



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Gan, J;Umboh, SW;Wang, H;Wirth, A;Zhang, Z

Title:

Optimal Dynamic Parameterized Subset Sampling

Date:

2024-11

Citation:

Gan, J., Umboh, S. W., Wang, H., Wirth, A. & Zhang, Z. (2024). Optimal Dynamic Parameterized Subset Sampling. *Proceedings of the ACM on Management of Data*, 2 (5), pp.1-26. <https://doi.org/10.1145/3695827>.

Persistent Link:

<https://hdl.handle.net/11343/365422>

License:

[CC BY](#)



DL Latest updates: <https://dl.acm.org/doi/10.1145/3695827>

Published: 07 November 2024

RESEARCH-ARTICLE

[Citation in BibTeX format](#)

## Optimal Dynamic Parameterized Subset Sampling

JUNHAO GAN, University of Melbourne, Melbourne, VIC, Australia

SEEUN WILLIAM UMBOH, University of Melbourne, Melbourne, VIC, Australia

HANZHI WANG, University of Copenhagen, Copenhagen, Hovedstaden, Denmark

ANTHONY IAN WIRTH, The University of Sydney, Sydney, NSW, Australia

ZHUO ZHANG, University of Melbourne, Melbourne, VIC, Australia

Open Access Support provided by:

University of Melbourne

The University of Sydney

University of Copenhagen

# Optimal Dynamic Parameterized Subset Sampling\*

JUNHAO GAN, The University of Melbourne, Australia

SEEUN WILLIAM UMBOH, The University of Melbourne, Australia and ARC Training Centre in Optimisation Technologies, Integrated Methodologies, and Applications (OPTIMA), Australia

HANZHI WANG<sup>†</sup>, BARC, University of Copenhagen, Denmark

ANTHONY WIRTH, The University of Sydney, Australia

ZHUO ZHANG, The University of Melbourne, Australia

In this paper, we study the *Dynamic Parameterized Subset Sampling* (DPSS) problem in the *Word RAM model*. In DPSS, the input is a set,  $S$ , of  $n$  items, where each item,  $x$ , has a *non-negative integer weight*,  $w(x)$ . Given a pair of query parameters,  $(\alpha, \beta)$ , each of which is a *non-negative rational number*, a *parameterized subset sampling query* on  $S$  seeks to return a subset  $T \subseteq S$  such that each item  $x \in S$  is selected in  $T$ , independently, with probability  $p_x(\alpha, \beta) = \min \left\{ \frac{w(x)}{\alpha \sum_{x \in S} w(x) + \beta}, 1 \right\}$ . More specifically, the DPSS problem is defined in a dynamic setting, where the item set,  $S$ , can be updated with insertions of new items or deletions of existing items.

Our first main result is an optimal algorithm for solving the DPSS problem, which achieves  $O(n)$  pre-processing time,  $O(1 + \mu_S(\alpha, \beta))$  expected time for each query parameterized by  $(\alpha, \beta)$ , *given on-the-fly*, and  $O(1)$  time for each update; here,  $\mu_S(\alpha, \beta)$  is the expected size of the query result. At all times, the worst-case space consumption of our algorithm is linear in the current number of items in  $S$ .

Our second main contribution is a hardness result for the DPSS problem when the item weights are  $O(1)$ -word *float numbers*, rather than integers. Specifically, we reduce Integer Sorting to the deletion-only DPSS problem with float item weights. Our reduction shows that an optimal algorithm for deletion-only DPSS with float item weights (achieving all the same bounds as aforementioned) implies a  $O(N)$ -expected-time algorithm for sorting  $N$  integers. The latter remains an important open problem. Moreover, a deletion-only DPSS algorithm which supports float item weights, with complexities worse, by at most a factor of  $o(\sqrt{\log \log N})$ , than the optimal counterparts, would already improve the current-best integer sorting algorithm [18].

Last but not least, a key technical ingredient for our first main result is a set of exact and efficient algorithms for generating Bernoulli (of certain forms) and Truncated Geometric random variates in  $O(1)$  expected time with  $O(n)$  worst-case space in the Word RAM model. Generating Bernoulli and geometric random variates efficiently is of great importance not only to sampling problems but also to encryption in cybersecurity. We believe that our new algorithms may be of independent interests for related research.

CCS Concepts: • **Theory of computation** → **Database theory; Proof complexity; Randomness, geometry and discrete structures**; • **Mathematics of computing** → **Probabilistic algorithms**.

\*The full version of this paper can be found at <http://arxiv.org/abs/2409.18036>

<sup>†</sup>Work partially done at Renmin University of China and The University of Melbourne.

---

Authors' Contact Information: Junhao Gan, The University of Melbourne, Melbourne, Australia, [junhao.gan@unimelb.edu.au](mailto:junhao.gan@unimelb.edu.au); Seeun William Umboh, The University of Melbourne, Melbourne, Australia and ARC Training Centre in Optimisation Technologies, Integrated Methodologies, and Applications (OPTIMA), Melbourne, Australia, [william.umboh@unimelb.edu.au](mailto:william.umboh@unimelb.edu.au); Hanzhi Wang, BARC, University of Copenhagen, Copenhagen, Denmark, [hzwang.helen@gmail.com](mailto:hzwang.helen@gmail.com); Anthony Wirth, The University of Sydney, Sydney, Australia, [anthony.wirth@sydney.edu.au](mailto:anthony.wirth@sydney.edu.au); Zhuo Zhang, The University of Melbourne, Melbourne, Australia, [zhuo.zhang@student.unimelb.edu.au](mailto:zhuo.zhang@student.unimelb.edu.au).

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2836-6573/2024/11-ART209

<https://doi.org/10.1145/3695827>

Additional Key Words and Phrases: Subset Sampling, Word RAM Model, Exact Random Variate Generation

### ACM Reference Format:

Junhao Gan, Seoun William Umboh, Hanzhi Wang, Anthony Wirth, and Zhuo Zhang. 2024. Optimal Dynamic Parameterized Subset Sampling. *Proc. ACM Manag. Data* 2, 5 (PODS), Article 209 (November 2024), 26 pages. <https://doi.org/10.1145/3695827>

## 1 Introduction

Sampling is arguably one of the most important techniques for solving problems on massive data. It can effectively reduce the problem size while still providing reliable estimation for certain metrics and largely retaining the statistical properties of the original data. As such, sampling has been widely adopted for various problems in different research fields, including anomaly detection [21, 25, 33], feature selection [23, 28], sketching algorithms [26], network measurement [20, 29, 30], distributed computing [9, 11], cybersecurity [3, 22], machine learning [17, 24], and computational geometry [8, 19]. In general, sampling methods fall into two categories: *Weighted Sampling* and *Subset Sampling*. In *Weighted Sampling*, the input is a set of  $n$  items,  $S$ , and each item,  $x$ , has a non-negative weight,  $w(x)$ . It aims to sample *just one* item such that each item  $x$  is chosen with probability  $\frac{w(x)}{\sum_{x \in S} w(x)}$ . On the other hand, in the vanilla setting of *Subset Sampling*, the input is a set  $S$  of  $n$  items, each with a *probability*,  $p(x)$ , the goal is to sample a *subset*  $T \subseteq S$  such that each item  $x$  has probability  $p(x)$  to be selected in  $T$ , independently.

A natural extension of *Subset Sampling* is to consider the *dynamic* setting, where the item set  $S$  can be updated by either insertions of new items or deletions of existing ones. This is called the *Dynamic Subset Sampling* (DSS) problem. Yi et al. [32] proposed an *optimal* algorithm, called ODSS, for solving the DSS problem in a special Real RAM model<sup>1</sup> which allows both logarithm and rounding operations can be performed in  $O(1)$  time. Specifically, their ODSS algorithm achieves  $O(n)$  preprocessing time,  $O(1 + \mu)$  expected query time, where  $\mu$  is the expected size of the sampling result, and supports each update in  $O(1)$  time.

**Dynamic Parameterized Subset Sampling.** In this paper, we study a more general setting of the *Subset Sampling* problem, called *Dynamic Parameterized Subset Sampling* (DPSS), in the Word RAM model. The DPSS problem is defined as follows. Consider a set  $S$  of  $n$  items, which is subject to updates by either item insertions or deletions; each item  $x \in S$  has a *non-negative integer* weight  $w(x)$ . Each *Parameterized Subset Sampling* (PSS) query is associated with a pair of non-negative rational parameters  $(\alpha, \beta)$ , and asks for a subset  $T \subseteq S$  such that each item  $x \in S$  is selected in  $T$ , independently, with probability  $p_x(\alpha, \beta) = \min \left\{ \frac{w(x)}{\alpha \sum_{x \in S} w(x) + \beta}, 1 \right\}$ .

*Distinctions between DPSS and DSS.* Comparing to the aforementioned DSS problem, there are several crucial distinctions from DPSS. First, in DSS, the sampling probability of each item is known in advance, while in DPSS, the sampling probabilities of the items are parameterized by the query parameters  $(\alpha, \beta)$  *on the fly* during query time. Second, in DSS, the sampling probabilities of the items are *decoupled*, in the sense that updates on the sampling probability of items would not affect each other. However, in contrast, in DPSS, the sampling probabilities of the items are *inter-correlated* such that an update to an item's weight would affect the sampling probabilities of all other items. In other words, the impact of each update to the item set  $S$  in DPSS is a lot larger. These distinctions, therefore, make DPSS more challenging than the DSS problem. To see this, the existing optimal ODSS algorithm requires  $\Omega(n)$  time to support an update in the DPSS setup even just with fixed and known query parameters  $(\alpha, \beta)$  at all times.

<sup>1</sup>Although Yi et al. [32] claimed that their algorithm can work on the Word RAM model, the technical details of the generation of Truncated Geometric random variates have been overlooked.

Applications of DPSS. Subset sampling is a versatile algorithmic tool that has found multifaceted applications in influence maximization [15, 16], local clustering [29], graph neural networks [29], bipartite matching [5], and computational epidemiology [13], among many others. Several of these applications align well with the DPSS setting, where the item sampling probabilities may simultaneously change after a single update, further illustrating that current algorithms (such as DSS algorithms) are insufficient to handle such cases. For example, when solving influence maximization [15, 16, 34] on *dynamic* networks, computing the reverse reachable (RR) sets by subset sampling can be formalized as instances of DPSS. In the local clustering [2, 27, 29, 31] scenario, Wang et al. [29] uses subset sampling techniques to accelerate the computation of random-walk probabilities on static unweighted graphs. When this idea is applied to *dynamic weighted networks*, efficient DPSS algorithms would be useful in this scenario. Further discussions on the applications of DPSS can be found in the full version [1] of this paper.

**Main Result 1: An Optimal DPSS Algorithm.** Our first contribution is an optimal algorithm for solving the DPSS problem in the Word RAM model. The main result is summarized as follows:

**THEOREM 1.1.** *Consider a set  $S$  of  $n$  items; there exists a data structure for solving the Dynamic Parameterized Subset Sampling (DPSS) problem, which achieves:*

- *PRE-PROCESSING TIME:*  $O(n)$  worst-case time to construct the data structure on  $S$ ;
- *QUERY TIME:*  $O(1 + \mu_S(\alpha, \beta))$  expected time to return a sample for each Parameterized Subset Sampling (PSS) query on  $S$  with parameters  $(\alpha, \beta)$ , where  $\mu_S(\alpha, \beta)$  is the expected size of the sample;
- *UPDATE TIME:*  $O(1)$  worst-case time for each item insertion or deletion.

*At all times, the worst-case space consumption of such a data structure is bounded by  $O(n)$ , where  $n$  denotes the current cardinality of  $S$ .*

An Overview of Our Algorithm. Given a pair of Parameterized Subset Sampling (PSS) query parameters,  $(\alpha, \beta)$ , the crucial idea of our algorithm is to convert the PSS query on  $S$  into a  $O(1)$  number of PSS query instances (with possibly different query parameters on possibly different item sets), of which the PSS results can be converted back into a feasible PSS result on  $S$  by *rejection sampling*. And each of these  $O(1)$  query instances can be solved optimally in  $O(1 + \mu)$  expected time, where  $\mu$  is its expected sample size. To achieve such a conversion, and to solve those converted PSS query instances optimally, we propose a novel data structure, called *Hierarchy + Adapter + Lookup Table* (HALT). As its name suggests, it consists of three main components: (i) a three-level sampling hierarchy, (ii) a set of dynamic adapters, and (iii) a static lookup table.

Specifically, the sampling hierarchy converts the PSS query on  $S$  into  $O(1)$  PSS query instances on other item sets, while the lookup table can answer PSS queries on a special family of item sets. Last, but not least, as the lookup table is static, while the PSS queries are parameterized on the fly, our dynamic adapters play a crucial role to bridge the hierarchy and the lookup table. Specifically, the adapters, dynamically translate the PSS query instances, during query time, into a form for which our lookup table is applicable.

**Main Result 2: A Hardness Result on DPSS with Float Item Weights.** While our HALT algorithm is optimal for DPSS with integer item weights, interestingly, if the item weights are floating-point numbers, each represented by a  $O(1)$ -word exponent and a  $O(1)$ -word mantissa, the DPSS problem suddenly becomes a lot more difficult. In fact, this difficulty arises even in the *deletion-only* setting where the updates are deletions only. We demonstrate this by showing that DPSS can be used to obtain an algorithm for Integer Sorting:

**THEOREM 1.2.** *Suppose there exists an algorithm for deletion-only DPSS on a set of  $N$  items with float weights that has pre-processing time  $t_p(N)$ , expected query time  $(1 + \mu) \cdot t_q(N)$ , and deletion*

time  $t_{del}(N)$ . Then the Integer Sorting problem for  $N$  integers can be solved in  $t_p(N) + O(N \cdot (t_q(N) + t_{del}(N)))$  expected time on the  $d$ -bit Word RAM model with  $d \in \Omega(\log N)$ .

In particular, this implies that if there is an optimal algorithm for deletion-only DPSS—i.e., one with  $t_p(N) \in O(N)$ , and  $t_q(N), t_{del}(N) \in O(1)$ —then sorting  $N$  integers (each represented with  $O(1)$  words) can be achieved in  $O(N)$  expected running time on the  $d$ -bit Word RAM model with  $d \in \Omega(\log N)$ . However, solving the Integer Sorting problem in  $O(N)$  expected time for every word-length  $d \in \Omega(\log N)$  remains an open problem [4]. Moreover, an algorithm for deletion-only DPSS (with float item weights) with  $t_p(N) \in o(N \cdot \sqrt{\log \log N})$ , and  $t_q(N), t_{del}(N) \in o(\sqrt{\log \log N})$  already improves the current-best  $O(N \cdot \sqrt{\log \log N})$  expected running time bound of Han and Thorup [18].

**Main Result 3: An Efficient Algorithm for Truncated Geometric Variate Generation.** Our third contribution is an efficient algorithm for generating *Truncated Geometric* random variates in the Word RAM model. Specifically, for  $p \in (0, 1)$  and a positive integer  $n$ , a truncated geometric distribution,  $\text{T-Geo}(p, n)$ , returns a value  $i \in \{1, \dots, n\}$  with probability  $\frac{p(1-p)^{i-1}}{1-(1-p)^n}$ . As can be seen from the literature, c.f., the aforementioned ODSS algorithm and from our HALT algorithm, generating truncated geometric random variates is of great importance and a crucial building block in many sampling-related algorithms. Somewhat surprisingly, given its importance, we are not aware of any previous work on generating truncated geometric variates in  $O(1)$  expected time with  $O(n)$  worst-case space in the Word RAM model. We give the following theorem to make it documented:

**THEOREM 1.3.** *Let  $p \in (0, 1)$  be a rational number whose numerator and denominator fit in  $O(1)$  words, and  $n$  be a positive integer fitting in  $O(1)$  words. There exists an algorithm that generates a random variate from  $\text{T-Geo}(p, n)$  in  $O(1)$  expected time and  $O(n)$  space in the worst case in the Word RAM model.*

Considering the importance of efficient geometric random variate generation to many research fields, we believe that our new algorithm is of independent interest.

The proofs of some statements (indicated by [\*]) have been deferred to the appendix. Meanwhile, the pseudocodes of our DPSS algorithms can be found in the appendix.

## 2 Preliminaries

In this section, we first introduce the *Word RAM model* [12] and then formally define the problem of *Dynamic Parameterized Subset Sampling* (DPSS).

### 2.1 The Word RAM Model

In the Word RAM model [12], each *word* consists of  $d$  bits. Each memory address is represented by  $O(1)$  words and hence, there are at most  $2^{O(d)}$  memory addresses, each corresponding to a unique memory cell. Each memory cell stores a word of content which can be considered as an integer in the range of  $\{0, 1, \dots, 2^d - 1\}$ . The content of each memory cell can be accessed by its memory address in  $O(1)$  time. Furthermore, given two one-word integers  $a$  and  $b$ , each of the following basic arithmetic operations on  $a$  and  $b$  can be performed in  $O(1)$  time: addition ( $a + b$ ), subtraction ( $a - b$ ), multiplication ( $a \cdot b$ ), division with rounding ( $\lfloor a/b \rfloor$ ), comparison ( $>, <, =$ ), truncation, and bit operations (e.g., finding the index of the highest or lowest non-zero bit in  $a$ ). Moreover, every *long integer* is represented by an array of words, and every *float number* is represented by its exponent and mantissa, where each takes  $O(1)$  words. Finally, as a convention, we assume that a uniformly random word of  $d$  bits can be generated in  $O(1)$  time. Throughout this paper, the space consumption of an algorithm is measured in words. The following fact about the Word RAM model is useful; a detailed implementation can be found in the full version [1].

**FACT 1.** *In the  $d$ -bit Word RAM model, consider the integer universe  $U = \{0, 1, \dots, d - 1\}$ , and a set  $I \subseteq U$  of  $n$  integers; there exists a data structure that maintains all the integers in  $I$  in a sorted linked list and supports each of the following operations in  $O(1)$  worst-case time:*

- *each update (either an insertion of a new integer or a deletion of an existing integer) to  $I$ ;*
- *finding the predecessor or successor for any integer  $q \in U$  in  $I$ .*

*Such a data structure consumes  $O(n)$  space at all times, where  $n$  is the current cardinality of  $I$ .*

## 2.2 Dynamic Parameterized Subset Sampling

Consider a set of  $n$  items  $S = \{x_1, \dots, x_n\}$ ; each item has a *weight*  $w(x_i)$  which is a *non-negative* integer. Without loss of generality, we assume that the largest possible value of  $n$ , denoted by  $n_{\max}$  and the largest possible value of the item weights, denoted by  $w_{\max}$  can be represented with one word (of  $d$  bits) in the Word RAM model. Therefore,  $d \in \Omega(\log(n_{\max} \cdot w_{\max}))$  holds all the time.

Define  $W_S(\alpha, \beta) = \alpha \cdot \sum_{x \in S} w(x) + \beta$  as the *parameterized total weight function* of all the items in  $S$  by two parameters  $\alpha$  and  $\beta$ , where both  $\alpha$  and  $\beta$  are *non-negative rational numbers* and each can be represented by a  $O(1)$ -word numerator and a  $O(1)$ -word denominator. Thus, the parameterized total weight  $W_S(\alpha, \beta)$  of  $S$  fits in  $O(1)$  words. Furthermore, if the value of  $\sum_{x \in S} w(x)$  is pre-computed and maintained, the value of  $W_S(\alpha, \beta)$  can be computed in  $O(1)$  time.

Given a pair of parameters  $(\alpha, \beta)$ , a *parameterized subset sampling query* on  $S$  asks for a subset  $T \subseteq S$  such that each item  $x \in S$  is selected in  $T$  independently, with probability  $p_x(\alpha, \beta) = \min\left\{\frac{w(x)}{W_S(\alpha, \beta)}, 1\right\}$ . Moreover, the expected size of  $T$  is defined as  $\mu_S(\alpha, \beta) = \sum_{x \in S} p_x(\alpha, \beta)$ . Observe that since the parameters  $(\alpha, \beta)$  are given on-the-fly in a parameterized subset sampling query, the sampling probability of each item can be vastly different with different parameters  $(\alpha, \beta)$ .

In this paper, we study the problem of parameterized subset sampling on a *dynamic setting*, where the item set  $S$  can be updated by insertions of new items (associated with a weight) and deletions of existing items from  $S$ . We call this problem the *Dynamic Parameterized Subset Sampling* (DPSS) problem. It is worth mentioning that in the DPSS problem, the sampling probability of each item is not only affected by the parameters  $(\alpha, \beta)$ , but also by the updates on  $S$ . To see this, consider a fixed pair  $(\alpha, \beta)$ ; before and after the insertion of an item with a very large weight, the value of  $W_S(\alpha, \beta)$  can change significantly and hence, such an update dramatically changes the sampling probabilities of all the items.

Our goal is to design a data structure for solving the DPSS problem such that: (i) it can be constructed in  $O(n)$  time and consumes  $O(|S|)$  space at all times, (ii) it can answer every parameterized subset sample query in  $O(1 + \mu_S(\alpha, \beta))$  expected time, and (iii) it supports each update to  $S$  in  $O(1)$  worst-case time.

## 3 Sampling Random Variates in the Word RAM Model

In this section, we discuss how to exactly and efficiently sample, in the Word RAM model, the key random variates in our optimal DPSS algorithm. Due to the importance of random variate generation, we believe that these new algorithms would be of independent interest to other problems.

As we see in the next section, our algorithm needs to generate a number of variants of *Bernoulli* and *Geometric* random variates. We explore each in turn.

### 3.1 Bernoulli Random Variate Generation in the Word RAM Model

Consider a real number  $p \in (0, 1)$ , the *Bernoulli distribution* parameterized by  $p$ , denoted by  $\text{Ber}(p)$ , takes values in  $\{0, 1\}$  with probability  $\Pr[\text{Ber}(p) = 1] = 1 - \Pr[\text{Ber}(p) = 0] = p$ . In particular, our algorithm needs to generate the following three types of Bernoulli random variates,  $\text{Ber}(p)$ :

**type (i)**  $p$  is a *rational number* with numerator and denominator both fitting in  $O(1)$  words;

**type (ii)**  $p = p^*$ , where  $p^* = \frac{1-(1-q)^n}{n \cdot q}$ , where  $q$  is a rational number represented by a pair of  $O(1)$ -word numerator and denominator,  $n$  is a  $O(1)$ -word integer, and  $n \cdot q \leq 1$ ;

**type (iii)**  $p = \frac{1}{2p^*}$ , with  $p^*$  as in type (ii).

For type (i), Bringmann and Friedrich [7] proved the following result:

**FACT 2 ([7]).** *If  $p$  is a rational number with numerator and denominator being  $O(1)$ -word integers, then  $\text{Ber}(p)$  can be evaluated in  $O(1)$  expected running time, and  $O(1)$  worst-case space.*

To the best of our knowledge, there is no previous work showing how to generate a realization of  $\text{Ber}(p)$  of either type (ii) or type (iii) in  $O(1)$  expected time with  $O(n)$  worst-case space. However, as we see in Section 4, these two types of Bernoulli random variates play a crucial role in our optimal DPSS algorithm. We prove the following theorem:

**THEOREM 3.1.** *Consider a Bernoulli random variate,  $\text{Ber}(p)$ , in the form of either  $p = p^*$  or  $p = \frac{1}{2p^*}$ , with  $p^* = \frac{1-(1-q)^n}{n \cdot q}$ , where  $q$  is a rational number represented by a pair of  $O(1)$ -word numerator and denominator,  $n$  is a  $O(1)$ -word integer, and  $n \cdot q \leq 1$ ;  $\text{Ber}(p)$  can be generated in  $O(1)$  expected time with  $O(n)$  worst-case space.*

**Definition 3.2 (i-Bit Approximation).** Consider a value  $p$ , possibly comprising an infinite sequence of bits. An  $i$ -bit approximation of  $p$ , denoted by  $\tilde{p}_i$ , is an approximation value of the  $i$  most-significant bits in  $p$ , i.e., satisfies  $|\tilde{p}_i - p| \leq 2^{-i}$ , while  $\tilde{p}_i$  itself consists of at most  $i + 1$  bits.

We prove Theorem 3.1 via the following algorithmic framework of Bringmann and Friedrich [7]:

**FACT 3 ([7, 10]).** *Consider a Bernoulli random variate,  $\text{Ber}(p)$  with rational parameter  $p \in (0, 1)$ . Suppose the numerator and the denominator of  $p$  both can be represented with  $O(n)$  words, and  $p$  can be evaluated in  $O(\text{poly}(n))$  time and the  $i$ -bit approximation,  $\tilde{p}_i$ , of  $p$  can be evaluated in  $O(\text{poly}(i))$  time. Then  $\text{Ber}(p)$  can be generated in  $O(1)$  expected time with  $O(n)$  worst-case space in the Word RAM model.*

**LEMMA 3.3. [\*]** *These two claims hold simultaneously: (i) the numerator and the denominator of  $p = p^*$  both can be represented with  $O(n)$  words and  $p$  can be evaluated in  $O(\text{poly}(n))$  time; and (ii) there exists an algorithm that can evaluate each  $i$ -bit approximation  $\tilde{p}_i$  of  $p = p^*$ , satisfying all conditions stated in Theorem 3.1, in  $O(\text{poly}(i))$  time.*

**LEMMA 3.4. [\*]** *These two claims hold simultaneously: (i) the numerator and the denominator of  $p = \frac{1}{2p^*}$  both can be represented with  $O(n)$  words and  $p$  can be evaluated in  $O(\text{poly}(n))$  time; and (ii) there exists an algorithm that can evaluate each  $i$ -bit approximation  $\tilde{p}_i$  of  $p = \frac{1}{2p^*}$ , satisfying all conditions stated in Theorem 3.1, in  $O(\text{poly}(i))$  time.*

### 3.2 Geometric Random Variate Generation in the Word RAM Model

Consider a parameter  $p \in (0, 1)$ . The *Geometric distribution*  $\text{Geo}(p)$  takes discrete values in  $\mathbb{N}$ . Specifically,  $\Pr[\text{Geo}(p) = i] = p(1-p)^{i-1}$ , for  $i \in \mathbb{N}$ . As the value of a geometric variate can be arbitrarily large (though it is not really likely), thus in theory, it is impossible to generate variates from a geometric distribution with worst-case bounded space.

**Bounded Geometric Variates.** Nonetheless, in many application scenarios, including our optimal DPSS algorithm, geometric random variates within a bounded range are sufficient. This is known as the *Bounded Geometric distribution*. Given an integer  $n \in \mathbb{N}$ , the bounded geometric distribution, denoted by  $\text{B-Geo}(p, n)$ , is defined as  $\min\{n, \text{Geo}(p)\}$  with distribution:

$$\Pr[\text{B-Geo}(p, n) = i] = \begin{cases} p(1-p)^{i-1} & i \in \{1, \dots, n-1\}; \\ (1-p)^{n-1} & i = n. \end{cases}$$

Bringmann and Friedrich [7] studied the generation of bounded geometric random variates in the Word RAM model and gave the following theorem:

**FACT 4 ([7]).** *Let  $p \in (0, 1)$  be a rational number with numerator and denominator fitting in  $O(1)$  words and  $n$  be a positive integer fitting in  $O(1)$  words. A random variate from  $\text{B-Geo}(p, n) = \min\{n, \text{Geo}(p)\}$  can be realized in  $O(1)$  expected time and  $O(n)$  space in the worst case.*

**Truncated Geometric Variates.** Besides the Bounded Geometric variates, the so-called *Truncated Geometric* random variates are also of great importance to many applications. As we see in Section 4, it is a crucial building block in our optimal DPSS algorithm. Given  $p \in (0, 1)$  and  $n \in \mathbb{N}$ , the truncated geometric distribution, denoted by  $\text{T-Geo}(p, n)$ , defined on  $\{1, \dots, n\}$ , has distribution:

$$\Pr[\text{T-Geo}(p, n) = i] = \frac{p(1-p)^{i-1}}{1 - (1-p)^n}, \text{ for } i \in \{1, \dots, n\}.$$

The probability function of truncated geometric distribution can be interpreted as a conditional probability as follows. Consider sampling each item from a bucket of  $n$  items with probability  $p$  independently. Then  $\text{T-Geo}(p, n)$  is identically distributed as the smallest index of the sampled items, conditioned on at least one item being sampled.

**Proof of Theorem 1.3.** We consider the following possible cases:

**Case 1 ( $n \leq 2$ ):** This is the trivial case to generate  $i \sim \text{T-Geo}(p, n)$ . If  $n = 1$ , then  $\text{T-Geo}(p, n) = 1$  always holds. If  $n = 2$ , it can be verified that  $\text{T-Geo}(p, n)$  returns either  $i = 1$  with probability  $\frac{1}{2-p}$  or  $i = 2$  with probability  $\frac{1-p}{2-p}$ . In this case,  $\text{T-Geo}(p, n)$  is identically distributed as  $\text{Ber}(\frac{1-p}{2-p}) + 1$ . By Fact 2, this can be achieved in  $O(1)$  expected running time with  $O(1)$  worst-case space. Thus, Theorem 1.3 holds.

**Case 2.1 ( $n \geq 3$  and  $np \geq 1$ ):** In this case, an  $i \sim \text{T-Geo}(p, n)$  can be realized by repeatedly generating  $i \sim \text{B-Geo}(p, n+1)$  until  $i \leq n$  holds. In each generation,  $i \leq n$  happens with probability  $1 - (1-p)^n > 1 - \frac{1}{e}$  because  $np \geq 1$ . Thus,  $i$  is realized from  $\text{T-Geo}(p, n)$  in  $O(1)$  trials in expectation. By Fact 4, Theorem 1.3 holds.

**Case 2.2 ( $n \geq 3$  and  $np < 1$ ):** To handle this case, we propose an algorithm which runs as follows:

- $i \leftarrow 0$ ;
- while  $i \leq n$  do:
  - $i \leftarrow i + \text{B-Geo}(\frac{2}{n}, n+1)$ ;
  - if  $i \leq n$  and  $\text{Ber}((1-p)^{i-1}) = 1$ 
    - \* generate  $c \leftarrow \text{Ber}(\frac{1}{2p^*})$ , with  $p^* = \frac{1-(1-p)^n}{np}$ ;
    - \* if  $c = 1$ , **return**  $i$  as a realization of a random variate from  $\text{T-Geo}(p, n)$ ;
- start over from the first step with  $i \leftarrow 0$ ;

To complete the proof of Theorem 1.3, we show the correctness and running time cost of the above algorithm.

**Correctness.** Observe that the  $\text{B-Geo}(\frac{2}{n}, n+1)$  generation in the while-loop simulates the subset sampling process of choosing  $i$ , for each index  $i$ , independently with probability  $2/n$ . Thus, each  $i$  has probability  $2/n$  of being sampled in this step. Next,  $i$  is provisionally accepted with probability  $(1-p)^{i-1}$  via a Bernoulli trial, and  $i$  is further eventually accepted with probability  $\frac{1}{2p^*}$ . Therefore, combining all these probabilities, we have:

$$\Pr[i \text{ is returned as a } \text{T-Geo}(p, n) \text{ variate}] = \frac{2}{n} \cdot (1-p)^{i-1} \cdot \frac{1}{2 \cdot \frac{1-(1-p)^n}{np}} = \frac{p(1-p)^{i-1}}{1 - (1-p)^n}.$$

Hence, our algorithm for  $np < 1$  case correctly realizes a T-Geo( $p, n$ ) random variate.

The Running Time. First, it can be verified that, B-Geo( $2/n, n + 1$ ) returns  $i \leq n$  within  $O(1)$  trials in expectation. Next, we prove that, given an  $i$ , our algorithm accepts  $i$  with at least constant probability. To see this, given an  $i$ , it is accepted by the first Bernoulli trial with probability  $(1 - p)^{i-1} \geq (1 - p)^n \geq (1 - \frac{1}{n})^n \geq \frac{8}{27}$ , where the second last inequality comes from the fact that  $pn < 1$ , while the last inequality is by the fact that  $n \geq 3$  and the function  $(1 - \frac{1}{n})^n$  monotonically increases with  $n$ . Furthermore, for the second Bernoulli trial,  $i$  is accepted with probability  $\frac{1}{2p^*} \geq \frac{1}{2}$ , because  $p^* = \frac{1-(1-p)^n}{np} \leq 1$ . Therefore, in expectation, within  $O(1)$  trials in total, an  $i$  will be returned by our algorithm. Finally, by Facts 2 and 4, and by our Theorem 3.1, all the random variates used in each trial can be generated in  $O(1)$  expected time with  $O(n)$  worst-case space.

## 4 Our Optimal DPSS Algorithm

In this section, we show our optimal dynamic parameterized subset sampling algorithm. For ease of exposition, and based on the results in Section 3, we note that the evaluation of every random variate used in our algorithm can be performed in  $O(1)$  expected time with  $O(n)$  worst-case space.

### 4.1 A One-Level Bucket-Grouping Structure

We first introduce a one-level *Bucket-Grouping Structure* (BG-Str) which is a key technique to our *three-level sampling hierarchy*, and which reveals the intuitions and the correctness of our algorithm. In the following, we deliberately use  $X = \{x_1, \dots, x_N\}$  to denote the item set, to distinguish it from the *original* input item set,  $S$ : the BG-Str will be applied recursively to construct the sampling hierarchy. To avoid the complicated *floor* and *ceiling* operations on the logarithmic values, without loss of generality, we assume that the size,  $N$ , of an item set is always a *power-of-16* number. Otherwise, we can *conceptually* pad with *dummy* items, each with *zero* weight, to make  $N$  a power-of-16 number. This would only increase  $N$  by a constant factor of 16 in the worst case, and hence, would not affect our theoretical results. Thanks to this, the values of  $\log_2 N$ ,  $\log_2 \log_2 N$  and  $\log_2 \log_2 \log_2 N$  are all integers.

**The Data Structure.** Consider a PSS instance  $X = \{x_1, \dots, x_N\}$ , the one-level bucket-grouping structure on  $X$ , denoted by BG-Str( $X$ ), can be constructed in four steps:

- Step 1: Total Weight Calculation. Compute  $W_X = \sum_{x \in X} w(x)$ . Observe that  $W_X$  is just a sum of  $N$  one-word integers, and hence,  $W_X$  can be stored in  $O(1)$  words.
- Step 2: Item Bucketing. Assign the items in  $X$  to *buckets* by their weights. Bucket  $B_X(i)$  is assigned all the items  $x \in X$  such that  $2^i \leq w(x) < 2^{i+1}$ , for  $i = 0, 1, \dots, \lfloor \log_2 w_{\max} \rfloor$ , where  $w_{\max}$  is the largest possible weight of the items. All the *non-empty* buckets (containing at least one item from  $X$ ) are maintained in a *sorted* linked list by their bucket index,  $i$ , in ascending order.
- Step 3: Bucket Grouping. Organize all the buckets into *groups*, where the group with index,  $j$ , denoted by  $G_X(j)$ , is a collection of all the non-empty buckets  $B_X(i)$  whose bucket indices  $i$  satisfy  $\lfloor \frac{i}{\log_2 N} \rfloor = j$ , where the largest possible value of a group index,  $j$ , is  $j_{\max} = \lfloor \frac{\lfloor \log_2 w_{\max} \rfloor}{\log_2 N} \rfloor$ . It is worth mentioning that each group corresponds to exactly  $\log_2 N$  consecutive possible bucket indices. Analogous to the buckets, all the *non-empty* groups (containing at least one non-empty bucket) are maintained in a *sorted* linked list by their group index,  $j$ , in ascending order.
- Step 4: Next-Level Instance Construction. For each non-empty group,  $G_X(j)$ , construct a *next-level PSS instance* on a set,  $Y_j$ , of *next-level* items by the following steps:
  - for each non-empty bucket  $B_X(i) \in G_X(j)$ , create a *next-level item*, denoted by  $y_i$ , with weight  $w(y_i) = 2^{i+1} \cdot |B_X(i)|$ .

- denote the set of these next-level items by  $Y_j$ , and the PSS instance on  $Y_j$  is called the *next-level instance* of  $G_X(j)$ .

LEMMA 4.1. *The one-level bucket-grouping structure,  $\text{BG-Str}(X)$ , can be constructed in  $O(N)$  worst-case time with  $O(N)$  space consumption.*

PROOF. The construction of  $\text{BG-Str}(X)$  mainly involves the constructions and maintenance of *sorted linked lists* of non-empty buckets (and non-empty groups) of length at most  $\lfloor \log_2 w_{\max} \rfloor \leq d$ . According to Fact 1 in the Word RAM model, each of them can be constructed in *linear* time with *linear* space to its size. As the sum of all their sizes is bounded by  $O(N)$  (because they are non-empty), the overall construction cost and space consumption of them are bounded by  $O(N)$ .  $\square$

**Answering a PSS query on  $X$  with  $\text{BG-Str}(X)$ .** Consider the bucket-grouping structure  $\text{BG-Str}(X)$ , constructed on the item set  $X$ . Given a PSS query with parameters  $(\alpha, \beta)$ , the buckets  $B_X(i)$  can be categorized into three types. It is worth mentioning that  $\text{BG-Str}(X)$  itself does not have the concept of the query parameters  $(\alpha, \beta)$  when constructed. All the following concepts are defined *dynamically on the fly* during query time by the given parameters  $(\alpha, \beta)$ .

- $(\alpha, \beta)$ -*Insignificant Buckets*.  $B_X(i)$  is an  $(\alpha, \beta)$ -*insignificant bucket* if its bucket index  $i$  satisfies  $\frac{2^{i+1}}{W_X(\alpha, \beta)} \leq \frac{1}{N^2}$ , where recall that  $W_X(\alpha, \beta) = \alpha \cdot W_X + \beta$ . Since  $B_X(i)$  contains items with weights in the range  $[2^i, 2^{i+1})$ , each item  $x$  in a non-empty  $(\alpha, \beta)$ -insignificant bucket  $B_X(i)$  will have probability  $p_x(\alpha, \beta) < \frac{1}{N^2}$  to be selected in the sampling result.
- $(\alpha, \beta)$ -*Certain Buckets*.  $B_X(i)$  is an  $(\alpha, \beta)$ -*certain bucket* if its bucket index  $i$  satisfies  $\frac{2^i}{W_X(\alpha, \beta)} \geq 1$ . As a result, each item in a non-empty  $(\alpha, \beta)$ -certain bucket  $B_X(i)$  is included in the sampling result with probability 1.
- $(\alpha, \beta)$ -*Significant Buckets*. If  $B_X(i)$  is neither of the above two types, then  $B_X(i)$  is an  $(\alpha, \beta)$ -significant bucket. By definition, each item in such a non-empty bucket  $B_X(i)$  will have probability in the range of  $(\frac{1}{2N^2}, 1)$  to be sampled.

Moreover, if all the possible buckets (not necessarily non-empty) in a group  $G_X(j)$  are  $(\alpha, \beta)$ -insignificant (resp.,  $(\alpha, \beta)$ -certain), then  $G_X(j)$  is an  $(\alpha, \beta)$ -*insignificant group* (resp.,  $(\alpha, \beta)$ -*certain group*). If neither of these cases holds,  $G_X(j)$  is a  $(\alpha, \beta)$ -*significant group*. Observe that for every group, the difference between the largest and smallest bucket indices is at most  $\log_2 N - 1$ . Thus, no group can simultaneously contain an  $(\alpha, \beta)$ -certain bucket and an  $(\alpha, \beta)$ -insignificant bucket.

LEMMA 4.2. *Given a PSS query with parameters  $(\alpha, \beta)$ , there can be at most three  $(\alpha, \beta)$ -significant groups in  $\text{BG-Str}(X)$ .*

PROOF. By definition, each  $(\alpha, \beta)$ -significant group contains at least one possible  $(\alpha, \beta)$ -significant bucket. Furthermore, observe that the union of the probability ranges of all possible  $(\alpha, \beta)$ -significant buckets is  $(\frac{1}{2N^2}, 1)$ . As a result, there can be at most  $1 + 2 \log_2 N$  such buckets. Since each group contains  $\log_2 N$  possible buckets with consecutive bucket indices, therefore, all these  $(\alpha, \beta)$ -significant buckets must be contained in at most three  $(\alpha, \beta)$ -significant groups.  $\square$

**The Query Algorithm (Algorithm 1).** The basic idea is to answer a PSS query with parameters  $(\alpha, \beta)$  on  $X$  by solving the following three types of PSS query instances:

- *insignificant instance*: one PSS query with parameters  $(\alpha, \beta)$  on the set of all the items from  $X$  that are in the buckets in all the  $(\alpha, \beta)$ -insignificant groups;
- *certain instance*: one PSS query with parameters  $(\alpha, \beta)$  on the set of all the items from  $X$  that are in the buckets in all the  $(\alpha, \beta)$ -certain groups;

- *next-level instances*: one PSS query with parameters  $(0, W_X(\alpha, \beta))$  on the next-level item set  $Y_j$  for each  $(\alpha, \beta)$ -significant group  $G_X(j)$ . As each sampled next-level item in  $Y_j$  is essentially a bucket in  $G_X(j)$ , the items of  $X$  in the bucket can be sampled with *rejection sampling*.

The detailed steps of the query algorithm are shown in Algorithm 1.

---

**Algorithm 1:** One-Level PSS Query Algorithm: Query(BG-Str( $X$ ),  $(\alpha, \beta)$ ,  $\ell_{\text{cur}}$ )

---

**Data:** a Bucket-Grouping Structure BG-Str( $X$ ), PSS query parameters  $(\alpha, \beta)$ , the current recursion level  $\ell_{\text{cur}}$

**Result:** a PSS result, with parameters  $(\alpha, \beta)$ ,  $T \subseteq X$

```

1  $N \leftarrow |X|$ ;
2  $j_1 \leftarrow \left\lfloor \frac{\lceil \log_2 \frac{W_X(\alpha, \beta)}{N^2} \rceil - 1}{\log_2 N} \right\rfloor$ , the maximum possible  $(\alpha, \beta)$ -insignificant group index;
3  $T_1 \leftarrow \text{QueryInsignificant}(\text{BG-Str}(X), (\alpha, \beta), j_1 \cdot \log_2 N)$  // Algorithm 2;
4  $j_2 \leftarrow \lfloor \frac{\lceil \log_2 W_X(\alpha, \beta) \rceil}{\log_2 N} \rfloor$ , the minimum possible  $(\alpha, \beta)$ -certain group index;
5  $T_2 \leftarrow \text{QueryCertain}(\text{BG-Str}(X), (\alpha, \beta), j_2 \cdot \log_2 N)$  // Algorithm 3;
6  $T_3 \leftarrow \emptyset$ ;
7 for  $j$  from  $j_1 + 1$  to  $j_2 - 1$  do
8    $T'_3 \leftarrow \text{QueryASignificantGroup}(\text{BG-Str}(X), (\alpha, \beta), j, \ell_{\text{cur}})$  // Algorithm 4;
9    $T_3 \leftarrow T_3 \cup T'_3$ ;
10  $T \leftarrow T_1 \cup T_2 \cup T_3$ ;
11 return  $T$ ;
```

---

The query algorithm handles these three different types of instances with Algorithms 2, 3, and 4, respectively. The main methodology in our algorithm is *rejection sampling*. Specifically, for each item  $x$  with sampling probability  $p_x(\alpha, \beta)$ , we first sample  $x$  as a *potential item* with some probability  $p' \geq p_x(\alpha, \beta)$ , and then, we accept this item  $x$  as a PSS result with probability  $\frac{p_x(\alpha, \beta)}{p'}$  independently. In this way, item  $x$  will be sampled with probability  $p_x(\alpha, \beta)$  independently as required in the PSS query.

Next, we claim that that all the mathematical computations involved in Algorithm 1 can be evaluated in  $O(1)$  time in the Word RAM model. To prove this, it suffices to show the following Claim 1 whose proof can be found in the full version [1]. As a result, the largest index  $j_1$  of the insignificant groups and the smallest index  $j_2$  of the certain groups can be computed in  $O(1)$  time in the Word RAM model, as  $W_X(\alpha, \beta)$  is a rational number that satisfies the conditions in Claim 1.

CLAIM 1 (FULL VERSION [1]). *Consider a rational number  $x > 0$  with its numerator and denominator being  $O(1)$ -word integers, then  $\lfloor \log_2(x) \rfloor$  and  $\lceil \log_2(x) \rceil$  can be calculated in  $O(1)$  time.*

*Handling the Insignificant Instance (Algorithm 2).* By definition, each item  $x$  in the buckets in all those  $(\alpha, \beta)$ -insignificant groups has probability at most  $\frac{1}{N^2}$  to be sampled. Thus, we first generate a *bounded geometric variate*  $k \sim \text{B-Geo}(\frac{1}{N^2}, N + 1)$  to simulate the process of flipping a Bernoulli coin  $\text{Ber}(\frac{1}{N^2})$  for each item, one by one, and  $k$  is the index of the first item to be sampled. Clearly, if  $k > N$ , then nothing will be sampled and thus done. Otherwise, Algorithm 2 first collects all the items in this insignificant instance and stores them in an array  $A$ . If  $k > |A|$ , then it is done. But if  $k \leq |A|$ , the item  $x$  at the  $k^{\text{th}}$  index in  $A$ , i.e.,  $A[k]$ , is a potential item, and it is accepted in the sampling result with probability  $\frac{p_x(\alpha, \beta)}{1/N^2}$ . Thus, this item  $x$  is equivalently sampled with probability

$p_x(\alpha, \beta)$ . For each of the remaining items  $x$ , at index starting from  $k + 1$  in the array  $A$ , Algorithm 2 flips a Bernoulli coin with probability  $p_x(\alpha, \beta)$  to accept  $x$ . Therefore, each item  $x$  in the array  $A$  will be correctly sampled, independently, with probability  $p_x(\alpha, \beta)$ .

In terms of query efficiency, recall that the bounded geometric variate  $k \sim \text{B-Geo}(\frac{1}{N^2}, N + 1)$  can be generated in  $O(1)$  expected time. If  $k > N$ , the running time cost is  $O(1)$ . Otherwise, the process takes  $O(N)$  time. So the overall expected running time is bounded, thus leading to Lemma 4.3.

$$\Pr[k \leq N] \cdot O(N) + \Pr[k > N] \cdot O(1) \leq \left(1 - \left(1 - \frac{1}{N^2}\right)^N\right) \cdot O(N) + O(1) \leq \frac{O(N)}{N} + O(1) = O(1).$$

LEMMA 4.3. *Algorithm 2 returns a correct PSS result with parameters  $(\alpha, \beta)$  on the set of items in the insignificant instance in  $O(1)$  expected time.*

Handling the Certain Instance (Algorithm 3). For the certain instance, each item  $x$  has a sampling probability  $p_x(\alpha, \beta) = 1$ . Therefore, Algorithm 3 simply collects and returns all these items in the corresponding buckets, giving the following lemma.

LEMMA 4.4. *Algorithm 3 returns a correct PSS result with parameters  $(\alpha, \beta)$  on the set of items in the certain instance in  $O(1 + \ell)$  time, where  $\ell$  is the number of items sampled.*

Handling Each Significant Group (Algorithm 4). We now discuss the details for handling each  $(\alpha, \beta)$ -significant group  $G_X(j)$ . The core idea of Algorithm 4 is to first obtain a subset  $T_{Y_j}$  of the next-level item set  $Y_j$  such that each next-level item  $y_i \in Y_j$  is sampled independently with probability  $\min\left\{1, \frac{w(y_i)}{W_X(\alpha, \beta)}\right\}$ . To achieve this, Algorithm 4 invokes a PSS query with parameters  $(0, W_X(\alpha, \beta))$  on  $Y_j$ . As each next-level item  $y_i$  corresponds to a bucket  $B_X(i)$ ,  $T_{Y_j}$  is treated as a list of the buckets corresponding to its next-level items. And each bucket  $B_X(i) \in T_{Y_j}$  is called a *candidate bucket* because it has a chance to contain a *potential item*  $x \in B_X(i)$  which is sampled with probability of  $p = \min\left\{1, \frac{2^{i+1}}{W_X(\alpha, \beta)}\right\}$ . If a candidate bucket  $B_X(i)$  is confirmed via rejection sampling to have at least a potential item in it, then  $B_X(i)$  is called a *promising bucket*.

The next step is to extract items from the candidate bucket. As this step will be reused later, we summarize the steps in Algorithm 5. Specifically, Algorithm 5 accepts each candidate bucket  $B_X(i) \in T_{Y_j}$  as a promising bucket by the following steps, depending on two cases. Denote by  $n_i = |B_X(i)|$  the number of items in  $B_X(i)$ .

**Case 1** ( $p \cdot n_i \geq 1$ ): In this case, bucket  $B_X(i)$  indeed has probability of 1 to be included in  $T_{Y_j}$ , i.e.,  $\Pr[B_X(i) \in T_{Y_j}] = 1$ . Algorithm 5 generates a random variate  $k \sim \text{B-Geo}(p, n_i + 1)$ . If  $k > n_i$ , which happens with probability  $q = (1 - p)^{n_i}$ , then reject  $B_X(i)$  because  $\Pr[B_X(i) \in T_{Y_j}] \cdot q = (1 - p)^{n_i}$  is exactly the probability that  $B_X(i)$  has no potential items. Otherwise,  $B_X(i)$  is accepted as a promising bucket and the  $k^{\text{th}}$  item  $x$  in  $B_X(i)$  is a potential item sampled with probability  $p$ .

**Case 2** ( $p \cdot n_i < 1$ ): In this case,  $\Pr[B_X(i) \in T_{Y_j}] = p \cdot n_i$ . Algorithm 5 generates a Bernoulli coin  $c \sim \text{Ber}\left(\frac{1 - (1 - p)^{n_i}}{p \cdot n_i}\right)$ . If  $c = 0$ , then reject  $B_X(i)$ . Otherwise,  $B_X(i)$  is accepted as a promising bucket, observing that this event happens with probability  $\Pr[B_X(i) \in T_{Y_j}] \cdot \Pr[c = 1] = 1 - (1 - p)^{n_i}$ , which is exactly the probability of  $B_X(i)$  containing at least one potential item. However, the most tricky aspect is to find the first potential item,  $x$ , in  $B_X(i)$  conditioned on the fact that  $B_X(i)$  is a promising bucket, such that  $x$  is still equivalently sampled with probability of  $p$ . To overcome this, Algorithm 5 generates a *truncated geometric* random variate  $k \sim \text{T-Geo}(p, n_i)$ , where the truncated geometric distribution guarantees to return an index  $k$  such that (i)  $k \leq n_i$ , and (ii)  $\Pr[k = j] = \frac{p(1-p)^{j-1}}{1 - (1-p)^{n_i}}$ . As discussed in Section 3, the value of  $k$  has the same distribution as the smallest sampled index

in  $[1, n_i]$ , each sampled with probability  $p$ , conditioned on that at least one index being sampled. Therefore, the  $k^{\text{th}}$  item in  $B_X(i)$  has a probability of  $p$  to be sampled as a potential item.

As a result, in either case, if  $B_X(i)$  is confirmed as a promising bucket, Algorithm 5 always gives the index  $k$  of the first potential item  $x$  in  $B_X(i)$  with correct probability,  $p$ . Algorithm 5 further generates a Bernoulli coin  $c \sim \text{Ber}\left(\frac{p_X(\alpha, \beta)}{p}\right)$  (Line 12) to accept  $x$  and add it to the PSS result,  $T_3$ . Then the algorithm repeatedly generates the index of the next potential item with bounded geometric distribution  $\text{B-Geo}(p, n_i + 1)$  and performs the same rejection sampling process until the index is  $> n_i$ .

**LEMMA 4.5.** *Algorithm 4 returns a correct PSS result with parameters  $(\alpha, \beta)$  on the set of items in the given  $(\alpha, \beta)$ -significant group  $G_X(j)$ .*

**PROOF.** The correctness of Algorithm 4 follows from the discussion above. That is, each bucket  $B_X(i) \in Y_j$  has a probability of  $1 - (1 - p)^{n_i}$  to be accepted as a promising bucket independently; and each item in a promising bucket  $B_X(i)$  has correct probability  $p$  to be identified as a potential item independently, where  $n_i = |B_X(i)|$  and  $p = \min\{1, \frac{2^{i+1}}{W_X(\alpha, \beta)}\}$ . Furthermore, each potential item,  $x$ , is accepted in the returned sampling result independently with probability  $\frac{p_X(\alpha, \beta)}{p}$ . Therefore, each item is sampled in the returned sampling result independently with probability  $p_X(\alpha, \beta)$ .  $\square$

**THEOREM 4.6 (CORRECTNESS).** *Given a PSS query with parameters  $(\alpha, \beta)$  on  $X$ , let  $T \subseteq X$  be the sampling result returned by Algorithm 1. Every item  $x \in X$  is selected in  $T$  independently with probability  $p_X(\alpha, \beta)$ .*

**PROOF.** This theorem follows immediately from Lemmas 4.3, 4.4 and 4.5.  $\square$

**THEOREM 4.7 (QUERY COMPLEXITY).** *Suppose that every PSS query with parameters  $(0, W_X(\alpha, \beta))$  on the next-level item set  $Y_j$  of a  $(\alpha, \beta)$ -significant group can be answered in  $O(1 + \mu_{Y_j})$  time in expectation, where  $\mu_{Y_j}$  is the expected size of the sampling result on  $Y_j$ . The overall expected running time of Algorithm 1 is bounded by  $O(1 + \mu_X(\alpha, \beta))$ .*

**PROOF.** Observe that Algorithm 1 invokes Algorithm 2 and Algorithm 3 each once, and Algorithm 4 at most three times according to Lemma 4.2. Moreover, by Lemma 4.3, the expected running time of Algorithm 2 for the insignificant instance is  $O(1)$ . And by Lemma 4.4, Algorithm 3 takes  $O(1 + \ell)$  time, where  $\ell$  is the number of items in the certain instance. Therefore, it suffices to show that the expected running time of Algorithm 4 for each significant group  $G_X(j)$  is bounded by  $O(1 + \mu)$ , where  $\mu$  is the expected size of the PSS result on the items in  $G_X(j)$  with parameters  $(\alpha, \beta)$ .

First, by the assumption in the theorem statement, the list  $T_{Y_j}$  of candidate buckets can be obtained in  $O(1 + \mu_{Y_j})$  expected time, where  $\mu_{Y_j}$  is equivalent to the expected number of candidate buckets, as discussed earlier. Consider the following claim.

**CLAIM 2.** *Each candidate bucket  $B_X(i) \in T_{Y_j}$  has  $\Omega(1)$  probability to be accepted as a promising.*

If Claim 2 holds, then  $\mu_{Y_j}$  can be bounded by  $O(\mu_{\text{prom}})$ , where  $\mu_{\text{prom}}$  is the expected number of promising buckets. Moreover, by definition, each promising bucket  $B_X(i)$  contains at least one potential item sampled with probability  $p = \min\{1, \frac{2^{i+1}}{W_X(\alpha, \beta)}\}$ . Thus,  $\mu_{\text{prom}}$  is bounded by the expected number of potential items. Furthermore, each potential item  $x \in B_X(i)$  is accepted to be in the sampling result with probability  $\frac{p_X(\alpha, \beta)}{p} \geq \frac{1}{2}$ . This is because, by the fact that  $x \in B_X(i)$ , we have  $w(x) \in [2^i, 2^{i+1})$ . Therefore,  $\mu_{Y_j}$  is then bounded by the expected number of promising buckets, then bounded by the expected number of potential items, and thus, bounded by the expected size

of the PSS result,  $O(\mu)$ . And hence, all the processing costs on the candidate buckets, promising buckets and the potential items can be charged to the output cost, in expectation.

Now, we **prove** Claim 2. Consider a candidate bucket  $B_X(i) \in T_{Y_j}$ ; let  $n_i = |B_X(i)|$  and  $p = \min\{1, \frac{2^{i+1}}{W_X(\alpha, \beta)}\}$ . Recall that Algorithm 4 considers two cases:

**Case 1** ( $p \cdot n_i \geq 1$ ): In this case,  $B_X(i)$  is rejected with probability  $(1-p)^{n_i} \leq (1-\frac{1}{n_i})^{n_i} \leq \frac{1}{e}$ . Thus,  $B_X(i)$  is accepted as a promising bucket with probability at least  $1 - \frac{1}{e} \in \Omega(1)$ .

**Case 2** ( $p \cdot n_i < 1$ ): Recall that, in this case,  $B_X(i)$  is accepted with probability  $\frac{1-(1-p)^{n_i}}{p \cdot n_i}$ . We claim that this probability is at least  $1 - \frac{1}{e}$ . To see this, consider the function  $f(a, b) = \frac{1-(1-b)^{a/b}}{a}$  for  $a = p \cdot n_i < 1$  and  $b = p \in (0, 1)$ . It can be verified that  $f(a, b)$  decreases monotonically with respect to  $a$  by analyzing the partial derivative of  $f$  with respect to  $a$ . The detailed analysis can be found in the full version [1]. Thus, we have  $f(a, b) > 1 - (1-b)^{1/b} > 1 - \frac{1}{e} \in \Omega(1)$ . As a result,  $B_X(i)$  is accepted as a promising bucket with probability at least  $\Omega(1)$ .

Combining the above two cases, Claim 2 holds. Moreover, our results on random variate generation in Section 3 imply that all the random variates used in Algorithm 4 can be generated in  $O(1)$  expected time. Putting all the above together, Theorem 4.7 follows.  $\square$

Theorem 4.7 implies that, as long as we can answer the PSS queries optimally on the next-level instances, we can answer the PSS query on  $X$  optimally. Moreover, observe that the next-level item set size is at most  $\log_2 N$ , logarithmic in the size of  $X$ . This motivates an idea of recursively applying the same bucket-grouping technique on next-level item set  $Y_j$  of each group  $G_X(j)$ . When the next-level instance is “small” enough, we use a *lookup table* to answer the PSS queries on them optimally. While this idea looks natural, certain non-trivial technical challenges remain to be addressed. First, by Lemma 4.2, each recursion level increases the number of next-level instances by a factor of 3. Thus, the recursion cannot be so deep that it exceeds  $O(1)$  recursion levels. Otherwise, a super-constant factor would be introduced to the query time complexity. Second, having only  $O(1)$  recursion levels does not reduce the problem size to  $O(1)$ . Thus, it is not trivial to design a lookup table with space consumption bounded by  $O(N)$ . Third, the next-level instances are *dynamic*, depending on the query parameters  $(\alpha, \beta)$ . However, the lookup table is *static*. As a result, reducing the next-level instances to problems where the lookup table is applicable during the query time is a challenging task.

## 4.2 The Three-Level Sampling Hierarchy

Given a set  $S$  of  $n$  items, the three-level sampling hierarchy of  $S$  is constructed as follows:

- **Level-1:** Construct  $\text{BG-Str}(S)$ ;
- **Level-2:** For each group  $G_S(j)$  in  $\text{BG-Str}(S)$ :
  - construct  $\text{BG-Str}(Y_j)$  on the next-level item set  $Y_j$  of  $G_S(j)$ ;
  - denote each bucket with bucket index  $i$  by  $B_{Y_j}(i)$ , and each group with group index  $k$  by  $G_{Y_j}(k)$ ;
  - denote the next-level item corresponding to bucket  $B_{Y_j}(i)$  by  $z_i$  and the next-level item set of group  $G_{Y_j}(k)$  by  $Z_k$ .
- **Level-3:** For each group  $G_{Y_j}(k)$  in  $\text{BG-Str}(Y_j)$ ,
  - \* construct  $\text{BG-Str}(Z_k)$  on the next level item set  $Z_k$  of  $G_{Y_j}(k)$ ;
  - \* denote each bucket with bucket index  $i$  by  $B_{Z_k}(i)$  and its corresponding next-level item by  $v_i$ ;
  - \* denote the set of all the next-level items of the buckets by  $V_k$ ; the PSS instance on  $Z_k$  is called a *final-level instance*.

**LEMMA 4.8.** *Given a set  $S$  of  $n$  items, the three-level sampling hierarchy can be constructed in  $O(n)$  time with  $O(n)$  space consumption.*




Fig. 1. Visualization of the three-level sampling hierarchy. Here, circles represent items and ovals represent buckets. The diagram illustrates the structure of a level-1 group  $G_X(j)$ , where each bucket in  $B_X(i)$  corresponds to an item  $y_i \in Y_j$ . The  $\text{BG-Str}(Y_j)$  contains multiple buckets in multiple groups. Each group, e.g.  $G_{Y_j}(k)$  corresponds to a next-level item set, e.g.  $Z_k$ . In level-3, each item set  $Z_k$  generates one final-level instance  $V_k$ , in which each item corresponds to a bucket in level-3.

**PROOF.** This lemma follows from Lemma 4.1 and the facts that, (i) at each level, the total number of the (next-level) items is at most  $n$ , and (ii) there are only three levels.  $\square$

**The Overall Query Algorithm.** To perform a PSS query with parameters  $(\alpha, \beta)$  on  $S$ , the overall query algorithm is as simple as invoking Algorithm 1,  $\text{Query}(\text{BG-Str}(S), (\alpha, \beta), \ell_{\text{cur}} = 1)$ .

**LEMMA 4.9 (CORRECTNESS).** *Every PSS query with parameters  $(\alpha, \beta)$  on  $S$  can be correctly answered with the three-level sampling hierarchy.*

**PROOF.** Follows immediately by applying Theorem 4.6 inductively.  $\square$

**LEMMA 4.10 (QUERY COMPLEXITY).** *If there exists an algorithm that can solve every final-level instance on  $Z_k$  optimally in  $O(1 + \mu_{Z_k})$  expected time, then the PSS query with parameters  $(\alpha, \beta)$  on  $S$  can be answered optimally in  $O(1 + \mu_S(\alpha, \beta))$  expected time.*

**PROOF.** First, by Lemma 4.2, there are at most 3 next-level instances on  $Y_j$ 's at level-1. By the same lemma, each of these instances on  $Y_j$  can produce at most 3 next-level instances on  $Z_k$ 's at level-2. And each instance on  $Z_k$  corresponds to a final-level instance  $V_k$ . Therefore, there are in total at most 9 final-level instances.

Second, by mathematical induction from the final level, with Theorem 4.7, the query complexity of the one-level algorithm, and by the fact that there are at most 9 final-level instances, the overall expected query time complexity follows.  $\square$

Thus, the problem reduces to solving the final-level PSS instances optimally. Our approach is to use a lookup table. However, the basic idea of a lookup table is to *hard code* all the possible sampling results for all possible *input configurations* for a special family of *static subset sampling* problems. Therefore, three main challenges must be addressed. The first challenge is to *efficiently*

construct a suitable static subset sampling instance from a final-level PSS instance, where the latter is dynamic and parameterized on the fly by  $(\alpha, \beta)$ . The second challenge is how to “encode” the configuration of static subset sampling instance such that it can be evaluated and located in the lookup table in  $O(1)$  time. Third, the overall space consumption of both the lookup table and the “encoder” must be bounded by  $O(n)$ . Next, we introduce the remaining two components in our HALT structure, i.e., the *adapters* and the *lookup table*, to address these technical challenges.

### 4.3 The Lookup Table in HALT

In order to set up the context of the goal which the adapters aim to achieve, we first introduce the lookup table in our HALT structure.

Define  $n_0$  be the size of the item set  $S$  when the HALT structure is constructed. Our lookup table aims to solve a *special static subset sampling* (4S) problem in the following form:

- the input item set  $V$  contains *exactly*  $K$  items  $\{v_1, v_2, \dots, v_K\}$ , where  $K = 2 \cdot \log_2 \log_2 \log_2 n_0$ ;
- the sampling probability of the  $i^{\text{th}}$  item  $v_i$  in  $V$  is in the form of  $p_i = \min\{1, \frac{2^{i+1} \cdot c_i}{m^2}\}$ , where  $c_i$  is an integer in the range of  $[0, m]$  and  $m$  is an integer bounded by  $O(\log \log n_0)$ ; as a result, the probability of each item must be an integer multiple of  $\frac{1}{m^2}$ .

Therefore, every possible input to the above 4S problem can be uniquely specified by a  $K$ -dimensional vector  $\vec{c} = (c_1, c_2, \dots, c_K)$ , where each  $c_i$  is an integer in  $[0, m]$ . Each  $\vec{c}$  is called an *input configuration* of the 4S problem. The total number of possible input configurations is  $(m + 1)^K$ .

Consider a fixed input configuration  $\vec{c}$ . Every possible subset sampling result can be encoded by a  $K$ -bit string  $r$ , where the  $i^{\text{th}}$  bit  $r[i] = 1$  indicates the item  $v_i$  is in the sampling result and  $r[i] = 0$  means  $v_i$  is not. Thus, the probability of obtaining a subset sampling result  $r$  can be calculated as:

$$\Pr(r) = \prod_{i=1}^K (r[i] \cdot p_i + (1 - r[i]) \cdot (1 - p_i)) .$$

Since  $p_i$  is an integer multiple of  $\frac{1}{m^2}$ ,  $\Pr(r)$  is an integer multiple of  $\frac{1}{(m^2)^K}$ . Furthermore, observe that there are exactly  $2^K$  possible subset sampling results,  $r$ . Summing  $\Pr(r)$  over all  $r$ 's, we must have  $\sum_{r \in \{0,1\}^K} \Pr(r) = 1$ .

Therefore, all the possible subset sampling results for a fixed input configuration,  $\vec{c}$ , can be encoded in an array  $A(\vec{c})$  of  $(m^2)^K$  cells and each cell is a  $K$ -bit string. Specifically, each subset sampling result  $r$  is stored in exactly  $\Pr(r) \cdot (m^2)^K$  cells in the array  $A(\vec{c})$ . As a result, given an input configuration  $\vec{c}$  to the 4S problem, a subset sampling result for  $\vec{c}$  can be obtained by picking a cell in  $A(\vec{c})$  uniformly at random and returning the  $K$ -bit string  $r$  as the subset sampling result. Clearly, each subset sampling result  $r$  will be returned with probability exactly  $\Pr(r)$  independently.

Figure 2 shows the detailed structure of the lookup table.

**The Lookup Table and Query Algorithm.** Our lookup table  $\mathcal{T}$  contains  $(m + 1)^K$  rows, where each row is an array  $A(\vec{c})$  corresponding to an input configuration  $\vec{c}$ . Given an input configuration  $\vec{c}$ , the query algorithm works as follows:

- locate the corresponding row  $A(\vec{c})$  in  $\mathcal{T}$  of  $\vec{c}$ ;
- uniformly at random select an index in  $A(\vec{c})$  and return the  $K$ -bit string in the cell at the index as the subset sampling result.

LEMMA 4.11. *Each input configuration  $\vec{c}$  can be represented in  $O(1)$  words.*

PROOF. Since each  $c_i$  is an integer in  $[0, m]$ , it can be encoded in  $\lceil \log_2(m + 1) \rceil$  bits. Thus,  $\vec{c}$  can be encoded in  $K \cdot \lceil \log_2(m + 1) \rceil \in O(\log n_0) \subseteq O(d)$  bits, equivalently,  $O(1)$  words, where  $K \in O(\log \log \log n_0)$  and  $m \in O(\log \log n_0)$ .  $\square$




Fig. 2. Visualization of a lookup table. It contains  $(m+1)^K$  rows, each row corresponds to a configuration, and each row contains  $(m^2)^K$  cells. Each cell holds a  $K$ -bit string, e.g.  $r_0, r_1, \dots$ . As shown in the figure, for a given configuration  $\vec{c} = \{c_1, c_2, \dots, c_K\}$ , if  $\Pr(r_0) = \frac{6}{(m^2)^K}$ , then there are 6 cells storing  $r_0$  in the corresponding row.

LEMMA 4.12. *The lookup table  $\mathcal{T}$  answers each subset sampling query for each configuration  $\vec{c}$  for the 4S problem in  $O(1 + \mu(\vec{c}))$  expected time, where  $\mu(\vec{c})$  is the expected size of the subset sample.*

PROOF. Since each  $\vec{c}$  can be encoded in  $O(1)$  words, locating the row  $A(\vec{c})$  in the lookup table  $\mathcal{T}$  can be performed in  $O(1)$  time. By our assumption at the beginning of this section, generating an index of  $A(\vec{c})$  uniformly at random can be achieved in  $O(1)$  expected time. Retrieving the subset sampling result from the  $K$ -bit string  $r$  can be performed in  $O(1 + |r|)$  time in the Word RAM model, where  $|r|$  is the number of items in the sampling result. Therefore, the overall query time is bounded by  $O(1 + \mu(\vec{c}))$  in expectation.  $\square$

LEMMA 4.13. *The lookup table  $\mathcal{T}$  can be constructed in  $O(n_0)$  time and the space consumption of  $\mathcal{T}$  is bounded by  $O(n_0)$ .*

PROOF. For an input configuration  $\vec{c}$ , there are  $2^K$  possible subset sampling results  $r$ . The probability of each  $r$  can be computed in  $O(K)$  time. Furthermore,  $A(\vec{c})$  contains  $m^{2K}$  cells and each cell stores  $K$  bits. The total running time for constructing  $A(\vec{c})$  is thus bounded by  $O(K \cdot 2^K + K \cdot m^{2K}) = O(K \cdot m^{2K})$  time. Summing over all  $(m+1)^K$  rows, the overall construction time of  $\mathcal{T}$  is bounded by  $O((m+1)^K \cdot K \cdot m^{2K})$  time. As  $K \in O(\log \log \log n_0)$  and  $m \in O(\log \log n_0)$ , this construction time is thus bounded by  $O(n_0)$ . In terms of space consumption,  $\mathcal{T}$  contains in total  $(m+1)^K \cdot K \cdot m^{2K}$  bits, which is bounded by  $O(n_0)$  words.  $\square$

#### 4.4 Bridging the Hierarchy and the Lookup Table via Adapters

It remains to show how to convert a final-level PSS query instance into an input configuration for the 4S problem in  $O(1)$  time on the fly during query time. Once this is accomplished, then Lemma 4.10 implies that the PSS query  $(\alpha, \beta)$  on  $S$  can be answered optimally.

For this purpose, we introduce an *adapter* for each final-level PSS query instance. And the space consumption of each such adapter is bounded by  $O(1)$  words. As there are at most  $O(n_0)$  final-level PSS query instances, the total space consumption of the adapters is bounded by  $O(n_0)$ . Recall that  $n_0$  is the size of the item set  $S$  when the HALT structure is constructed, and  $n$  is the *current* size of  $S$ . Without loss of generality, we assume that  $n_0/2 \leq n \leq 2n_0$ . If this is not the case, we can rebuild the HALT structure on the current  $S$  to make this assumption hold. More details about global rebuilding and de-amortization are discussed in Section 4.5.

Recall that in the three-level hierarchy, at Level-1, the next-level item set  $Y_j$  of each group  $G_S(j)$  in  $\text{BG-Str}(S)$  contains at most  $\log_2 n_0$  items; and at Level-2, the next-level item set  $Z_k$  of each group  $G_{Y_j}(k)$  in  $\text{BG-Str}(Y_j)$  has size at most  $\log_2 \log_2 n_0$ . Furthermore, at Level-3, each item  $v_i$  in the item set  $V_k$  in the final-level instance corresponds to a non-empty bucket,  $B_{Z_k}(i)$ , of items in  $Z_k$ , where each item  $z \in Z_k$  corresponds to a bucket in group  $G_{Y_j}(k)$ . Thus, we have  $|V_k| \leq |Z_k| \leq m$ , where  $m = \log_2 \log_2 n_0$ . And the weight of each item  $v_i \in V_k$  is  $w(v_i) = 2^{i+1} \cdot |B_{Z_k}(i)| \leq 2^{i+1} \cdot m$ .

**A Simple but Space-Inefficient Adapter Implementation.** Define  $m = \log_2 \log_2 n_0$ . Consider the final-level instance on  $Z_k$  as aforementioned. A simple implementation of an adapter is to maintain an array  $\mathcal{A}$  of length  $L$ , where  $L$  is number of possible bucket indices of the items in  $Z_k$ , and for  $i \in [0, L - 1]$ ,  $\mathcal{A}[i]$  stores the size of (possibly empty) bucket  $B_{Z_k}(i)$ . Moreover, as  $|Z_k| \leq m$ , each bucket size must be an integer in  $[0, m]$ .

*The Query Algorithm for a Final-Level Instance.* Given a PSS query with parameters  $(0, W_S(\alpha, \beta))$  on  $V_k$ , the following query algorithm aims to return a PSS result  $T \subseteq Z_k$  with parameters  $(0, W_S(\alpha, \beta))$ .

**QueryFinalLevel**( $\text{BG-Str}(Z_k)$ ,  $(0, W_S(\alpha, \beta))$ ).

- compute the *largest* possible index  $i_1$  of bucket  $B_{Z_k}(i_1)$  such that  $\frac{2^{i_1+1}}{W_S(\alpha, \beta)} \leq \frac{2}{m^2}$ ;
- compute the *smallest* possible index  $i_2$  of bucket  $B_{Z_k}(i_2)$  such that  $\frac{2^{i_2}}{W_S(\alpha, \beta)} \geq 1$ ;
- $T_1 \leftarrow \text{QueryInsignificant}(\text{BG-Str}(Z_k), (0, W_S(\alpha, \beta)), i_1)$  // Algorithm 2;
- $T_2 \leftarrow \text{QueryCertain}(\text{BG-Str}(Z_k), (0, W_S(\alpha, \beta)), i_2)$  // Algorithm 3;
- process all the remaining buckets  $B_{Z_k}(i)$  with index  $i \in [i_1 + 1, i_2 - 1]$ :
  - collect all the bucket sizes in  $\mathcal{A}[i_1+1, i_2-1]$  as an input configuration  $\vec{c}$  for the 4S problem; specifically, recall that the  $j^{\text{th}}$  item in the 4S problem is selected with probability  $p_j = \min\left\{1, \frac{2^{j+1} \cdot \vec{c}[j]}{m^2}\right\}$ , where  $\vec{c}[j] = |B_{Z_k}(j + i_1)|$ ;
  - use the lookup table  $\mathcal{T}$  to obtain a subset sampling result  $r_3$  for the 4S problem with input configuration  $\vec{c}$ ; note that  $r_3$  is a bit string of length  $(i_2 - i_1 - 1)$ , where the  $j^{\text{th}}$  bit in  $r_3$  being 1 indicates the bucket with index  $(j + i_1)$  is selected;
  - initialize an empty set  $T'_3$ ;
  - for each index  $j \in \{1, 2, \dots, i_2 - i_1 - 1\}$  such that  $r_3[j] = 1$ , add the item  $v \in V_k$  corresponding bucket  $B_{Z_k}(j + i_1)$  to  $T'_3$  with probability  $\frac{\min\left\{1, \frac{w(v)}{W_S(\alpha, \beta)}\right\}}{p_j}$ ;
  - $T_3 \leftarrow \text{ExtractItems}(\text{BG-Str}(Z_k), T'_3, (0, W_S(\alpha, \beta)))$  // Algorithm 5
- return  $T = T_1 \cup T_2 \cup T_3$ ;

LEMMA 4.14. *The configuration  $\vec{c}$  is a legal 4S problem input defined for the lookup table  $\mathcal{T}$ .*

PROOF. From the above query algorithm,  $\vec{c}$  contains exactly  $i_2 - 1 - (i_1 + 1) + 1 = i_2 - i_1 - 1$  integers. Observe that all the buckets with indices in the range  $[i_1 + 1, i_2 - 1]$  must correspond to a power-of-two probability in the range  $[\frac{1}{m^2}, 1]$ . Therefore, there are at most  $\lceil 2 \log_2 m \rceil \in O(\log m) = O(\log \log \log n_0)$  possible buckets. And hence,  $\vec{c}$  contains exactly  $K \in O(\log \log \log n_0)$  integers. Moreover, each integer is indeed a bucket size which is in the range  $[0, m]$ . Therefore,  $\vec{c}$  is a legal input configuration to the 4S problem.  $\square$

LEMMA 4.15 (CORRECTNESS).  *$T$  is a correct PSS result on  $Z_k$  with parameters  $(0, W_S(\alpha, \beta))$ .*

PROOF. The correctness of the sampling result  $T_1 \cup T_2$  follows immediately from the correctness of our algorithms for the insignificant instance and the certain instance. By the correctness of the lookup table and the rejection sampling approach, each item  $v \in V_k$  corresponding to a non-empty bucket  $B_{Z_k}(i)$  with index in  $\mathcal{A}[i_1 + 1, i_2 - 1]$  is sampled with probability of  $\min\left\{1, \frac{w(v)}{W_S(\alpha, \beta)}\right\}$

independently. Thus, by the correctness of Algorithm 5 which extracts the items from a given candidate bucket list, the sampling result  $T_3$  is correct, and hence,  $T$  is correct.  $\square$

**LEMMA 4.16 (QUERY COMPLEXITY).** *The overall running time of the above final-level query algorithm is bounded by  $O(1 + \mu_{Z_k})$  in expectation, where  $\mu_{Z_k}$  is the expected size of PSS query result on  $Z_k$  with parameters  $(0, W_S(\alpha, \beta))$ .*

**PROOF.** First, the expected query running time for obtaining items in  $T_1 \cup T_2$  follows from the complexity of the corresponding algorithms.

Next, we analyse the running time cost of obtaining  $T_3$ . By the definition of  $i_1$  and  $i_2$ , the length of the subarray  $\mathcal{A}[i_1 + 1, i_2 - 1]$  is bounded  $i_2 - i_1 - 1 \in O(\log m)$ . Each bucket size is an integer at most  $m$ ; thus, each bucket size can be encoded in  $O(\log m)$  bits. Therefore,  $\mathcal{A}[i_1 + 1, i_2 - 1]$  takes  $O(\log^2 m) = O((\log \log \log n_0)^2) \subset O(d)$  bits, equivalently,  $O(1)$  words. As a result, the input configuration  $\vec{c}$  can be constructed in  $O(1)$  time. Moreover, the lookup table returns a sampling result  $r_3$  for the 4S problem on  $\vec{c}$  in  $O(1)$  time. In the Word RAM model, the next non-zero bit  $j$  in  $r_3$  can be computed in  $O(1)$  time. Thus, the overall expected running time of rejection sampling is bounded by  $O(J)$ , where  $J$  is the number of non-zero bits in  $r_3$ . To complete the proof of this lemma, it remains to show that each potential item  $v$  suggested by a non-zero bit in  $r_3$  has at least  $\Omega(1)$  probability to be accepted in  $T_3$ . If this is the case, then  $O(J)$  as well as the overall expected running time of our query algorithm, is bounded by  $O(1 + \mu_{V_k})$ .

Consider a non-zero bit  $r_3[j] = 1$  at index  $j$ ; let  $v \in V_k$  be the potential item corresponding to the bucket  $B_{Z_k}(j + i_1)$  which is sampled by the lookup table with probability  $p_j = \min\{1, \frac{2^{j+i_1} \cdot |B_{Z_k}(j+i_1)|}{m^2}\}$ . Denote  $p'_j = \min\{1, \frac{w(v)}{W_S(\alpha, \beta)}\}$ , where  $w(v) = 2^{j+i_1+1} \cdot |B_{Z_k}(j+i_1)|$ . Then  $v$  is accepted with probability  $\frac{p'_j}{p_j} \geq \frac{2^{i_1} \cdot m^2}{W_S(\alpha, \beta)} \geq \frac{1}{2}$ , because, by the definition of  $i_1$ , we have  $\frac{2^{i_1+1} m^2}{W_S(\alpha, \beta)} \geq 1$ . Moreover, according to the previous running time analysis on Algorithm 5, it takes  $O(1 + \mu_{Z_k})$  expected time to extract the items from the bucket candidate list  $T'_3$ . Therefore, Lemma 4.16 follows.  $\square$

The above lemmas show that the simple implementation with an array  $\mathcal{A}$  for the adapter is correct and has the desired query complexity. However, its space consumption is too large. This is because the length,  $L$ , of  $\mathcal{A}$  is the number of all possible bucket indices on item set,  $Z_k$ . The weight of item  $z \in Z_k$  can range from  $[1, n_{\max} \cdot w_{\max}]$ , and hence,  $L = \log_2(n_{\max} \cdot w_{\max}) = d$ . Moreover, the bucket size stored in each cell of  $\mathcal{A}$  is an integer in  $[0, m]$ , which in turn takes  $\lceil \log_2(m+1) \rceil \in O(\log m)$  bits to encode. Therefore, the total space consumption of  $\mathcal{A}$  becomes  $O(d \log m)$  bits, equivalently,  $O(\log m) \subseteq O(\log \log \log n_0)$  words. As aforementioned, there can be up to  $O(n_0)$  final-level instances, each which needs an adapter. Thus, the overall space consumption of all the adapters would be  $O(n_0 \log \log \log n_0)$  words, exceeding our promised space bound  $O(n_0)$ .

**A Compact Representation of  $\mathcal{A}$ .** While the total number,  $L$ , of possible bucket indices can be as large as  $d$ , our crucial observation is that, indeed, only a “small” consecutive index sub-range  $[l_1, l_2] \subseteq [0, L - 1]$  can possibly have non-empty buckets of the items in  $Z_k$ . As a result, it suffices to store the sizes of the buckets with indices in  $[l_1, l_2]$ , that is, just to keep a *sub-array*  $\mathcal{A}[l_1, l_2]$  of the conceptual array  $\mathcal{A}$ , along with the value of  $l_1$  (to record the starting index of the sub-range). Hence, the space consumption can be reduced to  $O((l_2 - l_1 + 1) \cdot \log m + d)$  bits, where the  $d$  bits are to store the value of  $l_1$ . This sub-array  $\mathcal{A}[l_1, l_2]$  along with the value of  $l_1$  together is called a *compact representation* of  $\mathcal{A}$ .

When answering a PSS query with the previous query algorithm with a compact representation of  $\mathcal{A}$ , one can simply treat every bucket met during the query with index outside the range  $[l_1, l_2]$  as empty, i.e., of size 0. The correctness and query complexity of the query algorithm still holds. Next, we show the length of  $\mathcal{A}[l_1, l_2]$  is bounded by  $O(\log \log n_0)$ .

LEMMA 4.17. *For every final-level PSS instance on  $Z_k$ , the space consumption of the compact representation of the adapter  $\mathcal{A}$  is bounded by  $O(1)$  words.*

PROOF. Recall that each item  $v_i \in V_k$  corresponds to a bucket  $B_{Z_k}(i)$  of items in  $Z_k$  at bucket index  $i$ . Our goal is to show that those non-empty buckets  $B_{Z_k}(i)$  must have bucket index  $i$  in range  $[l_1, l_2]$ , for some  $l_1$  and  $l_2$ .

To prove this, we show that the *minimum possible* weight and the *maximum possible* weight of the items in  $Z_k$  would not differ by too much. To see this, observe that the item set  $Z_k$  is constructed from the buckets in the group  $G_{Y_j}(k)$ , where  $G_{Y_j}(k)$  contains buckets  $B_{Y_j}(i)$  with bucket indices from  $k_1 = k \cdot \log_2 |Y_j|$  to  $k_2 = (k+1) \cdot \log_2 |Y_j|$ . Therefore, the minimum possible weight of a bucket in  $G_{Y_j}(k)$  is  $2^{k_1+1}$ , while the maximum possible weight of a bucket in  $G_{Y_j}(k)$  is  $2^{k_2+1} \cdot |Y_j|$ . And each of the buckets corresponds to an item  $z_i \in Z_k$  with the same weight. As a result, the *possible weight ratio* of the items in  $Z_k$ , i.e. the ratio between the maximum and minimum possible weights, can be computed as

$$\frac{2^{k_2+1} \cdot |Y_j|}{2^{k_1+1}} = 2^{k_2-k_1} \cdot |Y_j| = 2^{\log_2 |Y_j|} \cdot |Y_j| = |Y_j|^2 \leq (\log_2 n_0)^2,$$

where  $|Y_j| \leq \log_2 n_0$ . Thus, there can be at most  $\log_2 ((\log_2 n_0)^2) = 2 \log_2 \log_2 n_0 \in O(\log \log n_0)$  possible buckets on the items in  $Z_k$ . Moreover, since the smallest possible bucket weight of  $B_{Z_k}(i)$  is  $2^{k_1+1}$ , the smallest possible bucket index for  $B_{Z_k}(i)$  is  $l_1 = k_1 + 1 = k \cdot \log_2 |Y_j| + 1$ . And hence,  $l_2 \leq l_1 + 2 \log_2 \log_2 n_0$ .

By the previous discussion, a compact representation of  $\mathcal{A}$  consists of a sub-array  $\mathcal{A}[l_1, l_2]$ , where each cell stores an integer in range  $[0, m]$ , and an integer value  $l_1$ . The total space consumption is thus bounded by  $(l_2 - l_1 + 1) \cdot \lceil \log_2(m+1) \rceil + d \in O(\log \log n_0 \cdot \log \log \log n_0 + d)$  bits which is bounded by  $O(1)$  words.  $\square$

#### 4.5 Handling Updates

Next, we discuss how to update our HALT structure on the item set  $S$ . Specifically, each update on  $S$  is either an insertion of a new item  $x$  with weight  $w(x)$  or a deletion of an existing item  $x$  from  $S$ . For simplicity, we assume that  $\frac{1}{2}n_0 \leq n \leq 2n_0$ , where  $n$  is the current size of  $S$ .

The HALT structure on  $S$  can be maintained as follows. First, observe that the lookup table in the HALT structure on  $S$  does not need to be updated. Second, the three-level sampling hierarchy is basically a hierarchy of sorted lists of buckets and groups with length at most  $O(d)$ . Each update operation (either insertion or deletion) on these sorted lists can be easily done in  $O(1)$  time in the Word RAM Model. The only thing to be careful about is that each update operation at the previous level in the hierarchy could trigger a *weight update* of a next-level item, which would require two update operations to achieve: delete the item first and then insert the item again with the new weight. As the hierarchy has only three levels, therefore, there can be in total at most  $O(1)$  update operations to the hierarchy. Moreover, there is also only  $O(1)$  final-level instances would be affected by an update to  $S$ . Their adapters can be updated simply by updating the corresponding bucket sizes in the sub-array  $\mathcal{A}[l_1, l_2]$ . Therefore, the overall worst-case update time is  $O(1)$ .

To handle the case when  $n < \frac{1}{2}n_0$  or  $n > 2n_0$ , we adopt the standard global rebuilding technique to re-construct the entire HALT structure on the current item set  $S$  and update  $n_0 \leftarrow n$ . Recall that the HALT structure on  $S$  can be constructed in  $O(n_0)$  time with  $O(n_0)$  space consumption. As there must be at least  $\Omega(n_0)$  updates to trigger the reconstruction, the rebuilding cost can be charged to those updates making each update cost become  $O(1)$  amortized. Moreover, the space consumption is bounded by  $O(n)$  all the times.

This  $O(1)$  amortized update time can be easily “de-amortized” by applying the same technique for the de-amortization for dynamic arrays, just increasing the space consumption by a constant factor. Thus, we have the following theorem.

**THEOREM 4.18.** *The HALT structure on item set  $S$  can be maintained in  $O(1)$  worst-case time for each update to  $S$ , and consumes  $O(n)$  space at all times, where  $n$  is the current size of  $S$ .*

## 5 A Hardness Result on the DPSS Problem with Float Item Weights

We now prove Theorem 1.2. Let *DPSS-ALG* be a deletion-only DPSS algorithm for float item weights with pre-processing time  $\mathcal{P}(N)$ , query time  $\mathcal{Q}(N)$ , and deletion time  $\mathcal{D}(N)$ . As a black box, we apply the algorithm to execute Integer Sorting.

Consider a set of  $N$  integers  $I = \{a_1, \dots, a_N\}$ , each of which is represented by one word of  $d$  bits. We can assume w.l.o.g. that the integers are distinct<sup>2</sup>. The integers in  $I$  can be sorted in descending order by the following algorithm:

- for each integer  $a_i \in I$ , create an item  $x_i$  with weight  $w(x_i) = 2^{a_i}$ , represented by a float number;
- initialize  $S$  to be the set of all these  $N$  items;
- initialize an empty linked list,  $R$ , of the integers in  $I$ , which is maintained to be sorted, in descending order, by the Insertion Sort [14] algorithm;
- initialize *DPSS-ALG* on  $S$ ;
- while  $S$  is not empty, perform the following:
  - repeatedly invoke *DPSS-ALG* on  $S$  to perform a PSS query with parameters  $(1, 0)$  until the sampling result  $T \neq \emptyset$ ;
  - let  $x^*$  be the item in  $T$  with the largest weight and let  $w(x^*) = 2^{a^*}$ ;
  - invoke *DPSS-ALG* to delete  $x^*$  from  $S$ ;
  - invoke Insertion Sort to insert the weight exponent,  $a^*$ , of  $x^*$  to  $R$  (in descending order);
- return  $R$  as the sorted list of all the integers in  $I$  (in descending order);

**Correctness.** When the above algorithm terminates,  $R$  contains all the  $N$  integers in  $I$ . As  $R$  is maintained by the Insertion Sort algorithm,  $R$  is a correct sorted list of  $I$ .

**Running Time Analysis.** Next we analyse the overall expected running time of the above integer sorting algorithm. First, the construction cost of the item set  $S$  is bounded by  $O(N)$  and the pre-processing cost of *DPSS-ALG* is  $t_p(N)$ . Next, we focus on the running time in each iteration in the while-loop. Denote the item set at the start of the  $i^{\text{th}}$  iteration by  $S_i$ , where  $i \in \{1, \dots, N\}$  and  $S_1 = S$ .

**LEMMA 5.1.** *In the  $i^{\text{th}}$  iteration, in expectation, at most two PSS queries with parameters  $(1, 0)$  on  $S_i$  are invoked to obtain a non-empty subset sampling result,  $T$ .*

**PROOF.** To prove this lemma, it suffices to show that the largest item (i.e., with the largest item weight) in  $S_i$  will be selected in the sampling result  $T$  with probability at least  $\frac{1}{2}$ . As a result, with at most 2 trials, in expectation, the largest item would be sampled in  $T$ , and hence,  $T$  is non-empty.

To see this, first, observe that with the query parameters  $(1, 0)$ , we have  $W_{S_i}(1, 0) = \sum_{x \in S_i} w(x)$ . Second, by the construction of  $S$ , the weights of all the items are power-of-two numbers with distinct exponents. Therefore, it can be verified that the largest item in  $S_i$  must have weight  $\geq \frac{1}{2}W_{S_i}(1, 0)$ , and hence, the corresponding item has probability at least  $\frac{1}{2}$  to be sampled.  $\square$

**LEMMA 5.2.** *For a PSS query with parameters  $(1, 0)$  on  $S_i$ , the expected result size  $\mu_{S_i}(1, 0) = 1$ .*

<sup>2</sup>To achieve distinctness, to each integer, we can append a unique ID with a word in its binary representation.

PROOF. It follows from the fact that  $W_{S_i}(1, 0) = \sum_{x \in S_i} w(x)$  and  $\mu_{S_i}(1, 0) = \sum_{x \in S_i} \frac{w(x)}{W_{S_i}(1, 0)}$ .  $\square$

By Lemmas 5.1 and 5.2, in the  $i^{\text{th}}$  iteration, the running time for the steps of sampling a non-empty subset  $T$  and finding the largest item  $x^*$  in  $T$  is bounded by  $O(t_q(N))$  in expectation. Moreover, DPSS-ALG deletes  $x^*$  from  $S_i$  in  $t_{del}(N)$  worst-case time. As a result, summing over all  $N$  iterations, the overall expected cost for these steps is bounded by  $O(N \cdot (t_q(N) + t_{del}(N)))$ .

Next, we analyse the running time for the Insertion Sort.

LEMMA 5.3. *The expected running time for the Insertion Sort is  $O(N)$ .*

PROOF. The order of the integers in  $I$  is identical to the order of their corresponding weights. To ease the discussion, we treat the sorted linked list  $R$  of the integers in  $I$  as the sorted linked list of their corresponding items. It is known that the overall cost for Insertion Sort is bounded by  $O(N + \sum_{i=1}^N \#\text{Swap}_i)$ , where  $\#\text{Swap}_i$  is the number of *swaps* for the item  $x^*$  which is inserted to the sorted linked list  $R$  from the back in the  $i^{\text{th}}$  iteration. Equivalently,  $\#\text{Swap}_i$  is the number of items that have weights *smaller than*  $w(x^*)$  and were inserted to  $R$  *before* the  $i^{\text{th}}$  iteration.

Define  $\#\text{FutureSwap}_i$  as the number of items that have weights *greater than*  $w(x^*)$  and will be inserted to  $R$  *after* the  $i^{\text{th}}$  iteration. Observe that  $\sum_{i=1}^N \#\text{Swap}_i = \sum_{i=1}^N \#\text{FutureSwap}_i$ . Furthermore, define the *rank* of an item  $x \in S_i$ , denoted by  $\text{rank}_{S_i}(x)$ , as the number of items that have weight greater than  $w(x)$  in  $S_i$ . Then, we have  $\#\text{FutureSwap}_i = \text{rank}_{S_i}(x^*)$ , where  $x^*$  is the item inserted to  $R$  in the  $i^{\text{th}}$  iteration. This follows from the definition of  $\#\text{FutureSwap}_i$  and the fact that  $S_i \setminus \{x^*\}$  is exactly the set of items inserted to  $R$  in future iterations. As a result, it suffices to bound the expectation of  $\text{rank}_{S_i}(x^*)$ .

CLAIM 3. [\*] *In the  $i^{\text{th}}$  iteration, consider the largest item  $x^*$  in the sample set  $T \subseteq S_i$ . We have that  $\mathbb{E}[\text{rank}_{S_i}(x^*)] \in O(1)$ .*

Therefore, the running time of the Insertion Sort over all the  $N$  iterations is bounded by  $O(N)$  in expectation. This completes the proof of Lemma 5.3.  $\square$

Putting together all of the above analyses, the total expected running time of the algorithm is  $O(N) + t_p(N) + O(N \cdot (t_q(N) + t_{del}(N))) + O(N)$ , which is  $t_p(N) + O(N \cdot (t_q(N) + t_{del}(N)))$ .

Theorem 1.2 thus follows.

## 6 Conclusion

In this paper, we study the Dynamic Parameterized Subset Sampling (DPSS) problem in the Word RAM model. Our first main result is an optimal DPSS algorithm. We propose our novel data structure HALT that achieves the optimal bounds. Our second main result showing that if the item weights in the DPSS problem are allowed to be float numbers, then an deletion-only optimal DPSS algorithm can be used to solve Integer Sorting over the whole range of bit-length in  $O(N)$  expected time, where  $N$  is the number of integers to sort. Finally, our third main result is a new efficient method for sampling variates from the Truncated Geometric distribution in  $O(1)$  expected time in the Word RAM model. We expect considerable application of this technique for efficient random variate realization.

## Acknowledgements

In this work, Junhao Gan is supported in part by Australian Research Council (ARC) Discovery Early Career Research Award (DECRA) DE190101118. Hanzhi Wang is supported by the VILLUM Foundation grant 54451.

## A Pseudocode of Algorithms 2 to 5

---

**Algorithm 2:** QueryInsignificant(BG-Str( $X$ ),  $(\alpha, \beta)$ ,  $i_1$ )
 

---

**Data:** a Bucket-Group Structure BG-Str( $X$ ), PSS query parameters  $(\alpha, \beta)$ , the maximum possible bucket index  $i_1$  in the  $(\alpha, \beta)$ -insignificant groups

**Result:** a PSS result  $T_1 \subseteq X$  within all the  $(\alpha, \beta)$ -insignificant groups

```

1 initialize  $T_1 \leftarrow \emptyset$  and  $N \leftarrow |X|$ ;
2  $k \leftarrow \text{B-Geo}(1/N^2, N + 1)$ ;
3 if  $k \leq N$  then
4    $A \leftarrow$  the set of all the items in the non-empty buckets (with bucket index  $\leq i_1$ );
5   if  $|A| < k$  then
6      $\perp$  return  $\emptyset$ ;
7    $x \leftarrow A[k]$ , the  $k^{\text{th}}$  item in  $A$ ;
8    $T_1 \leftarrow T_1 \cup \{x\}$  with probability  $\frac{p_x(\alpha, \beta)}{1/N^2}$ ;
9   for  $i \leftarrow \{k + 1, \dots, |A|\}$  do
10     $x \leftarrow A[i]$ , the  $i^{\text{th}}$  item in  $A$ ;
11     $\perp$   $T_1 \leftarrow T_1 \cup \{x\}$  with probability  $p_x(\alpha, \beta)$ ;
12 return  $T_1$ ;

```

---



---

**Algorithm 3:** QueryCertain(BG-Str( $X$ ),  $(\alpha, \beta)$ ,  $i_2$ )
 

---

**Data:** a Bucket-Group Structure BG-Str( $X$ ), PSS query parameters  $(\alpha, \beta)$ , the minimum possible bucket index  $i_2$  in the  $(\alpha, \beta)$ -certain groups

**Result:** a PSS result  $T_2 \subseteq X$  within all the  $(\alpha, \beta)$ -certain groups

```

1 initialize  $T_2 \leftarrow \emptyset$ ;
2 for each non-empty buckets  $B_X(i)$  with bucket index  $i \geq i_2$  in BG-Str( $X$ ) do
3    $\perp$   $T_2 \leftarrow T_2 \cup B_X(i)$ ;
4 return  $T_2$ ;

```

---

## B Proofs for Section 3 (Sampling Random Variates in the Word RAM Model)

LEMMA 3.3. *These two claims hold simultaneously: (i) the numerator and the denominator of  $p = p^*$  both can be represented with  $O(n)$  words and  $p$  can be evaluated in  $O(\text{poly}(n))$  time; and (ii) there exists an algorithm that can evaluate each  $i$ -bit approximation  $\tilde{p}_i$  of  $p = p^*$ , satisfying all conditions stated in Theorem 3.1, in  $O(\text{poly}(i))$  time.*

PROOF. We first prove claim (i). Let  $A$  and  $B$  denote the numerator and denominator of  $q$ . By definition, both  $A$  and  $B$  fit in  $O(1)$  words. Then it can be verified that the numerator and the denominator of  $p^*$  can be computed as  $B^n - (B - A)^n$  and  $B^n - nAB^{n-1}$  in  $O(\text{poly}(n))$  time, respectively. And each of them fits in  $O(n)$  words as they are just multiplication results of  $n$  constant-word integers. As a result,  $p^* = \frac{1 - (1-q)^n}{n \cdot q}$  is a rational number with  $O(n)$ -word numerator and denominator and it can be evaluated in  $O(\text{poly}(n))$  time.

Next we prove claim (ii). According to the binomial theorem,  $1 - (1 - q)^n = \sum_{j=1}^n \binom{n}{j} (-1)^{j+1} q^j$ . Thus,  $p^* = \frac{1 - (1-q)^n}{n \cdot q}$  can be written as  $p^* = \sum_{j=1}^n a_j$ , where  $a_j = \frac{(-1)^{j+1} q^{j-1} (n-1)!}{j!(n-j)!}$ . Note that

**Algorithm 4:** QueryASignificantGroup(BG-Str( $X$ ),  $(\alpha, \beta)$ ,  $j$ ,  $\ell_{\text{cur}}$ )

---

**Data:** a Bucket-Group Structure BG-Str( $X$ ), PSS query parameters  $(\alpha, \beta)$ , an  $(\alpha, \beta)$ -significant group index  $j$ , the current recursion level  $\ell_{\text{cur}}$

**Result:** a PSS result  $T_3 \subseteq X$  within the  $(\alpha, \beta)$ -significant group  $G_X(j)$

- 1 **if**  $G_X(j)$  is empty **then**
- 2    $\lfloor$  return  $\emptyset$ ;
- 3 // obtain a PSS result  $T_{Y_j}$  with parameters  $(0, W_X(\alpha, \beta))$  on the next-level item set  $Y_j$  of  $G_X(j)$ ;
- 4 **if**  $\ell_{\text{cur}} = 1$  **then**
- 5    $\lfloor$   $T_{Y_j} \leftarrow \text{Query}(\text{BG-Str}(Y_j), (0, W_X(\alpha, \beta)), \ell_{\text{cur}} + 1)$  // recursively invoke Algorithm 1;
- 6 **else**
- 7    $\lfloor$   $T_{Y_j} \leftarrow \text{QueryFinalLevel}(\text{BG-Str}(Y_j), (0, W_X(\alpha, \beta)))$  // in Section 4.4;
- 8  $T_3 \leftarrow \text{ExtractItems}(\text{BG-Str}(X), T_{Y_j}, (0, W_X(\alpha, \beta)))$  // Algorithm 5;
- 9 **return**  $T_3$ ;

---

**Algorithm 5:** ExtractItems(BG-Str( $X$ ),  $T_{Y_j}$ ,  $(0, W_X(\alpha, \beta))$ )

---

**Data:** a Bucket-Group Structure BG-Str( $X$ ), a PSS result  $T_j$  on the next-level items, each corresponds to a candidate bucket of items of  $X$ , PSS query parameters  $(0, W_X(\alpha, \beta))$

**Result:** a PSS result  $T_3 \subseteq X$  from the items in the candidate bucket list  $T_{Y_j}$

- 1 initialize  $T_3 \leftarrow \emptyset$ ;
- 2 **for each** bucket  $B_X(i) \in T_{Y_j}$  **do**
- 3    $p \leftarrow \min\{1, \frac{2^{i+1}}{W_X(\alpha, \beta)}\}$  and  $n_i \leftarrow |B_X(i)|$ ;
- 4   **if**  $p \cdot n_i \geq 1$  **then**
- 5      $\lfloor$   $k \leftarrow \text{B-Geo}(p, n_i + 1)$ ;
- 6     **else**
- 7       **if**  $\text{Ber}(\frac{1-(1-p)^{n_i}}{p \cdot n_i}) = 0$  **then**
- 8          $\lfloor$  continue to process next bucket in  $T_{Y_j}$ ;
- 9          $\lfloor$   $k \leftarrow \text{T-Geo}(p, n_i)$ ;
- 10    **while**  $k \leq n_i$  **do**
- 11      $x \leftarrow$  the  $k^{\text{th}}$  item in  $B_X(i)$ ;
- 12      $T_3 \leftarrow T_3 \cup \{x\}$  with probability  $\frac{p_x(\alpha, \beta)}{p}$ ;
- 13      $k \leftarrow k + \text{B-Geo}(p, n_i + 1)$ ;
- 14 **return**  $T_3$ ;

---

$\frac{(n-1)!}{(n-j)!} \leq n^{j-1}$  and  $nq \leq 1$ . We can bound the absolute value of  $a_j$  with  $|a_j| \leq \frac{1}{j!} \leq \frac{1}{2^{j-1}}$ . Further,  $|\sum_{j=i+3}^n a_j| \leq \sum_{j=i+3}^n |a_j| \leq \frac{1}{2^{i+1}}$ . In other words, the sum of  $a_j$  for  $1 \leq j \leq i+2$  achieves a good approximation of  $p^*$ . Denoting  $\bar{p}_i = \sum_{j=1}^{i+2} a_j$ , we have  $|p^* - \bar{p}_i| \leq \frac{1}{2^{i+1}}$ .

Now we discuss how to compute  $\bar{p}_i$  in the Word RAM model. We perform the computation with *working precision*  $r$ , which means we work with floating point numbers and each of these floating point numbers contains an exact exponent encoded by an array of words and a mantissa which is an array of  $\lceil r/d \rceil$  words. Each mathematical operation between two floating point numbers under

working precision  $r$  can be done in  $O(\text{poly}(r))$  time. Since  $\bar{p}_i$  can be evaluated with  $O(\text{poly}(i))$  such floating point mathematical operations, thus an approximation (with working precision  $r$ ) of  $\bar{p}_i$ , denoted by  $\tilde{p}_i(r)$ , can be computed in  $O(\text{poly}(i, r))$  time. Analogous to the analysis in Appendix C of [7], it suffices to choose  $r = O(\text{poly}(i))$  to achieve an absolute error  $|\tilde{p}_i(r) - \bar{p}_i| \leq \frac{1}{2^{i+1}}$ .

As a result,  $|\tilde{p}_i(r) - p^*| \leq |\tilde{p}_i(r) - \bar{p}_i| + |p^* - \bar{p}_i| \leq \frac{1}{2^{i+1}} + \frac{1}{2^{i+1}} = \frac{1}{2^i}$ . Therefore,  $\tilde{p}_i = \tilde{p}_i(r)$  is a valid  $i$ -bit approximation of  $p^*$ , and the overall computational cost is bounded by  $O(\text{poly}(i))$ .  $\square$

**LEMMA 3.4.** *These two claims hold simultaneously: (i) the numerator and the denominator of  $p = \frac{1}{2p^*}$  both can be represented with  $O(n)$  words and  $p$  can be evaluated in  $O(\text{poly}(n))$  time; and (ii) there exists an algorithm that can evaluate each  $i$ -bit approximation  $\tilde{p}_i$  of  $p = \frac{1}{2p^*}$ , satisfying all conditions stated in Theorem 3.1, in  $O(\text{poly}(i))$  time.*

**PROOF.** Analogous to the proof of Lemma 3.3, the numerator and the denominator of  $p = \frac{1}{2p^*}$  fit in  $O(n)$  words and  $p$  can be evaluated in  $\text{poly}(n)$  time. Thus, claim (i) holds.

To prove claim (ii), by Lemma 3.3, for an arbitrary integer  $r \geq 3$ , an  $r$ -bit approximation  $\tilde{p}_r^*$  of  $p^*$  can be computed in  $O(\text{poly}(r))$  time, such that  $|\tilde{p}_r^* - p^*| \leq \frac{1}{2^r}$ . By Newton's method [6], an approximate reciprocal of  $2\tilde{p}_r^*$ , denoted by  $u(r)$ , can be computed in  $O(\text{poly}(i))$  time with error at most  $\frac{1}{2^{i+1}}$ , namely,  $|u(r) - \frac{1}{2\tilde{p}_r^*}| \leq \frac{1}{2^{i+1}}$ .

Therefore, the absolute error,  $|p - u(r)|$ , is bounded by  $|\frac{1}{2p^*} - u(r)| \leq |\frac{1}{2p^*} - \frac{1}{2\tilde{p}_r^*}| + |u(r) - \frac{1}{2\tilde{p}_r^*}| \leq \frac{|2\tilde{p}_r^* - 2p^*|}{2p^* \cdot 2\tilde{p}_r^*} + \frac{1}{2^{i+1}} \leq \frac{1}{2^{r-1}} + \frac{1}{2^{i+1}}$ , where the last inequality comes from the upper bound of  $p^* \geq 1 - 1/e \geq \frac{1}{2}$  and  $\tilde{p}_r^* \geq p^* - \frac{1}{2^r} \geq \frac{1}{2}$ .

Thus, it suffices to choose  $r = i + 2$  to obtain an  $i$ -bit approximation  $u(r)$  of  $p$ , and the overall computational cost is bounded by  $O(\text{poly}(i))$ .  $\square$

## C Proofs for Section 5 (A Hardness Result on the DPSS Problem with Float Item Weights)

**CLAIM 3.** *In the  $i^{\text{th}}$  iteration, consider the largest item  $x^*$  in the sample set  $T \subseteq S_i$ . We have that  $\mathbb{E}[\text{rank}_{S_i}(x^*)] \in O(1)$ .*

**PROOF.** As discussed earlier, the item weights are power-of-two numbers with distinct exponents. We have the following observations. First, the largest item  $y_1 \in S_i$  would be sampled in the PSS query on  $S_i$  with parameter  $(1, 0)$  with probability at least  $\frac{1}{2}$  and at most 1. Thus, the probability of the  $j^{\text{th}}$  largest item  $y_j$  in  $S_i$  to be sampled is at most  $\frac{1}{2^{j-1}}$ , for all  $j \in \{1, \dots, |S_i|\}$ . Second, the  $j^{\text{th}}$  largest item  $y_j$  in  $S_i$  has  $\text{rank}_{S_i}(y_j) = j - 1$ .

Moreover, the probability that the  $j^{\text{th}}$  largest item  $y_j$  in  $S_i$  is the largest item in the sample set  $T$  is at most the probability of  $y_j$  being sampled conditioned on  $T$  being non-empty. Therefore, the expected rank of  $x^*$  satisfies:

$$\begin{aligned} \mathbb{E}[\text{rank}_{S_i}(x^*)] &\leq \sum_{j=1}^{|S_i|} \text{rank}_{S_i}(y_j) \cdot \frac{\Pr[y_j \text{ being sampled} \wedge T \text{ is non-empty}]}{\Pr[T \text{ is non-empty}]} \\ &\leq \sum_{j=1}^{|S_i|} (j-1) \cdot \frac{2}{2^{j-1}} \in O(1), \end{aligned}$$

where the last inequality is by the fact that  $T$  is non-empty with probability at least  $\frac{1}{2}$ .  $\square$

## References

- [1] 2024. Optimal Dynamic Parameterized Subset Sampling. arXiv:2409.18036 [cs.DS]. Available at <https://arxiv.org/abs/2409.18036>.
- [2] Reid Andersen, Fan Chung, and Kevin Lang. 2006. Local graph partitioning using pagerank vectors. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. IEEE, 475–486.
- [3] Adrian Baldwin, Iffat Gheyas, Christos Ioannidis, David Pym, and Julian Williams. 2017. Contagion in cyber security attacks. *Journal of the Operational Research Society* 68, 7 (2017), 780–791.
- [4] Djamel Belazzougui, Gerth Stølting Brodal, and Jesper Sindahl Nielsen. 2014. Expected Linear Time Sorting for Word Size  $\Omega(\log^2 n \log \log n)$ . In *Scandinavian Workshop on Algorithm Theory*. Springer, 26–37.
- [5] Sayan Bhattacharya, Peter Kiss, Aaron Sidford, and David Wajc. 2023. Near-Optimal Dynamic Rounding of Fractional Matchings in Bipartite Graphs. *arXiv preprint arXiv:2306.11828* (2023).
- [6] Richard P Brent and Paul Zimmermann. 2010. *Modern computer arithmetic*. Vol. 18. Cambridge University Press.
- [7] Karl Bringmann and Tobias Friedrich. 2013. Exact and efficient generation of geometric random variates and random graphs. In *International Colloquium on Automata, Languages, and Programming*. Springer, 267–278.
- [8] Mark De Berg. 2000. *Computational geometry: algorithms and applications*. Springer Science & Business Media.
- [9] Hossein Esfandiari, Vahab Mirrokni, and Peilin Zhong. 2021. Almost linear time density level set estimation via dbscan. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 7349–7357.
- [10] Philippe Flajolet and Nasser Saheb. 1986. The complexity of generating an exponentially distributed variate. *Journal of Algorithms* 7, 4 (1986), 463–488.
- [11] James Fleming and Mark Cannon. 2018. Stochastic MPC for additive and multiplicative uncertainty using sample approximations. *IEEE Trans. Automat. Control* 64, 9 (2018), 3883–3888.
- [12] Michael L Fredman and Dan E Willard. 1993. Surpassing the information theoretic bound with fusion trees. *Journal of computer and system sciences* 47, 3 (1993), 424–436.
- [13] Timothy C Germann, Kai Kadau, Ira M Longini Jr, and Catherine A Macken. 2006. Mitigation strategies for pandemic influenza in the United States. *Proceedings of the National Academy of Sciences* 103, 15 (2006), 5935–5940.
- [14] Sandeep Kaur Gill, Virendra Pal Singh, Pankaj Sharma, and Durgesh Kumar. 2019. A comparative study of various sorting algorithms. *International Journal of Advanced Studies of Scientific Research* 4, 1 (2019).
- [15] Qintian Guo, Sibowang, Zhewei Wei, and Ming Chen. 2020. Influence Maximization Revisited: Efficient Reverse Reachable Set Generation with Bound Tightened. In *SIGMOD Conference*. ACM, 2167–2181.
- [16] Qintian Guo, Sibowang, Zhewei Wei, Wenqing Lin, and Jing Tang. 2022. Influence maximization revisited: efficient sampling with bound tightened. *ACM Transactions on Database Systems (TODS)* 47, 3 (2022), 1–45.
- [17] Mayssa Hajar, Mohamad El Badaoui, Amani Raad, and Frederic Bonnardot. 2019. Discrete random sampling: Theory and practice in machine monitoring. *Mechanical Systems and Signal Processing* 123 (2019), 386–402.
- [18] Yijie Han and Mikkel Thorup. 2002. Sorting integers in  $O(n\sqrt{\log \log n})$  expected time and linear space. In *IEEE Symposium on Foundations of Computer Science (FOCS'02)*.
- [19] David Haussler and Emo Welzl. 1986. Epsilon-nets and simplex range queries. In *Proceedings of the second annual symposium on Computational geometry*. 61–71.
- [20] Taher Haveliwala et al. 1999. *Efficient computation of PageRank*. Technical Report. Citeseer.
- [21] Brendan Juba, Christopher Musco, Fan Long, Stelios Sidiropoulos-Douskos, and Martin C Rinard. 2015. Principled Sampling for Anomaly Detection.. In *NDSS*.
- [22] Alexander D Kent. 2016. Cyber security data sources for dynamic network research. In *Dynamic Networks and Cyber-Security*. World Scientific, 37–65.
- [23] Vipin Kumar and Sonajharia Minz. 2014. Feature selection. *SmartCR* 4, 3 (2014), 211–229.
- [24] Edo Liberty, Kevin Lang, and Konstantin Shmakov. 2016. Stratified sampling meets machine learning. In *International conference on machine learning*. PMLR, 2320–2329.
- [25] Jianning Mai, Chen-Nee Chuah, Ashwin Sridharan, Tao Ye, and Hui Zang. 2006. Is sampled data sufficient for anomaly detection?. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. 165–176.
- [26] Florin Rusu and Alin Dobra. 2009. Sketching sampled data streams. In *2009 IEEE 25th International Conference on Data Engineering*. IEEE, 381–392.
- [27] Daniel A Spielman and Shang-Hua Teng. 2004. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. 81–90.
- [28] Loan Thi Vu, Lien Thi Vu, Nga Thu Nguyen, Phuong Thi Thuy Do, and Daniel Dao. 2019. Feature selection methods and sampling techniques to financial distress prediction for Vietnamese listed companies. *Investment Management and Financial Innovations* 16, 1 (2019), 276–290.
- [29] Hanzhi Wang, Mingguo He, Zhewei Wei, Sibowang, Ye Yuan, Xiaoyong Du, and Ji-Rong Wen. 2021. Approximate graph propagation. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 1686–1696.

- [30] Wenpu Xing and Ali Ghorbani. 2004. Weighted pagerank algorithm. In *Proceedings. Second Annual Conference on Communication Networks and Services Research, 2004*. IEEE, 305–314.
- [31] Renchi Yang, Xiaokui Xiao, Zhewei Wei, Sourav S Bhowmick, Jun Zhao, and Rong-Hua Li. 2019. Efficient estimation of heat kernel pagerank for local clustering. In *Proceedings of the 2019 International Conference on Management of Data*. 1339–1356.
- [32] Lu Yi, Hanzhi Wang, and Zhewei Wei. 2023. Optimal Dynamic Subset Sampling: Theory and Applications. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3116–3127.
- [33] Houssam Zenati, Manon Romain, Chuan-Sheng Foo, Bruno Lecouat, and Vijay Chandrasekhar. 2018. Adversarially learned anomaly detection. In *2018 IEEE International conference on data mining (ICDM)*. IEEE, 727–736.
- [34] Lingling Zhang, Hong Jiang, Ye Yuan, and Guoren Wang. 2024. Influence Maximization in Hypergraphs by Stratified Sampling for Efficient Generation of Reverse Reachable Sets. *arXiv preprint arXiv:2406.01911* (2024).

Received May 2024; accepted August 2024