  buckets. This time, we have  $p' = \Theta(m^2)$  and  $k' \in [p']$ . Let  $A_i$  be the number of keys mapped to bucket  $i$ . The hashing guarantees that for a random pair  $(k', p')$ , the expectation of each  $A_i^2$  is bounded by  $O(1)$ . Similarly, we can represent  $x^{(1)}$  further as a pair such that the first component is the hash value in  $[m]$ , and the second component is the quotient function value, which is at most  $O(m)$ .

The third level hashing then hashes all keys in the same bucket to different integers. It is applied on  $x^{(1)}$ :  $g_{k_i, p_i}(x^{(1)}) = (k_i x^{(1)} \bmod p_i) \bmod A_i^2$ , for  $p_i = \Theta(m^2)$  and  $k_i \in [p_i]$  such that all keys in the bucket are mapped to different integers. It turns out that a random pair  $(k_i, p_i)$  works with constant probability.

The data structure stores the following for the hash functions:

1. the top-level hash functions  $(k, p)$  and  $(k', p')$ ,
2. a list of  $O(\lg m)$  (random) choices for the third-level hash functions  $(k_1, p_1), (k_2, p_2), \dots$ ,
3. for each bucket  $i$ , the index  $\pi_i$  of the first hash function in the list that works.

It turns out that it is possible to use only  $O(m)$  bits to store the indices  $\pi_i$ . This is because each second-level hash function works with constant probability, the entropy of each  $\pi_i$  is a constant. We can use the Huffman coding for each  $\pi_i$  to achieve constant bits per index (which turns out to be the unary representation of  $\pi_i$ ).

These hash functions map all  $m$  input keys to  $O(m)$  buckets with *no* collisions. By storing a rank data structure (e.g., [Př08]) among the  $O(m)$  buckets using  $O(m)$  bits of space, we further map all the non-empty buckets to  $[m]$ . Finally, we store for each bucket, the *quotient functions* of the input key mapped to it. Hence, it takes  $\lg(V/m^2) + \lg m + O(1) = \lg(V/m) + O(1)$  bits to encode each key. Thus, the total space is  $m \lg(V/m) + O(m + \lg \lg V) = \lg \binom{V}{m} + O(m + \lg \lg V)$  bits.

This data structure supports membership queries, and naturally defines a bijection  $h$  between  $S$  and  $[m]$ , namely  $h(x)$  simply being the bucket  $x$  is mapped to. To generalize the data structure and define an efficiently computable bijection  $\bar{h}$  between  $[V] \setminus S$  and  $[V - m]$ , we apply an approach similar to Section 6.2. To this end, we first store the number of keys  $m'$  in  $[V - m]$ . This is also the number of non-keys in  $\{V - m, \dots, V - 1\}$ . We are going to store a mapping that maps all  $m'$  non-keys in  $\{V - m, \dots, V - 1\}$  to all  $m'$  keys in  $[V - m]$ .

We then store the above data structure for all keys in  $[V - m]$ , using

$$m' \lg((V - m)/m') + O(m' + \lg \lg V) \leq \lg \binom{V}{m} + O(m + \lg \lg V)$$

bits, which defines a bijection  $h'$  between  $S \cap [V - m]$  and  $[m']$ . Note that this data structure also allows us to “randomly access” all keys. That is, given an index  $i \in [m']$ , it returns a key  $x_i$ , such that  $\{x_1, \dots, x_{m'}\}$  is the set of all  $m'$  keys in  $[V - m]$ . Then, we store a rank data structure for  $\{V - m, \dots, V - 1\}$ , such that given an  $x \in \{V - m, \dots, V - 1\}$ , the query algorithm returns if  $x$  is a key, as well as its rank over the set of keys (or non-keys). Hence, it maps all keys in  $\{V - m, \dots, V - 1\}$  to  $[m - m']$  and all non-keys to  $[m']$ . The total space is  $\text{OPT}_{V,m} + O(m + \lg \lg V)$ .

For each  $x \in S$ , we define  $h(x)$  as follows.

- if  $x < V - m$ , let  $h(x) := h'(x)$ ;
- if  $x \geq V - m$ , let  $h(x)$  be  $m' - 1$  plus the rank of  $x$  in  $S \cap \{V - m, \dots, V - 1\}$ .

For  $x \notin S$ , we define  $\bar{h}(x)$  as follows.

- if  $x < V - m$ , let  $h(x) := x$ ;
- if  $x \geq V - m$ , suppose the rank of  $x$  in  $\{V - m, \dots, V - 1\} \setminus S$  is  $i$ , then let  $h(x) := x_i$ .

Having stored the above data structures,  $h(x)$  or  $\bar{h}(x)$  can be computed in constant time.