



HAL
open science

Dynamic Local Searchable Symmetric Encryption

Brice Minaud, Michael Reichle

► **To cite this version:**

Brice Minaud, Michael Reichle. Dynamic Local Searchable Symmetric Encryption. Crypto 2022 - 42nd Annual International Cryptology Conference, Aug 2022, Santa Barbara, United States. 10.1007/978-3-031-15985-5_4 . hal-03863896

HAL Id: hal-03863896

<https://hal.science/hal-03863896>

Submitted on 21 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dynamic Local Searchable Symmetric Encryption

Brice Minaud* and Michael Reichle*

*École normale supérieure, PSL University, CNRS, Inria, France.
{brice.minaud,michael.reichle}@ens.fr

Abstract. In this article, we tackle for the first time the problem of *dynamic* memory-efficient Searchable Symmetric Encryption (SSE). In the term “memory-efficient” SSE, we encompass both the goals of *local* SSE, and *page-efficient* SSE. The centerpiece of our approach is a new connection between those two goals. We introduce a map, called the Generic Local Transform, which takes as input a *page-efficient* SSE scheme with certain special features, and outputs an SSE scheme with strong *locality* properties. We obtain several results.

- First, we build a *dynamic* SSE scheme with storage efficiency $\mathcal{O}(1)$ and page efficiency only $\tilde{\mathcal{O}}(\log \log(N/p))$, where p is the page size, called **LayeredSSE**. The main technique behind **LayeredSSE** is a new weighted extension of the two-choice allocation process, of independent interest.
- Second, we introduce the Generic Local Transform, and combine it with **LayeredSSE** to build a *dynamic* SSE scheme with storage efficiency $\mathcal{O}(1)$, locality $\mathcal{O}(1)$, and read efficiency $\tilde{\mathcal{O}}(\log \log N)$, under the condition that the longest list is of size $\mathcal{O}(N^{1-1/\log \log \lambda})$. This matches, in every respect, the purely *static* construction of Asharov et al. from STOC 2016: dynamism comes at no extra cost.
- Finally, by applying the Generic Local Transform to a variant of the Tethys scheme by Bossuat et al. from Crypto 2021, we build an unconditional static SSE with storage efficiency $\mathcal{O}(1)$, locality $\mathcal{O}(1)$, and read efficiency $\mathcal{O}(\log^\varepsilon N)$, for an arbitrarily small constant $\varepsilon > 0$. To our knowledge, this is the construction that comes closest to the lower bound presented by Cash and Tessaro at Eurocrypt 2014.

1 Introduction

Searchable Symmetric Encryption. In Searchable Symmetric Encryption (SSE), a client outsources the storage of a set of documents to an untrusted server. The client wishes to retain the ability to search the documents, by issuing search queries to the server. In the setting of *dynamic* SSE, the client may also issue update queries, in order to modify the contents of the database, for instance by adding or removing entries. The server must be able to correctly process all queries, while learning as little information as possible about the client’s data and queries. SSE is relevant in many cloud storage scenarios: for example, in cases such as outsourcing the storage of a sensitive database, or offering an encrypted messaging service, some form of search functionality may be highly desirable.

In theory, SSE is a special case of computation on encrypted data, and could be realized using generic solutions, such as Fully Homomorphic Encryption. In practice, such approaches incur a large performance penalty. Instead, SSE schemes typically aim for high-performance solutions, scalable to large real-world databases. Towards that end, SSE trades off security for efficiency. The server is allowed to learn some information about the client’s data. For example, SSE schemes typically leak to the server the repetition of queries (*search pattern*), and the identifiers of the documents that match a query (*access pattern*). The security model of SSE is parametrized by a *leakage function*, which specifies the nature of the information leaked to the server.

Locality. In the case of single-keyword SSE, search queries ask for all documents that contain a given keyword. To realize that functionality, the server maintains an (encrypted) reverse index, where each keyword is mapped to the list of identifiers of documents that match the keyword. When the client wishes to search for the documents that match a given keyword, the client simply retrieves the corresponding list from the server. A subtle issue, however, is how the lists should be stored and accessed by the server.

The naive approach of storing one list after the other is unsatisfactory: indeed, the position of a given list in memory becomes dependent on the lengths of other lists, thereby leaking information about those lists. A common approach to address that issue is to store each list element at a random location in memory. In that case, when retrieving a list, the server must visit as many random memory locations as the number of elements in the list. This is also undesirable, for a different reason: for virtually all modern storage media, accessing many random memory locations is much more expensive than visiting one continuous region. Because SSE relies on fast symmetric cryptographic primitives, the cost of memory accesses becomes the performance bottleneck. To capture that cost, [CT14] introduces the notion of *locality*: in short, the locality of an SSE scheme is the number of discontinuous memory locations that the server must access to answer a query.

The two extreme solutions outlined above suggest a conflict between security and locality. At Eurocrypt 2014, Cash and Tessaro showed that this conflict is inherent [CT14]: if a secure SSE scheme has constant storage efficiency (the size of the encrypted database is linear in the size of the plaintext database), and constant read efficiency (the amount of data read by the server to answer a search query is linear in the size of the plaintext answer), then it cannot have constant locality.

Local SSE constructions. Since then, many SSE schemes with constant locality have been proposed, typically at the cost of superconstant read efficiency. At STOC 2016, Asharov et al. presented a scheme with $\mathcal{O}(1)$ storage efficiency, $\mathcal{O}(1)$ locality, and $\tilde{\mathcal{O}}(\log N)$ read efficiency, where N is the size of the database [ANSS16]. At Crypto 2018, Demertzis et al. improved the read efficiency to $\mathcal{O}\left(\log^{2/3+\varepsilon} N\right)$ [DPP18]. Several trade-offs with $\omega(1)$ storage efficiency were also proposed in [DP17]. When the size of the longest list in the database is bounded, stronger results are known. When such an upper bound is required, we

will call the construction *conditional*. The first conditional SSE is due to Asharov et al., and achieves $\tilde{\mathcal{O}}(\log \log N)$ read efficiency, on the condition that the size of the longest list is $\mathcal{O}(N^{1-1/\log \log N})$. This was later improved to $\tilde{\mathcal{O}}(\log \log \log N)$ read efficiency, with a stronger condition of $\mathcal{O}(N^{1-1/\log \log \log N})$ on the size of the longest list.

Locality was introduced as a performance measure for memory accesses, assuming an implementation on Hard Disk Drives (HDDs). In [BBF⁺21], Bossuat et al. show that in the case of Solid State Drives (SSDs), such as flash disks, locality is no longer the relevant target. Instead, performance is mainly determined by the number of memory pages accessed, regardless of whether they are contiguous. In that setting the right performance metric is *page efficiency*. Page efficiency is defined as the number of pages read by the server to answer a query, divided by the number of pages needed to store the plaintext answer. The main construction of [BBF⁺21] achieves $\mathcal{O}(1)$ storage efficiency and $\mathcal{O}(1)$ page efficiency, assuming a client-side memory of $\omega(\log \lambda)$ pages.

To this day, a common point among all existing constructions, both local and page-efficient, is that they are purely *static*, as known techniques for sublogarithmic read efficiency and page efficiency do not apply to the dynamic setting. That may be because of the difficulty inherent in building local SSE, even in the static case (as evidenced, from the onset, by the impossibility result of Cash and Tessaro [CT14]). Nevertheless, many, if not most, applications of SSE require dynamism. This state of affairs significantly hinders the applicability of local and page-efficient SSE.

While one work [MM17] targets local SSE in a dynamic setting, and has constant storage efficiency and locality, it has read efficiency $\mathcal{O}(L \log W)$, where L is the maximum list size. Further, [MM17] employs an ORAM-variant which incurs a heavy computational overhead, in addition to the large read efficiency. When reinterpreting [MM17] in the context of page-efficiency, its guarantees improve to $\mathcal{O}(\log W)$ page efficiency and constant storage efficiency, but the heavy computational cost of ORAM remains.

1.1 Our Contributions

In this article, we consider the problem of dynamic memory-efficient SSE, by which we mean that we target both dynamic *page-efficient* SSE, and dynamic *local* SSE.

The centerpiece of our approach is a novel connection between these two goals. We introduce a map, called the Generic Local Transform, which takes as input a page-efficient SSE scheme with certain special features, and outputs a SSE scheme with strong locality properties. Our strategy will be to first build page-efficient schemes, then apply the Generic Local Transform to obtain local schemes. This approach turns out to be quite effective, and we present several results.

(1) **Dynamic page-efficient SSE.** We start by building a dynamic page-efficient SSE scheme, *LayeredSSE*. *LayeredSSE* achieves storage efficiency $\mathcal{O}(1)$,

and page efficiency $\tilde{\mathcal{O}}\left(\log \log \frac{N}{p}\right)$, where p is the page size. In line with prior work on memory-efficient SSE, the technical core of **LayeredSSE** is a new dynamic allocation scheme, **L2C**. **L2C** is a weighted variant of the so-called “2-choice” algorithm, notorious in the resource allocation literature. **L2C** is of independent interest: the two-choice allocation process is ubiquitous in various areas of computer science, such as load balancing, hashing, job allocation, or circuit routing (a survey of applications may be found in [RMS01]). Weighted variants have been considered in the past, but have so far required a *distributional* assumption [TW07, TW14] or presorting [ANSS16]. What we show is that by slightly tweaking the two-choice process, a dynamic and distribution-free result can be obtained (Theorem 1). Such a distribution-free result is necessary for cryptographic applications, where the adversary may influence the weights (as in our case). Other uses beyond cryptography are discussed in Appendix F.

(2) **Generic Local Transform.** We introduce the Generic Local Transform. On input any page-efficient scheme **PE-SSE** with certain special features, called *page-length-hiding* SSE, the Generic Local Transform outputs a local SSE scheme **Local[PE-SSE]**. Roughly speaking, if **PE-SSE** has client storage $\mathcal{O}(1)$, storage efficiency $\mathcal{O}(1)$, and page efficiency $\mathcal{O}(P)$, then **Local[PE-SSE]** has storage efficiency $\mathcal{O}(1)$, and read efficiency $\mathcal{O}(P)$. Regarding locality, the key feature is that if **PE-SSE** has locality $\mathcal{O}(L)$ *when querying lists of size at most one page*, then **Local[PE-SSE]** has locality $\mathcal{O}(L + \log \log N)$ *when querying lists of any size*. Thus, the **Local** construction may be viewed as bootstrapping a scheme with weak locality properties into a scheme with much stronger locality properties.

The Generic Local Transform also highlights an interesting connection between the goals of page efficiency and locality. Originally, locality and page efficiency were introduced as distinct performance criteria, targeting the two most widespread storage media, HDDs and SSDs respectively. It was already observed in [BBF⁺21] that a scheme with locality L and read efficiency R must have page efficiency at most $R + 2L$. In that sense, page efficiency is an “easier” goal. With the Generic Local Transform, surprisingly, we build a connection in the reverse direction: we use page-efficient schemes as building blocks to obtain local schemes. On a theoretical level, this shows a strong connection between the two goals. On a practical level, it provides a strategy to target both goals at once.

(3) **Dynamic local SSE.** By applying the Generic Local Transform to the **LayeredSSE** page-efficient scheme, we immediately obtain a dynamic SSE scheme **Local[LayeredSSE]**, with storage efficiency $\mathcal{O}(1)$, locality $\mathcal{O}(1)$, and read efficiency $\tilde{\mathcal{O}}(\log \log N)$. The construction is conditional: it requires that the longest list is of size $\mathcal{O}(N^{1-1/\log \log N})$. The asymptotic performance of **Local[LayeredSSE]** matches exactly the second *static* construction from [ANSS16], including the condition on maximum list size: dynamism comes at no extra cost. In particular, **Local[LayeredSSE]** matches the lower bound from [ASS21] for SSE schemes built using what [ASS21] refers to as “allocation schemes”—showing that the bound can be matched even in the dynamic setting.

(4) **Unconditional local SSE in the static setting.** The original 1-choice scheme from [ANSS16] achieves $\mathcal{O}(1)$ storage efficiency, $\mathcal{O}(1)$ locality, and $\tilde{\mathcal{O}}(\log N)$ read efficiency, unconditionally. The read efficiency was improved to $\mathcal{O}(\log^{2/3+\varepsilon} N)$ in [DPP18], for any constant $\varepsilon > 0$. This was, until now, the only SSE construction to achieve sublogarithmic efficiency unconditionally. By applying the Generic Local Transform to a variant of Tethys [BBF⁺21], in combination with techniques inspired by [DPP18], we obtain an unconditional static SSE scheme with storage efficiency $\mathcal{O}(1)$, locality $\mathcal{O}(1)$, and read efficiency $\mathcal{O}(\log^\varepsilon N)$, for any constant $\varepsilon > 0$. To our knowledge, this is the construction that comes closest to the impossibility result of Cash and Tessaro, stating that $\mathcal{O}(1)$ locality, storage efficiency, and read efficiency simultaneously is impossible.

Table 1 – Page-efficient SSE schemes. N denotes the total size of the database, W denotes the number of keywords, p is the number elements per page, $\varepsilon > 0$ is an arbitrarily small constant, and λ is the security parameter.

Schemes	Client st.	Page eff.	Storage eff.	Dynamism
$\Pi_{\text{pack}}, \Pi_{2\text{lev}}$ [CJJ ⁺ 14]	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(p)$	Static
TCA [ANSS16]	$\mathcal{O}(1)$	$\tilde{\mathcal{O}}(\log \log N)$	$\mathcal{O}(1)$	Static
Tethys [BBF ⁺ 21]	$\mathcal{O}(p \log \lambda)$	3	$3 + \varepsilon$	Static
IO-DSSE [MM17]	$\mathcal{O}(W)$	$\mathcal{O}(\log W)$	$\mathcal{O}(1)$	Dynamic
LayeredSSE	$\mathcal{O}(1)$	$\tilde{\mathcal{O}}\left(\log \log \frac{N}{p}\right)$	$\mathcal{O}(1)$	Dynamic

Table 2 – SSE schemes with constant locality and storage efficiency. N denotes the total size of the database, and $\varepsilon > 0$ is an arbitrarily small constant.

Schemes	Locality	Read eff.	St. eff.	Max list size	Dynamism
TCA [ANSS16]	$\mathcal{O}(1)$	$\tilde{\mathcal{O}}(\log \log N)$	$\mathcal{O}(1)$	$\mathcal{O}\left(N^{1-1/\log \log N}\right)$	Static
[ASS21]	$\mathcal{O}(1)$	$\tilde{\mathcal{O}}(\log \log \log N)$	$\mathcal{O}(1)$	$\mathcal{O}\left(N^{1-1/\log \log \log N}\right)$	Static
OCA [ANSS16]	$\mathcal{O}(1)$	$\tilde{\mathcal{O}}(\log N)$	$\mathcal{O}(1)$	Unconditional	Static
[DPP18]	$\mathcal{O}(1)$	$\tilde{\mathcal{O}}\left(\log^{2/3+\varepsilon} N\right)$	$\mathcal{O}(1)$	Unconditional	Static
Local[LayeredSSE]	$\mathcal{O}(1)$	$\tilde{\mathcal{O}}(\log \log N)$	$\mathcal{O}(1)$	$\mathcal{O}\left(N^{1-1/\log \log N}\right)$	Dynamic
UncondSSE	$\mathcal{O}(1)$	$\tilde{\mathcal{O}}(\log^\varepsilon N)$	$\mathcal{O}(1)$	Unconditional	Static

Remark on Forward Security. The SSE schemes built in this work have a standard “minimal” leakage profile during **Search**: namely, searches leak the search pattern, the access pattern and the length of the retrieved list of document identifiers. For our dynamic schemes, **Update** operations importantly leak no

information about unqueried keywords, but leak an identifier of the list being updated, as well as, in some cases, the length of the list. As a consequence, our dynamic schemes are not *forward-secure*. The underlying issue is that the goals of forward security and memory efficiency seem to be fundamentally at odds. Indeed, locality asks that identifiers associated to the same keywords must be stored close to each other; while forward-privacy requires that the location where a new identifier is inserted should be independent of the keyword it is associated with. That issue was already noted in [Bos16], who claims that “for dynamic schemes, locality and forward-privacy are two irreconcilable notions”. We refer the reader to [Bos16] for more discussion of the problem and leave further analysis of this issue for future work.

Note that SSE has a very varied range of uses cases, for example private database services, online messaging and encrypted text search. In practice, its security requirements depend entirely on the use case. There are use cases where forward secrecy is crucial. The argument for forward security that is often given in the literature (e.g. [Bos16, ?, ?, ?]) is to thwart file injection attacks in the style of [?]. Those attacks require injecting adversarially crafted entries into the target database. In an online messaging scenario, those attacks could be realistic, hence forward security is needed. In other cases, adversarial file injection is much less of a threat, and forward security can be reasonably dispensed with. For use cases where forward security is not required, we show that dynamism and memory efficiency are achievable at the same time.

Remark on the Focus on the Reverse Index. As most SSE literature, this work focuses on the (inverse) document index. The simplest usage scenario is to retrieve document indices from the index, then fetch those documents from a separate database. In reality, there are many other ways to use the index, for example by intersecting the document indices from several queries before fetching, fetching only some of the documents (see [?]), or building graph databases via several layers of inverse indices [?].

In most cases, the cost of fetching the actual documents is the same for the encrypted database as it is for the equivalent plaintext database: the efficiency overhead comes entirely from the inverse index. Schemes that hide access pattern or volume leakage are a possible exceptions but are out of the scope of this work.

2 Technical Overview

This work contains several results, tied together by the Generic Local Transform. As such, we believe it is beneficial to present them together within one paper. This requires introducing a number of different allocation mechanisms. We have endeavored to provide in this section a clear overview of those mechanisms. Formal specifications, theorems, and proofs will be presented in subsequent sections.

It is helpful to first recall a few well-studied allocation mechanisms. In what follows, “with overwhelming probability” is synonymous with “except with negligible probability” (in the usual cryptographic sense), whereas “with high proba-

bility” simply means with probability close to 1 in some sense, but not necessarily overwhelming.

One-choice allocation. In one-choice allocation, n balls are thrown into n bins. Each ball is inserted into a bin chosen independently and uniformly at random (by hashing an identifier of the ball). A standard analysis using Chernoff bounds shows that, at the outcome of the insertion process, the most loaded bin contains $\mathcal{O}(\log n)$ balls with high probability [JK77]. (And at most $\mathcal{O}(f(n) \log n)$ balls with overwhelming probability, for any $f = \omega(1)$.)

Two-choice allocation. Once again, n balls are thrown into n bins. For each ball, two bins are chosen independently and uniformly at random (e.g. by hashing an identifier of the ball). The ball is inserted into whichever of the two bins contains the fewest balls at the time of insertion. A celebrated result by Azar et al. shows that, at the outcome of the insertion process, the most loaded bin contains $\mathcal{O}(\log \log n)$ balls with high probability [ABKU94]. (It was later shown that the result holds with overwhelming probability [RMS01].)

2.1 Layered 2-Choice Allocation

Our first goal is to build a dynamic page-efficient scheme. Let us summarize what this entails, starting with the static case. As explained in the introduction, to realize single-keyword SSE, we want to store lists of arbitrary sizes on an untrusted server. Hiding the contents of the lists can be achieved in a straightforward way using symmetric encryption. The main challenge is how to store the lists in the server memory, in such a way that accessing one list does not reveal information about the lengths of other lists.

In the case of page-efficient schemes, this challenge may be summarized as follows. We are given a set of lists, containing N items in total. We are also given a page size p , which represents the number of items that can fit within a physical memory page. The memory of the server is viewed as an array of pages. We want to store the lists in the server memory, with three goals in mind.

1. In order to store all lists, we use $S \lceil N/p \rceil$ pages of server memory in total, where S is called the *storage efficiency* of the allocation scheme. We want S to be as small as possible.
2. Any list of length ℓ can be retrieved by visiting at most $P \lceil \ell/p \rceil$ pages in server memory, where P is called the *page efficiency* of the allocation scheme. We want P to be as small as possible.
3. Finally, the pages visited by the server to retrieve a given list should not depend on the lengths of other lists.

The first two goals are precisely the aim of bin packing algorithms. The third goal is a security goal: it stipulates that the pattern of memory accesses performed by the server should not leak certain information. As such, the goal relates to oblivious or data-independent algorithms. In [BBF⁺21], a framework for realizing the three goals was formalized as *Data-Independent Packing* (DIP).

To ease presentation, we will focus on the case where all lists are of size at most one page. If a list is of length more than one page, the general idea is that it will be split into chunks of one page, plus one final chunk of size at most one page; each chunk will then be treated as a separate list by the allocation scheme. We assume from now on that lists are of length less than one page.

In a nutshell, the idea proposed by [BBF⁺21] to instantiate a DIP scheme is to use weighted variant of cuckoo hashing [PR04]. In more detail, for each list, two pages are chosen uniformly at random, by hashing an identifier of the list. Each element of the list will then be stored in one of the two designated pages, or a stash. The stash is stored on the client side. In order to choose how each list is split between its three possible destinations (the two chosen pages, or the stash), [BBF⁺21] uses a maximum flow algorithm. The details of this algorithm are not relevant for our purpose. The important point is that when retrieving a list, the server accesses two uniformly random pages. Clearly, this reveals no information to the server about the lengths of other lists. The resulting algorithm, called Tethys, achieves storage efficiency $\mathcal{O}(1)$, page efficiency $\mathcal{O}(1)$, with client storage $\omega(\log \lambda)$ pages (used to store the stash).

In this paper, we wish to build a dynamic SSE. For that purpose, the underlying allocation scheme needs to allow for a new *update* operation. An update operation allows the client to add a new item to a list, increasing its length by one. The security goal remains essentially the same as in the static case: the pages accessed by the algorithm in order to update a given list should not depend on the lengths of other lists.

Tethys is not a suitable basis for a dynamic scheme, because it does not allow for an efficient data-independent update procedure: when inserting an element into a cell, the update procedure requires running a max flow algorithm. This either requires accessing other cells, with an access pattern that is intrinsically data-dependent, or performing a prohibitively expensive data-oblivious max flow computation each update. Instead, a natural idea is to use a weighted variant of the two-choice allocation scheme. With two-choice allocation, the access pattern made during an update is simple: only the two destination buckets associated to the list being updated need to be read. The new item is then inserted into whichever of the two buckets currently contains less items.

Instantiating that approach would require a weighted *dynamic* variant of two-choice allocation, along the following lines: given a multiset of list sizes $\{\ell_i : 1 \leq i \leq k\}$ with $\ell_i \leq p$ and $\sum \ell_i = N$, at the outcome of a two-choice allocation process into $\mathcal{O}(N/p)$ buckets, the most loaded bucket contains $\mathcal{O}(p \log \log N)$ items with overwhelming probability, even if the weight of balls is updated during the process. However, a result of that form appears to be a long-standing open problem (some related partial results are discussed in [BFHM08]). The two-choice process with weighted items has been studied in the literature [TW07, TW14, ANSS16], but to our knowledge, all existing results assume that (1) either the weight of the balls are sampled identically and independently from a sufficiently smooth distribution or (2) the balls are sorted initially and then allocated in decreasing order. Even disregarding constraints

on the distribution, in our setting, we cannot even afford to assume that list lengths are drawn independently: in the SSE security model, lists are chosen and updated *arbitrarily* by the adversary. Also, presorting the lists according to their length is not possible in a dynamic setting, as the list lengths can be changed via updates.

For our purpose, we require a *distribution-free* statement: we only know a bound p on the size of each list, and a bound N on the total size of all lists. We want an $\mathcal{O}(p \log \log N)$ upper bound on the size of the most loaded bucket that holds for *any* set of list sizes satisfying those constraints, even if list sizes are updated during the process. A result of that form is known for one-choice allocation processes [BFHM08] (with a $\mathcal{O}(p \log N)$ upper bound), but the same article shows that the same techniques cannot extend to the two-choice process.

To solve that problem, we introduce a *layered* weighted 2-choice allocation algorithm, L2C. L2C has the same basic behavior as a (weighted) two-choice algorithm: for each ball, two bins are chosen uniformly at random as possible destinations. The only difference is how the bin where the ball is actually inserted is selected among the two destination bins. The most natural choice would be to store the ball in whichever bin currently has the least load, where the *load* of a bin is the sum of the weights of the balls it currently contains. Instead, we use a slightly more complex decision process. In a nutshell, we partition the possible weights of balls into $\mathcal{O}(\log \log \lambda)$ subintervals, and the decision process is performed independently for balls in each subinterval. For the first subinterval (holding the smallest weights), we use a weighted one-choice process, while for the other subintervals, we use an unweighted two-choice process.

The point of this construction is that its analysis reduces to the analysis of the weighted one-choice process, and the unweighted two-choice process, for which powerful analytical techniques are known. We leverage those techniques to show that L2C achieves the desired distribution-free guarantees on the load of the most loaded bin. In practice, what this means is that we have an allocation algorithm that, for most intents and purposes, behaves like a weighted variant of two-choice allocation, and for which updates and distribution-free guarantees can be obtained relatively painlessly.

The LayeredSSE scheme is obtained by adding a layer of encryption and key management on top of L2C, using standard techniques from the SSE literature, although some care is required for updates. We refer the reader to Section 5 for more details.

2.2 Generic Local Transform

At Crypto 2018, Asharov et al. identified two main paradigms for building local SSE [ASS18]. The first is the *allocation* paradigm, which typically uses variants of multiple-choice allocation schemes, or cuckoo hashing. The second is the *pad-and-split* approach. The main difficulty of memory-efficient SSE is to pack together lists of different sizes. The idea of the pad-and-split approach is to store lists separately according to their size, which circumvents the issue. The simplest way to realize this is to pad all lists length to the next power of 2. This yields

$\log N$ possible values for list lengths. All lists of a given length can be stored together using, for instance, a standard hash table. Since we do not want to reveal the number of lists of each length, the hash table at each level needs to be dimensioned to be able to receive the entire database. As a result, a basic pad-and-split scheme has storage efficiency $\mathcal{O}(\log N)$, but easily achieves $\mathcal{O}(1)$ locality and read efficiency.

For the Generic Local Transform, we introduce the notion of *Overflowing SSE* (OSSE). An OSSE behaves like an SSE scheme in all aspects, except that, during its setup and during updates, it may refuse to store some list elements. Such elements are called *overflowing*. An OSSE is intended to be used as a subcomponent within an overarching SSE construction. The OSSE scheme is used to store part of the database, while overflowing elements are stored using a separate mechanism. The notion of OSSE was not formalized before, but in hindsight, the use of OSSE may be viewed as implicit in several existing constructions [DPP18, ASS18, BBF⁺21]. We choose to introduce it explicitly here for ease of exposition.

We are now in a position to explain the Generic Local Transform. The chief limitation of the pad-and-split approach is that it creates a $\log N$ overhead in storage. The high-level idea of the Generic Local Transform, then, is to use an OSSE to store all but a fraction $1/\log N$ of the database. Then a pad-and-split variant is used to store the $N/\log N$ overflowing elements. The intent is to benefit from the high efficiency of the pad-and-split approach, without having to pay for the $\log N$ storage overhead.

There is, however, a subtle but important issue with that approach. A given list may be either entirely stored within the OSSE scheme, or only partially stored, or not stored at all. In the OSSE scheme that we will later use (as well as OSSEs that were implicit in prior work), those three situations should be indistinguishable to the server, or else security breaks down. To address that issue, we proceed as follows.

Let us assume all lists have been padded to the next power of 2. For the pad-and-split part of the construction, we create $\log N$ SSE instances, one for each possible list size. We call each of these instances a *layer*. The overflowing elements of a list of size ℓ will be stored in the layer that handles lists of size ℓ , regardless of how many elements did overflow from the OSSE for that list.

The OSSE guarantees that the total number of overflowing items is at most $n = \mathcal{O}(N/\log N)$. Thus, if we focus on the layer that handles lists of size ℓ , the layer will receive at most n elements. These elements will be split into lists of size at most ℓ (corresponding to the set of overflowing elements, for each list of size ℓ in the original database). To achieve storage efficiency $\mathcal{O}(S)$ overall, we want the layer to store those lists using $\mathcal{O}(Sn)$ storage. To achieve read efficiency R , the layer should also be able to retrieve a given list by visiting at most $R\ell$ memory locations. This is where everything comes together: an SSE scheme satisfying those conditions is precisely a page-efficient SSE scheme with page size ℓ , storage efficiency S , and page efficiency R .

The page-efficient scheme used for each layer is also required satisfy a few extra properties: first, when searching for a list of size at most one page, the length of the list should not be leaked. We call this property *page-length-hiding*. (We avoid the term *length-hiding* to avoid confusion with volume-hiding SSE, which fully hides lengths.) All existing page-efficient constructions have that property. Second, we require the page-efficient scheme to have $\mathcal{O}(1)$ client storage. All constructions in this article satisfy that property, but the construction from [BBF⁺21] does not. Finally, we require the scheme to have locality $\mathcal{O}(1)$ when fetching a single page. All existing page-efficient constructions have this property. (The last two properties could be relaxed, at the cost of more complex formulas and statements.) We call an SSE scheme satisfying those three properties *suitable*.

Putting everything together, the Generic Local Transform takes as input a suitable *page-efficient* scheme, with storage efficiency S and page efficiency P . It outputs a *local* scheme with storage efficiency $S + S'$, read efficiency $P + R'$, and locality L' , where S' , R' , and L' are the storage efficiency, read efficiency, and locality of the underlying OSSE. It remains to explain how to build a local OSSE scheme with $\mathcal{O}(N/\log N)$ overflowing items, discussed next.

2.3 ClipOSSE: an OSSE scheme with $\mathcal{O}(N/\log N)$ Overflowing Items

At STOC 2016, Asharov et al. introduced so-called “2-dimensional” variants of one-choice and two-choice allocation, for the purpose of building local SSE. The one-choice variant works as follows. Consider an SSE database with N elements. Allocate $m = \tilde{\mathcal{O}}(N/\log N)$ buckets, initially empty. For each list of length ℓ in the database, choose one bucket uniformly at random. The first element of the list is inserted into that bucket. The second element of the list is inserted into the next bucket (assuming a fixed order of buckets, which wraps around when reaching the last bucket), the third one into the bucket after that, and so on, until all list elements have been inserted. Thus, assuming $\ell \leq m$, all list elements have been placed into ℓ consecutive buckets, one element in each. An analysis very similar to the usual analysis of the one-choice process shows that with overwhelming probability, the most loaded bucket receives at most $\tau = \tilde{\mathcal{O}}(\log N)$ elements. To build a static SSE scheme from this allocation scheme, each bucket is padded to the maximal size τ and encrypted. Search queries proceed in the natural way.

Such a scheme yields storage efficiency $\mathcal{O}(1)$, locality $\mathcal{O}(1)$ (since retrieving a list amounts to reading consecutive buckets), and read efficiency $\tilde{\mathcal{O}}(\log N)$ (since retrieving a list of length ℓ requires reading ℓ buckets, each of size $\tau = \tilde{\mathcal{O}}(\log N)$). To build ClipOSSE, we start from the same premise, but “clip” buckets at the threshold $\tau = \tilde{\mathcal{O}}(\log \log N)$. That is, each bucket can only receive up to τ elements. Elements that cannot fit are overflowing.

In the standard one-choice process, where n balls are thrown i.i.d. into n bins, it is not difficult to show that clipping bins at height $\tau = \mathcal{O}(\log \log n)$ results

in at most $\mathcal{O}(n/\log n)$ overflowing elements with overwhelming probability. In fact, by adjusting the multiplicative constant in the choice of τ , the number of overflowing elements can be made $\mathcal{O}(n/\log^d n)$ for any given constant d . We show that a result of that form still holds for (a close variant of) the 2-dimensional one-choice process outlined earlier. The result is conditional: it requires that the maximum list size is $\mathcal{O}(N/\text{polylog } N)$. (A condition of that form is necessary, insofar as the result fails when the maximum list size gets close to $N/\log N$.) The proof of the corresponding theorem is the most technically challenging part of this work, and relies on the combination of a convexity argument with a stochastic dominance argument. An overview of the proof is given in section 6.5, so we omit more discussion here.

In the end, ClipOSSE achieves storage efficiency $\mathcal{O}(1)$, locality $\mathcal{O}(1)$, and read efficiency $\mathcal{O}(\log \log N)$, with $\mathcal{O}(N/\log^d N)$ overflowing elements (for any fixed constant d of our choice), under the condition that the maximum list size is $\mathcal{O}(N/\text{polylog } N)$. All applications of the Generic Local Transform in this article use ClipOSSE as the underlying OSSE. (That is why we write Local[PE-SSE] for the Generic Local Transform applied to the page-efficient scheme PE-SSE, and do not put the underlying OSSE as an explicit parameter.)

2.4 Dynamic Local SSE with $\tilde{\mathcal{O}}(\log \log N)$ Overhead

By using the Generic Local Transform with ClipOSSE as the underlying OSSE, and LayeredSSE as the page-efficient scheme, we obtain Local[LayeredSSE]. The Local[LayeredSSE] scheme has storage efficiency $\mathcal{O}(1)$, locality $\mathcal{O}(1)$, and read efficiency $\tilde{\mathcal{O}}(\log \log N)$. This result follows from the main theorem regarding the Generic Local Transform, and does not require any new analysis.

Local[LayeredSSE] is a conditional scheme: it requires that the longest list is of length $\mathcal{O}(N^{1-1/\log \log \lambda})$. The reason is subtle. ClipOSSE by itself has a condition that the longest list is $\mathcal{O}(N/\text{polylog } N)$, which is less demanding. The reason for the condition comes down to the fact that LayeredSSE only achieves a negligible probability of failure as long as the number of pages in the scheme is at least $\Omega(\lambda^{1/\log \log \lambda})$. More generally, the same holds for the number of bins in two-choice allocation processes in general, even the standard, unweighted process. The condition is optimal: [ASS21] shows that any sublogarithmic “allocation-based” scheme must be conditional, and gives a bound on the condition. Local[PE-SSE] matches that bound.

2.5 Unconditional Static Local SSE with $\mathcal{O}(\log^\varepsilon N)$ Overhead

The (static) Tethys scheme from [BBF⁺21] achieves storage efficiency $\mathcal{O}(1)$ and page efficiency $\mathcal{O}(1)$ simultaneously. It is also page-length-hiding. Since we have the Generic Local Transform at our disposal, it is tempting to apply it to Tethys. There is, however, one obstacle: Tethys uses $\omega(p \log \lambda)$ client memory, in order to store a stash on the client side. For the Generic Local Transform, we need

$\mathcal{O}(1)$ client memory. To reduce the client memory of Tethys, a simple idea is to store the stash on the server side. Naively, reading the stash for every search would increase the page efficiency to $\omega(\log \lambda)$. To avoid this, we store the stash within an ORAM.

For that purpose, we need an ORAM with a failure probability of zero: indeed, since we may store as few as $\log \lambda$ elements in the ORAM, a correctness guarantee of the form $\text{negl}(n)$, where $n = \log \lambda$ is the number items in the ORAM, fails to be sufficient (it is not $\text{negl}(\lambda)$). We also need the ORAM to have $\mathcal{O}(1)$ locality. An ORAM with these characteristics was devised in [DPP18], motivated by the same problem. The ORAM from [DPP18] achieves read efficiency $\mathcal{O}(n^{1/3+\varepsilon})$, for any arbitrary constant $\varepsilon > 0$. It was already conjectured in [DPP18] that it could be improved to $\mathcal{O}(n^\varepsilon)$. We build that variant explicitly, and name it LocORAM. Roughly speaking, LocORAM is a variant of the Goldreich-Ostrovsky hierarchical ORAM, with a constant number of levels.

By putting the stash of Tethys within LocORAM on the server side, we naturally obtain a page-efficient SSE scheme `OramTethys`, with $\mathcal{O}(\log^\varepsilon \lambda)$ read efficiency, suitable for use within the Generic Local Transform. This yields a static local SSE for lists of size at most $N/\text{polylog } N$. To handle larger lists, borrowing some ideas from [DPP18], we group lists by size, and use again `OramTethys` to store them. In the end, we obtain an unconditional SSE with $\mathcal{O}(1)$ store efficiency, $\mathcal{O}(1)$ locality, and $\mathcal{O}(\log^\varepsilon \lambda)$ read efficiency.

Comparing with the $\mathcal{O}(\log^{2/3+\varepsilon} \lambda)$ construction from [DPP18], we note that the bottleneck of their construction comes from the allocation schemes the authors use for what they call “small” and “medium” lists. This is precisely the range where we use `Local[OramTethys]`. Our construction essentially removes that bottleneck, so that the $\mathcal{O}(\log^\varepsilon \lambda)$ read efficiency bottleneck now comes entirely from the ORAM component.

3 Preliminaries

Let $\lambda \in \mathbb{N}$ be the security parameter. For a probability distribution X , we denote by $x \leftarrow X$ the process of sampling a value x from the distribution. Further, we say that x is We denote by $[a, b]_{\mathbb{R}}$ the interval $\{x \in \mathbb{R} \mid a \leq x \leq b\}$ and extend this naturally to intervals of the form $[a, b)_{\mathbb{R}}, (a, b]_{\mathbb{R}}, (a, b)_{\mathbb{R}}$.

3.1 Symmetric Searchable Encryption

A database $\text{DB} = \{w_i, (\text{id}_1, \dots, \text{id}_{\ell_i})\}_{i=1}^W$ is a set of keyword-identifier pairs with W keywords. We assume that each keyword w_i is represented by a machine word of $\mathcal{O}(\lambda)$ bits. We write $\text{DB}(w_i) = (\text{id}_1, \dots, \text{id}_{\ell_i})$ for the list of identifiers matching w_i . Throughout the article, we set $N = \sum_{i=1}^W \ell_i$ and define p as the page size (which we treat as a variable, independent of the size of the database N).

A dynamic searchable symmetric encryption scheme Σ is a 4-tuple of PPT algorithms (`KeyGen`, `Setup`, `Search`, `Update`) such that

- $\Sigma.\text{KeyGen}(1^\lambda)$: Takes as input the security parameter λ and outputs client secret key K .
- $\Sigma.\text{Setup}(K, N, \text{DB})$: Takes as input the client secret key K , an upper bound on the database size N and a database DB . Outputs encrypted database EDB and client state st .
- $\Sigma.\text{Search}(K, w, \text{st}; \text{EDB})$: The client receives as input the secret key K , keyword w and state st . The server receives as input the encrypted database EDB . Outputs some data d and updated state st' for the client. Outputs updated encrypted database EDB' for the server.
- $\Sigma.\text{Update}(K, (w, L), \text{op}, \text{st}; \text{EDB})$: The client receives as input the secret key K , a pair (w, L) of keyword w and list L of identifiers, an operation $\text{op} \in \{\text{del}, \text{add}\}$ and state st . The server receives as input the encrypted database EDB . Outputs updated state st' for the client. Outputs updated encrypted database EDB' for the server.

In the following, we omit the state st and assume that it is implicitly stored and updated by the client. We say that Σ is *static*, if it does not provide an `Update` algorithm. Further, we assume that the keyword w is preprocessed via a PRF by the client, whenever the client sends w to the server in either `Search` or `Update`. This ensures that the server never has access to w in plaintext and unqueried keywords are distributed uniformly random in the view of the server.

Intuitively, the client uses `Setup` to encrypt and outsource a database DB to the server. Then, the client can search keywords w using `Search` and receives the list of matching identifiers $\text{DB}(w)$ from the server. The list $\text{DB}(w)$ can be updated via `Update`, provided that the size of the database stays below N . Note that we allow the client to add (or delete) multiple identifiers at once for a single keyword (which is required for the Generic Local Transform section 6).

Security. We now define correctness and semantic security of SSE. Intuitively, correctness guarantees that a search always retrieves all matching identifiers and semantic security guarantees that the server only learns limited information (quantified by a leakage function) from the client.

Definition 1 (Correctness). *A SSE scheme Σ is correct if for all databases DB and $N \in \mathbb{N}$, keys $K \leftarrow \Sigma.\text{KeyGen}(1^\lambda)$, $\text{EDB} \leftarrow \Sigma.\text{Setup}(K, \text{DB})$ and sequences of search, add or delete queries S , the search protocol returns the correct result for all queries of the sequence, if the size of the database remains at most N .*

We use the standard semantic security notion for SSE (see [CGKO06]). Security is parameterized by a leakage function $\mathcal{L} = (\mathcal{L}_{\text{Stp}}, \mathcal{L}_{\text{Srch}}, \mathcal{L}_{\text{Updt}})$, composed of the setup leakage \mathcal{L}_{Stp} , the search leakage $\mathcal{L}_{\text{Srch}}$, and the update leakage $\mathcal{L}_{\text{Updt}}$. We define two games, `SSEReal` and `SSEIdeal`. First, the adversary chooses a database DB . In `SSEReal`, the encrypted database EDB is generated by `Setup`(K, N, DB), whereas in `SSEIdeal` the encrypted database is simulated by a (stateful) simulator `Sim` on input $\mathcal{L}_{\text{Stp}}(\text{DB}, N)$. After receiving EDB , the adversary issues search and update queries. All queries are answered honestly in

SSEReAL. In SSEIDEAL, the search queries on keyword w are simulated by Sim on input $\mathcal{L}_{\text{Srch}}(w)$, and update queries for operation op , keyword w and identifier list L are simulated by Sim on input $\mathcal{L}_{\text{Updt}}(\text{op}, w, L)$. Finally, the adversary outputs a bit b .

We write $\text{SSEReAL}^{\text{adp}}$ and $\text{SSEIDEAL}^{\text{adp}}$ if the queries of the adversary were chosen adaptively, *i.e.* dependant on previous queries. Similarly, we write $\text{SSEReAL}^{\text{sel}}$ and $\text{SSEIDEAL}^{\text{sel}}$ if the queries are chosen selectively by the adversary, *i.e.* sent initially in conjunction with the database before receiving EDB.

Definition 2 (Semantic Security). *Let Σ be a SSE scheme and $\mathcal{L} = (\mathcal{L}_{\text{Stp}}, \mathcal{L}_{\text{Srch}}, \mathcal{L}_{\text{Updt}})$ a leakage function. Scheme Σ is \mathcal{L} -adaptively secure if for all PPT adversaries \mathcal{A} , there exists a PPT simulator Sim such that*

$$|\Pr[\text{SSEReAL}_{\Sigma, \mathcal{A}}^{\text{adp}}(\lambda) = 1] - \Pr[\text{SSEIDEAL}_{\Sigma, \text{Sim}, \mathcal{L}, \mathcal{A}}^{\text{adp}}(\lambda) = 1]| = \text{negl}(\lambda).$$

Similarly, scheme Σ is \mathcal{L} -selectively secure if for all PPT adversaries \mathcal{A} , there exists a PPT simulator Sim such that

$$|\Pr[\text{SSEReAL}_{\Sigma, \mathcal{A}}^{\text{sel}}(\lambda) = 1] - \Pr[\text{SSEIDEAL}_{\Sigma, \text{Sim}, \mathcal{L}, \mathcal{A}}^{\text{sel}}(\lambda) = 1]| = \text{negl}(\lambda).$$

Intuitively, semantic security guarantees that the interaction between client and server reveals no information to the server, except the leakage of the given query. The schemes from this article have common leakage patterns. We use the standard notions of query pattern qp and history Hist from [Bos16] to formalize this leakage: (1) The query pattern $\text{qp}(w)$ for a keyword w are the indices of previous search or update queries for keyword w . (2) The history $\text{Hist}(w)$ is comprised of the list of identifiers matching keyword w that were inserted during setup and the history of updates on keyword w , that is each deleted and inserted identifier. We can retrieve the number ℓ_i of inserted identifiers and the number d_i of deleted identifiers from $\text{Hist}(w)$ for each keyword.

We define two leakage patterns we use throughout the article. (1) We define *page-length hiding* leakage $\mathcal{L}_{\text{len-hid}}$. We set $\mathcal{L}_{\text{len-hid}} = (\mathcal{L}_{\text{Stp}}^{\text{len-hid}}, \mathcal{L}_{\text{Srch}}^{\text{len-hid}}, \mathcal{L}_{\text{Updt}}^{\text{len-hid}})$, where the setup leakage is $\mathcal{L}_{\text{Stp}}^{\text{len-hid}}(\text{DB}, N) = N$ is the maximal size N of the database, the search leakage $\mathcal{L}_{\text{Srch}}^{\text{len-hid}}(w) = (\text{qp}, \lceil \ell_i/p \rceil, \lceil d_i/p \rceil)$ is the query pattern and the number of pages required to store the inserted and deleted items, and the update leakage $\mathcal{L}_{\text{Updt}}^{\text{len-hid}}(\text{op}, w, L) = (\text{op}, \text{qp}, \lceil (\ell_i + |L|)/p \rceil, \lceil (d_i + |L|)/p \rceil, \lceil \ell_i/p \rceil, \lceil d_i/p \rceil)$ is the operation, the query pattern and the number of pages required to store the inserted and deleted items (before and after the update)¹. (2) Similarly, we define *length revealing* leakage $\mathcal{L}_{\text{len-rev}}$. We set $\mathcal{L}_{\text{len-rev}} = (\mathcal{L}_{\text{Stp}}^{\text{len-rev}}, \mathcal{L}_{\text{Srch}}^{\text{len-rev}}, \mathcal{L}_{\text{Updt}}^{\text{len-rev}})$ with $\mathcal{L}_{\text{Stp}}^{\text{len-rev}}(\text{DB}, N) = N$, $\mathcal{L}_{\text{Srch}}^{\text{len-rev}}(w) = (\text{qp}, |L'|, \ell_i, d_i)$ and lastly $\mathcal{L}_{\text{Updt}}^{\text{len-rev}}(\text{op}, w, L') = (\text{op}, \text{qp}, |L'|, \ell_i, d_i)$.

We will use $\mathcal{L}_{\text{len-hid}}$ and $\mathcal{L}_{\text{len-rev}}$ for both dynamic and static schemes. When we say that a static scheme is \mathcal{L} -semantically secure, for $\mathcal{L} \in \{\mathcal{L}_{\text{len-hid}}, \mathcal{L}_{\text{len-rev}}\}$, we

¹ Note that we allow for inserting more than one identifier per keyword in a single update operation in this work. Thus, the server will also learn (limited) information about the number $|L|$ of added or deleted identifiers.

simply ignore the update leakage. Note that both leakage patterns, $\mathcal{L}_{\text{len-hid}}$ and $\mathcal{L}_{\text{len-rev}}$, have standard setup and search leakage, common in most SSE schemes. The update leakage of $\mathcal{L}_{\text{len-hid}}$ and $\mathcal{L}_{\text{len-rev}}$ is similar to their search leakage, and reveals nothing about unqueried keywords. While the update leakage is not forward secure, similar leakage patterns are commonly considered in literature, for example [CJJ⁺14]. We hope our techniques pave the way for future work on dynamic schemes with forward security and memory efficiency.

Efficiency Measures. We recall the notions of locality, storage efficiency and read efficiency [CT14], and page efficiency [BBF⁺21] (and extend them to the dynamic SSE setting in a natural manner). In the following definitions, we set $K \leftarrow \text{KeyGen}(1^\lambda)$ and $\text{EDB} \leftarrow \text{Setup}(K, N, \text{DB})$ given database DB and upper bound N on the number of document identifiers. Also, $S = (\text{op}_i, \text{in}_i)_{i=1}^s$ is a sequence of search and update queries, where $\text{op}_i \in \{\text{add}, \text{del}, \perp\}$ is a operation and $\text{in}_i = (\text{op}_i, w_i, L_i, \text{st}_i, \text{EDB}_i)$ its input. Here, w_i is a keyword and L_i is a (added or deleted) list of identifiers, and after executing all previous operations op_j for $j \leq i$, st_i is the client state and EDB_i the encrypted database. We denote by DB_i the database after i operations. We assume that the total number of identifiers never exceeds N . (If $\text{op}_i = \perp$, the query is a search query and L_i is empty.)

Definition 3 (Read Pattern). *Regard server-side storage as an array of memory locations, containing the encrypted database EDB . When processing search query $\text{Search}(K, w_i, \text{st}_i; \text{EDB}_i)$ or update query $\text{Update}(K, (w_i, L_i), \text{op}_i, \text{st}_i; \text{EDB}_i)$, the server accesses memory locations m_1, \dots, m_h . We call these locations the read pattern and denote it with $\text{RdPat}(\text{op}_i, \text{in}_i)$.*

Definition 4 (Locality). *A SSE scheme has locality L if for any λ , DB , N , sequence S , and any i , $\text{RdPat}(\text{op}_i, \text{in}_i)$ consists of at most L disjoint intervals.*

Definition 5 (Read Efficiency). *A SSE scheme has read efficiency R if for any λ , DB , N , sequence S , and any i , $|\text{RdPat}(\text{op}_i, \text{in}_i)| \leq R \cdot P$, where P is the number of memory locations needed to store all (added and deleted) document indices matching keyword w_i in plaintext (by concatenating indices).*

Definition 6 (Storage Efficiency). *A SSE scheme has storage efficiency E if for any λ , DB , N , sequence S , and any i , $|\text{EDB}_i| \leq E \cdot |\text{DB}_i|$.*

Definition 7 (Page Pattern). *Regard server-side storage as an array of pages, containing the encrypted database EDB . When processing search query $\text{Search}(K, w_i, \text{st}_i; \text{EDB}_i)$ or update query $\text{Update}(K, (w_i, L_i), \text{op}_i, \text{st}_i; \text{EDB}_i)$, the read pattern $\text{RdPat}(\text{op}_i, \text{in}_i)$ induces a number of page accesses p_1, \dots, p_h . We call these pages the page pattern, denoted by $\text{PgPat}(\text{op}_i, \text{in}_i)$.*

Definition 8 (Page Cost). *A SSE scheme has page cost $aX + b$, where a, b are real numbers, and X is a fixed symbol, if for any λ , DB , N , sequence S , and any i , $|\text{PgPat}(\text{op}_i, \text{in}_i)| \leq aX + b$, where X is the number of pages needed to store document indices matching keyword w_i in plaintext.*

Definition 9 (Page Efficiency). *A SSE scheme has page efficiency P if for any λ , DB , N , sequence S , and any i , $|\text{PgPat}(\text{op}_i, \text{in}_i)| \leq P \cdot X$, where X is the number of pages needed to store document indices matching keyword w_i in plaintext.*

4 Layered Two-Choice Allocation

In this section, we describe layered two-choice allocation (L2C), a variant of two-choice allocation that allows to allocate n weighted balls (b_i, w_i) into m bins, where b_i is a unique identifier and $w_i \in [0, 1]_{\mathbb{R}}$ is the weight of the ball. (We often write ball b_i for short.) First, let $1 \leq \delta(\lambda) \leq \log(\lambda)$ be a function. We denote by $w = \sum_{i=1}^n w_i$ the sum of all weights and set $m = w/(\delta(\lambda) \log \log w)$. We will later choose $\delta(\lambda) = o(\log \log \lambda)$ such that allocation has negligible failure probability. In the overview, we set $\delta(\lambda) = 1$ and assume that $m = \Omega(\lambda)$ for simplicity (which suffices for negligible failure probability).

Overview of L2C. L2C is based on both *weighted* one-choice allocation (1C) and *unweighted* two-choice allocation (2C). On a high level, we split the set of possible weights $[0, 1]_{\mathbb{R}}$ into $\log \log m$ subintervals

$$[0, 1/\log m]_{\mathbb{R}}, (1/\log m, 2/\log m]_{\mathbb{R}}, \dots, (2^{\log \log m - 1}/\log m, 1]_{\mathbb{R}}.$$

In words, the first interval is of size $1/\log m$ and the boundaries between intervals grow by a factor 2 every time. We will allocate balls with weights in a given subinterval independently from the others.

Balls in the first subinterval have weights $w_i \leq 1/\log m$ and are thus small enough to apply weighted 1C. Intuitively, this suffices because one-choice (provably) performs worst for uniform weights of maximal size $1/\log m$. In that case, there are at most $n' = w \log m$ balls and we expect a bin to contain $n'/m = \log m \cdot \log \log w$ balls of uniform weight, since $m = w/(\log \log w)$. As each ball has weight $1/\log m$, the expected load per bin is $\log \log w$. This translates to a $\mathcal{O}(\log \log w)$ bound with overwhelming probability after applying a Chernoff's bound.

For the other intervals, applying unweighted and independent 2C per interval suffices, as the weights of balls differ at most by a factor 2 and there are only $\log \log m$ intervals. More concretely, let n_i be the number of balls in the i -th subinterval $A_i = (2^{i-1}/\log m, 2^i/\log m]_{\mathbb{R}}$ for $i \in \{1, \dots, \log \log m\}$. Balls with weights in subinterval A_i fill the bins with at most $\mathcal{O}(n_i/m + \log \log m)$ balls, independent of other subintervals. Note that we are working with small weights, and thus potentially have $\omega(m)$ balls. Thus, we need to extend existing 2C results to negligible failure probability in m for the heavily-loaded case (cf. lemma 5). As there are only $\log \log m$ subintervals, and balls in interval A_i have weight at most $2^i/\log m$, we can just sum the load of each subinterval and receive a bound

$$\sum_{i=1}^{\log \log m} \frac{2^i}{\log m} \mathcal{O}(n_i/m + \log \log m) = \mathcal{O}(w/m + \log \log m).$$

In total, we have $\mathcal{O}(w/m + \log \log m) = \mathcal{O}(\log \log w)$ bounds for the first and the remaining intervals. Together, this shows that all bins have load at most $\mathcal{O}(\log \log w)$ after allocating all n items. This matches the bound of standard 2C with unweighted balls if $m = \Omega(\lambda)$. For our SSE application, we want to allow for negligible failure probability with the least number of bins possible. We can set $\delta(\lambda) = \log \log \log(\lambda)$ and obtain a bin size of $\tilde{\mathcal{O}}(\log \log w)$ with overwhelming probability, if $m = \frac{w}{\delta(\lambda) \log \log w}$. The analysis is identical in this case.

Handling Updates. The described variant of L2C is static. That is, we have not shown a bound on the load of the most loaded bin if we add balls or update the weight of balls. Fortunately, inserts of new balls are trivially covered by the analysis sketched above, if m was chosen large enough initially in order to compensate for the added weight. Thus, we assume there is some upper bound w_{\max} on the total weights of added balls which is used to initially set up the bins. We can also update weights if we proceed with care.

For this, let b_i be some ball with weight w_{old} . We want to update its weight to $w_{\text{new}} > w_{\text{old}}$. If w_{old} and w_{new} reside in the subinterval, we can directly update the weight of b_i , as L2C ignores the concrete weight of balls inside a given subinterval for the allocation. Indeed, in the first interval, the bin in which b_i is inserted is determined by a single random choice, and for the remaining subintervals, the 2C process only considers the number of balls inside the same subinterval, ignoring concrete weights.

When w_{new} is larger than the bounds of the current subinterval, we need to make sure that the ball is inserted into the correct bin of its two choices. For this, the ball b_i is inserted into the bin with the lowest number of balls with weights inside the new subinterval. Even though the bin of b_i might change in this process, we still need to consider b_i as a ball of weight w_{old} in the old bin for subsequent ball insertions in the old subinterval. Thus, we mark the ball as *residual* ball but do not remove it from its old bin. That is, we consider it as ball of weight w_{old} for the 2C process but assume it is not identified by b_i anymore. As there are only $\log \log m$ different subintervals, storing the residual balls has a constant overhead. The full algorithm L2C is given in Algorithm 1. We parameterize it by a hash function H mapping uniformly into $\{1, \dots, m\}^2$. The random bin choices of a ball b_i are given by $\alpha_1, \alpha_2 \leftarrow H(b_i)$.

Load Analysis of L2C. Let either $\delta(\lambda) = 1$ or $\delta(\lambda) = \log \log \log \lambda$ and m sufficiently large such that $m^{-\Omega(\delta(\lambda) \log \log w)} = \text{negl}(\lambda)$. (Note that this is the probability that allocation of 1C and 2C fails.)

We need to show that after setup and during a (selective) sequence of operations, the most loaded bin has a load of at most $\mathcal{O}(\delta(\lambda) \log \log w_{\max})$, where w_{\max} is an upper bound on the total weight of the inserted balls. We sketch the proof here and refer to Appendix C for further details. First, we modify the sequence S such that we can reduce the analysis to only (sufficiently independent) L2C.InsertBall operations, while only increasing the final bin load by a constant factor. This is constant factor of the load is due to the additional weight of residual balls. Then, we analyze the load of the most loaded bin for the each

subinterval independently. This boils down to an analysis of a 1C process in the first subinterval and a 2C process in the remaining subintervals as in the overview of L2C (see Section 4). Summing up the independent bounds yields the desired result.

Theorem 1. *Let either $\delta(\lambda) = 1$ or $\delta(\lambda) = \log \log \log \lambda$. Let $w_{\max} = \text{poly}(\lambda)$ and $m = w_{\max}/(\delta(\lambda) \log \log w_{\max})$. We require that $m = \Omega(\lambda^{\frac{1}{\log \log \lambda}})$ if $\delta(\lambda) = \log \log \log \lambda$ or $m = \Omega(\lambda)$ otherwise. Let $\{(b_i, w_i)_{i=1}^n\}$ be balls with (pair-wise unique) identifier b_i and weight $w_i \in [0, 1]$. Further, let $S = (\text{op}_i, \text{in}_i)_{i=n+1}^{s+n}$ be a sequence of s insert or update operations $\text{op}_i \in \{\text{L2C.InsertBall}, \text{L2C.UpdateBall}\}$ with input $\text{in}_i = (b_i, w_i, B_{\alpha_{i,1}}, B_{\alpha_{i,2}})$ for inserts and $\text{in}_i = (b_i, o_i, w_i, B_{\alpha_{i,1}}, B_{\alpha_{i,2}})$ for updates. Here, b_i denotes the identifier of a ball with weight w_i and old weight $o_i \leq w_i$ before the execution of op_i . Also, the bins are chosen via $\alpha_{i,1}, \alpha_{i,2} \leftarrow \text{H}(b_i)$.*

Execute $(B_i)_{i=1}^m \leftarrow \text{L2C.Setup}(\{(b_i, w_i)_{i=1}^n\})$ and the operations $\text{op}_i(\text{in}_i)$ for all $i \in [n+1, n+s]$. We require that $\sum_{i=1}^{n+s} w_i - o_i \leq w_{\max}$, i.e. the total weight after all operations is at most w_{\max} .

Then it holds that throughout the process, the most loaded bin of B_1, \dots, B_m has at most load $\mathcal{O}(\delta(\lambda) \log \log w_{\max})$ except with negligible probability, if H is modeled as a random oracle.

5 Dynamic Page Efficient SSE – Overview

We introduce the SSE scheme LayeredSSE based on L2C. Essentially, we interpret lists L_i of identifiers matching keyword w_i as balls of a certain weight. Then, we use L2C to manage the balls in m encrypted bins, where each bin corresponds to a memory page, yielding page efficiency $\tilde{\mathcal{O}}(\log \log N/p)$ and constant storage efficiency. Let N be the maximal size of the database, $p \leq N^{1-1/\log \log \lambda}$ be the page size² and H be a hash function mapping into $\{1, \dots, m\}^2$ for $m = \lceil w_{\max}/(\log \log \log \lambda \cdot \log \log w_{\max}) \rceil$ and $w_{\max} = N/p$. Due to space limitations, we assume that each keyword has at most p associated keywords, and outline the scheme and its security analysis. We refer to Appendix A for details (without restrictions on the database³).

For convenience, we adapt the notation of L2C to lists of identifiers. A ball (w, L) of weight $|L|/p \in [0, 1]_{\mathbb{R}}$ is a list of (at most p) identifiers matching keyword w . The 2 bin choices α_1, α_2 for ball (w, L) are given via $(\alpha_1, \alpha_2) \leftarrow \text{H}(w)$. Now, L2C.Setup takes input balls $\{(w_i, L_i)\}_{i=1}^W$ and maximal weight w_{\max} , and allocates them as before into m bins. L2C.InsertBall receives ball (w, L) and two bins $(B_{\alpha_1}, B_{\alpha_2})$, and inserts (w, L) into either bin B_{α_1} or bin B_{α_2} as before.

² This condition is needed for the requirement $m \geq \lambda^{1/\log \log \lambda}$ of L2C which guarantees negligible failure probability (see Theorem 1). In practice, we have $p \ll N$.

³ For arbitrary lists sizes, we can split lists into sublists of size at most p and deal with each sublist separately as before. Some care has to be taken, for example with the random choices of the bins, but details are mostly straightforward.

Algorithm 1 Layered 2-Choice Allocation (L2C)**L2C.Setup**($\{(b_i, w_i)\}_{i=1}^n, w_{\max}$)

- 1: Receive n balls (b_i, w_i) , and maximal total weight w_{\max}
- 2: Initialize $m = \lceil w_{\max}/(\delta(\lambda) \log \log w_{\max}) \rceil$ empty bins B_1, \dots, B_m
- 3: **for all** $i \in \{1, \dots, n\}$ **do**
- 4: Set $\alpha_1, \alpha_2 \leftarrow \mathbf{H}(b_i)$
- 5: **InsertBall**($b_i, w_i, B_{\alpha_1}, B_{\alpha_2}$)
- 6: Return B_1, \dots, B_m

L2C.InsertBall($b_{\text{new}}, w_{\text{new}}, B_{\alpha_1}, B_{\alpha_2}$)

- 1: Receive bins $B_{\alpha_1}, B_{\alpha_2}$, and ball $(b_{\text{new}}, w_{\text{new}})$
- 2: Assert that α_1, α_2 are the choices given by $\mathbf{H}(b_{\text{new}})$
- 3: Split the set of possible weights $[0, 1]_{\mathbb{R}}$ into $\log \log m$ sub-intervals

$$[0, 1/\log m]_{\mathbb{R}}, (1/\log m, 2/\log m]_{\mathbb{R}}, \dots, (2^{\log \log m - 1}/\log m, 1]_{\mathbb{R}}$$

- 4: Choose $k \in \mathbb{N}$ minimal such that $w_{\text{new}} \leq 2^k/\log m$
- 5: **if** $k = 1$ **then**
- 6: Set $\alpha \leftarrow \alpha_1$
- 7: **else**
- 8: Let B_{α} be the bin with the least number of balls of weight in $\left(\frac{2^{k-1}}{\log m}, \frac{2^k}{\log m}\right]_{\mathbb{R}}$ among B_{α_1} and B_{α_2}
- 9: Insert ball b_{new} into bin B_{α}

L2C.UpdateBall($b_{\text{old}}, w_{\text{old}}, w_{\text{new}}, B_{\alpha_1}, B_{\alpha_2}$)

- 1: Receive bins $B_{\alpha_1}, B_{\alpha_2}$ that contain ball $(b_{\text{old}}, w_{\text{old}})$, and new weight $w_{\text{new}} \geq w_{\text{old}}$
- 2: Assert that α_1, α_2 are the choices given by $\mathbf{H}(b_{\text{old}})$
- 3: **if** $w_{\text{old}}, w_{\text{new}} \in \left(\frac{2^{k-1}}{\log m}, \frac{2^k}{\log m}\right]_{\mathbb{R}}$ for some k **then**
- 4: Update the weight of b_{old} to w_{new} directly
- 5: **else**
- 6: Mark b_{old} as residual ball (it is still considered as a ball of weight w_{old})
- 7: **InsertBall**($b_{\text{old}}, w_{\text{new}}, B_{\alpha_1}, B_{\alpha_2}$)

L2C.UpdateBall receives old ball (w, L) , identifiers L' and bins $(B_{\alpha_1}, B_{\alpha_2})$, and updates ball (w, L) to ball $(w, L \cup L')$ as before, while merging both identifier lists L and L' . (The weight of the updated ball is $|L \cup L'|/p \in [0, 1]_{\mathbb{R}}$.)

5.1 LayeredSSE

We describe **LayeredSSE**, focusing on insert operations. In Appendix A, we describe **LayeredSSE** in more detail, and show how to treat arbitrary list sizes, introduce delete operations and show how to obtain updates in 1 RTT. A detailed description of **LayeredSSE** is given in algorithm 2.

Setup. To setup the initial database $\text{DB} = (w, L_i)_{i=1}^W$, given upperbound N on the number of keyword-identifiers, allocate the balls (w, L_i) into m bins via **L2C**. Next, each bin is filled up to maximal size $p \cdot c \log \log \log(\lambda) \log \log(N/p)$, for some constant c . Finally, the encrypted bins are output.

Search. During a search operation on keyword w , the client retrieves encrypted bins $B_{\alpha_1}, B_{\alpha_2}$ for $(\alpha_1, \alpha_2) \leftarrow H(w)$ from the server.

Update. During an update operation to add identifier list L' to keyword w , the client retrieves $B_{\alpha_1}, B_{\alpha_2}$, decrypts both bins and retrieves ball (w, L) from the corresponding bin $B_\alpha \in \{B_{\alpha_1}, B_{\alpha_2}\}$. Then, she calls `L2C.UpdateBall` with old ball (w, L) , new identifiers L' and bins $B_{\alpha_1}, B_{\alpha_2}$ to insert the new identifiers L' into one of the bins. Finally, she reencrypts the bins and sends them to the server. The server then replaces the old bins with the updated bins.

Algorithm 2 LayeredSSE

LayeredSSE.KeyGen(1^λ)	LayeredSSE.Update($K, (w, L'), \text{add}; \text{EDB}$)
<p style="text-align: center;">Global parameters: constant $c \in \mathbb{N}$, page size p</p> <p>1: Sample K_{Enc} for Enc with input 1^λ</p> <p>2: return $K = K_{\text{Enc}}$</p>	<p><i>Client:</i></p> <p>1: return w</p>
<p style="text-align: center;">LayeredSSE.Setup(K, N, DB)</p> <p>1: Set $\tau \leftarrow p \cdot c \log \log \log(\lambda) \log \log(N/p)$</p> <p>2: Sample bins B_1, \dots, B_m via <code>L2C.Setup</code> with input $(\{(w_i, \text{DB}(w_i))\}_{i=1}^W, N/p)$</p> <p>3: Fill B_1, \dots, B_m up to size τ with zeros</p> <p>4: Set $B_i^{\text{enc}} \leftarrow \text{Enc}_{K_{\text{Enc}}}(B_i)$ for $i \in [1, m]$</p> <p>5: return $\text{EDB} = (B_1^{\text{enc}}, \dots, B_m^{\text{enc}})$</p>	<p><i>Server:</i></p> <p>1: Set $\alpha_1, \alpha_2 \leftarrow H(w)$</p> <p>2: return $B_{\alpha_1}^{\text{enc}}, B_{\alpha_2}^{\text{enc}}$</p>
<p style="text-align: center;">LayeredSSE.Search($K, w; \text{EDB}$)</p> <p><i>Client:</i></p> <p>1: return w</p>	<p><i>Client:</i></p> <p>1: Set $B_{\alpha_i} \leftarrow \text{Dec}_{K_{\text{Enc}}}(B_{\alpha_i}^{\text{enc}})$ for $i \in \{1, 2\}$</p> <p>2: Retrieve ball (w, L) from B_α for appropriate $\alpha \in \{\alpha_1, \alpha_2\}$</p> <p>3: Run <code>L2C.UpdateBall</code>$((w, L), L', B_{\alpha_1}, B_{\alpha_2})$</p> <p>4: Set $B_{\alpha_i}^{\text{new}} \leftarrow \text{Enc}_{K_{\text{Enc}}}(B_{\alpha_i})$ for $i \in \{1, 2\}$</p> <p>5: return $B_{\alpha_2}^{\text{new}}, B_{\alpha_2}^{\text{new}}$</p>
<p><i>Server:</i></p> <p>1: Set $\alpha_1, \alpha_2 \leftarrow H(w)$</p> <p>2: return $B_{\alpha_1}^{\text{enc}}, B_{\alpha_2}^{\text{enc}}$</p>	<p><i>Server:</i></p> <p>1: Replace $B_{\alpha_i}^{\text{enc}}$ with $B_{\alpha_i}^{\text{new}}$ for $i \in \{1, 2\}$</p>

5.2 Security and Efficiency

Correctness. LayeredSSE is correct as each keyword has two bins that contain its identifiers associated to it (and these bins are consistently retrieved and updated with L2C). If the hash function is modeled as a random oracle, the bin choices are uniformly random and Theorem 1 guarantees that bins do not overflow.

Selective Security. LayeredSSE is selectively secure and has standard setup leakage N , such as search and update leakage qp , where qp is the query pattern⁴. This can be shown with a simple hybrid argument, sketched here. For setup, the simulator `Sim` receives N , recomputes m and initializes m empty bins B_1, \dots, B_m of size $p \cdot c \log \log \log(\lambda) \log \log(N/p)$ each. `Sim` then outputs $\text{EDB}' = (\text{Enc}_{K'_{\text{Enc}}}(B_i)_{i=1}^m)$ for some sampled key K'_{Enc} . As Enc is IND-CPA secure (and bins do not overflow in the real experiment except with negligible probability), the output EDB' is indistinguishable from the output of Setup in the real

⁴ This is equivalent to page length hiding leakage $\mathcal{L}_{\text{len-hid}}$, as we only restrict ourselves to lists of size at most p .

experiment. For a search query on keyword w , Sim checks the query pattern \mathbf{qp} whether w was already queried. If w was not queried before, Sim a new uniformly random keyword w' . Otherwise, Sim responds with the same keyword w' from the previous query. As we assume that keywords are preprocessed by the client via a PRF, the keywords w and w' are indistinguishable. For an update query on keyword w , the client output in the first flow is the same as in a search query and thus, Sim can proceed as in search. For the second flow, Sim receives two bins $B_{\alpha_1}, B_{\alpha_2}$ from the adversary, directly reencrypts them and sends them back to the adversary. This behaviour is indistinguishable, as the bins are encrypted and again, bins do not overflow except with negligible probability.

Adaptive Security. For adaptive security, the adversary can issue search and update queries that depend on previous queries. As Theorem 1 assumes selectively chosen `InsertBall` and `UpdateBall` operations, there is no guarantee that bins do not overflow anymore in the real game. Thus, the adversary can potentially distinguish update queries of the simulated game from real update queries if she manages to overflow a bin in the real game, as she would receive bins with increased size only in latter case. Fortunately, we can just add a check in `Update` whether one of the bins overflows after the `L2C.UpdateBall` operation. In that case, the client reverts the update and send back the (reencrypted) original bins. Now, Theorem 1 still guarantees that bins overflow only with negligible probability after `Setup` and we can show that the simulated game is indistinguishable from the real game as before. Note that `LayeredSSE` is still correct after this modification, since updates that lead to overflows cannot occur by accident, but only if the client systemically adapts the choice of updates to the random coins used during previous update operations (see Theorem 1).

Note that when the client remarks that a bin overflowed in an `Update` in a real world environment, this is due malicious `Update` operations. The client can adapt his reaction accordingly, whereas the server learns no information about the attack without being notified by the client. We can show that `LayeredSSE` with the adjustment of `Update` is correct and $\mathcal{L}_{\text{len-hid}}$ -adaptively secure. The same simulator Sim suffices and we omit the details.

Efficiency. `LayeredSSE` has constant storage efficiency, as the server stores $m = \left\lceil (N/p) / (\log \log \log \lambda \cdot \log \log \frac{N}{p}) \right\rceil$ bins of $\mathcal{O}\left(p \log \log \log \lambda \cdot \log \log \frac{N}{p}\right)$ identifiers each. There is no client stash required. Each search and update query, the server looks up 2 bins, and thus `LayeredSSE` has $\tilde{\mathcal{O}}(\log \log(N/p))$ page efficiency. Note that `LayeredSSE` has $\mathcal{O}(1)$ locality if only lists up to size p are inserted.

Extensions. With some care, `LayeredSSE` can handle deletes and arbitrary lists (without sacrificing security and efficiency). We refer to Appendix A.3 for more details. The results are formalized in Theorem 2.

Theorem 2 (LayeredSSE). *Let N be an upper bound on the size of database DB and $p \leq N^{1-1/\log \log \lambda}$ be the page size. The scheme `LayeredSSE` is correct and $\mathcal{L}_{\text{len-hid}}$ -adaptively semantically secure if `Enc` is IND-CPA secure and `H` is modeled*

as a random oracle. It has constant storage efficiency and $\tilde{\mathcal{O}}(\log \log N/p)$ page efficiency. If only lists up to size p are inserted, LayeredSSE has constant locality.

6 The Generic Local Transform

In this section, we define the Generic Local Transform (GLT), creating a link between the two IO-efficiency goals of *locality* and *page efficiency*. Namely, the GLT builds an SSE scheme with good *locality* properties from an SSE scheme with good *page efficiency*. For a page-efficient scheme to be used within the GLT, it needs to have certain extra properties. We define such schemes as *suitable* page-efficient schemes in Section 6.1. Next, we introduce the useful notion of *overflowing* SSE. The GLT is then obtained by combining an overflowing SSE with a suitable page-efficient scheme. The OSSE we will use for that purpose, ClipOSSE, is presented in Section 6.2. Finally, the GLT is built from the previous components in Section 6.4. An overview of the correctness and security proofs is provided in section 6.5. Full proofs are available in Appendix E.

6.1 Preliminaries

Suitable page-efficient SSE.

The GLT will create many instances of the underlying page-efficient scheme, each with a different page size. For that reason, for the purpose of the GLT, we slightly extend the standard SSE interface defined in Section 3: namely, $\text{Setup}(K, N, \text{DB}, p)$ takes as an additional parameter the page size p . In addition, recall that, in Section 3, we have allowed the $\text{Update}(K, (w, L), \text{op}, \text{st}; \text{EDB})$ procedure to add a *set* of matching documents K to a given keyword w in a single call. Note that S is allowed to be empty, in which case nothing is added.

If a scheme instantiates that interface, and, in addition, satisfies the following three conditions, we will call such a scheme a *suitable* page-efficient SSE.

- The scheme has client storage $\mathcal{O}(1)$.
- The scheme has locality $\mathcal{O}(1)$ during searches and updates *when accessing a list of length at most one page*.
- The leakage of the scheme is page-length-hiding.

Overflowing SSE. We introduce the notion of *Overflowing* SSE. An Overflowing SSE (OSSE) has the same interface and functionality as a standard SSE scheme, except that during a **Setup** or **Update** operation, it may refuse to store some document identifiers. Those identifiers are called *overflowing*. At the output of the **Setup** and **Update** operations, the client returns the set of overflowing elements. Compared to standard SSE, the correctness definition is relaxed in the following way: during a **Search**, only matching identifiers that were *not* overflowing need to be retrieved.

The intention of an Overflowing SSE is that it may be used as a component within a larger SSE scheme, which will store the overflowing identifiers using

a separate mechanism. The use of an OSSE may be regarded as implicit in some prior SSE constructions. We have chosen to introduce the notion explicitly because it allows to cleanly split the presentation of the Generic Local Transform into two parts: an OSSE scheme that stores most of the database, and an array of page-efficient schemes that store the overflowing identifiers.

6.2 Dynamic Two-Dimensional One-Choice Allocation

The first component of the Generic Local Transform is an OSSE scheme, ClipOSSE. In line with prior work, we split the presentation of ClipOSSE into two parts: an allocation scheme, which specifies where elements should be stored; and the SSE scheme built on top of it, which adds a layer of encryption, key management, and other mechanisms needed to convert the allocation scheme into a full SSE.

The allocation scheme within ClipOSSE is called 1C-Alloc. Similar to [ANSS16], the allocation scheme is an abstract construct that defines the memory locations where items should be stored, but does not store anything itself. In the case of 1C-Alloc, items are stored within buckets, and the procedures return as output the indices of *buckets* where items should be stored. From the point of view of 1C-Alloc, each bucket has unlimited storage. In more detail, 1C-Alloc contains two procedures, `Fetch` and `Add`.

- `Fetch(m, w, ℓ)`: given a number of buckets m , a keyword w , and a list length ℓ , `Fetch` returns (a superset of) the indices of buckets where elements matching keyword w may be stored, assuming there are ℓ such elements.
- `Add(m, w, ℓ)`: given the same input, `Add` returns the index of the bucket where the *next* element matching keyword w should be inserted, assuming there are currently ℓ matching elements.

The intention is that `Add` is used during an SSE `Update` operation, in order to choose the bucket where the next list element is stored; while `Fetch` is used during a `Search` operation, in order to determine the buckets that need to be read to retrieve all list elements. 1C-Alloc will satisfy the correctness property given in Definition 10. Note that the number of buckets m is always assumed to be a power of 2.

Definition 10 (Correctness). *For all m, w, ℓ , if m is a power of 2, then*

$$\bigcup_{0 \leq i \leq \ell-1} \text{Add}(m, w, i) \subseteq \text{Fetch}(m, w, \ell).$$

To describe 1C-Alloc, it is convenient to conceptually group buckets into superbuckets. For $\ell = 2^i \leq m$, an ℓ -*superbucket* is a collection of ℓ consecutive buckets, with indices of the form $k \cdot \ell, k \cdot \ell + 1, \dots, (k + 1) \cdot \ell - 1$, for some $k \leq m/\ell$. A 1-superbucket is the same as a bucket. Notice that for a given ℓ , the ℓ -superbuckets do not overlap. They form a partition of the set of buckets. For $\ell > 1$, each ℓ -superbucket contains exactly two $\ell/2$ -superbuckets.

Let H be a hash function, whose output is assumed to be uniformly random in $\{1, \dots, m\}$. 1C-Alloc works as follows. Fix a keyword w and length $\ell \leq m$

(the case $\ell > m$ will be discussed later). Let $\ell' = 2^{\lceil \log \ell \rceil}$ be the smallest power of 2 larger than ℓ . On input w and ℓ , `1C-Alloc.Fetch` returns the (unique) ℓ' -superbucket that contains $H(w)$.

Algorithm 3 Dynamic Two-Dimensional One-Choice Allocation (1C-Alloc)

<code>1C-Alloc.Fetch</code> (m, w, ℓ)	<code>1C-Alloc.Add</code> (m, w, ℓ)
1: $\ell' \leftarrow 2^{\lceil \log \ell \rceil}$	1: $\ell \leftarrow \ell \bmod m$
2: if $\ell' \geq m$ then	2: $\ell' \leftarrow 2^{\lceil \log(\ell+1) \rceil}$
3: return $\{0, \dots, m-1\}$	3: $i \leftarrow \lfloor H(w)/\ell' \rfloor$
4: else	4: if $\lfloor 2H(w)/\ell' \rfloor \bmod 2 = 0$ then
5: $i \leftarrow \lfloor H(w)/\ell' \rfloor$	5: return $\ell' \cdot i + \ell$
6: return $\{\ell' \cdot i, \dots, \ell' \cdot i + \ell' - 1\}$	6: else
	7: return $\ell' \cdot i + \ell - \ell'/2$

Meanwhile, `1C-Alloc.Add` is designed in order to ensure that the first ℓ successive locations returned by `Add` for keyword w are in fact included within the ℓ' -superbucket above $H(w)$ (that is, in order to ensure correctness). For the first list element (when $\ell = 0$), `Add` returns the bucket $H(w)$; for the second element, it returns the other bucket contained inside the 2-superbucket above $H(w)$. More generally, if S is the smallest superbucket above $H(w)$ that contains at least $\ell+1$ buckets, `Add` returns the leftmost bucket within S that has not yet received an element. In practice, the index of that bucket can be computed easily based on ℓ and the binary decomposition of $H(w)$, as done in Algorithm 3. (In fact, the exact order in which buckets are selected by `Add` is irrelevant, as long as it selects distinct buckets, and correctness holds.)

When the size of the list ℓ grows above the number of buckets m , `Fetch` returns all buckets, while `Add` selects the same buckets as it did for $\ell \bmod m$.

6.3 Clipped One-Choice OSSE

ClipOSSE is the OSSE scheme obtained by storing lists according to `1C-Alloc`, using $m = O(N/\log \log N)$ buckets, with each bucket containing up to $\tau = \lceil \alpha \log \log N \rceil$ items, for some constant α . Buckets are always padded to the threshold τ and encrypted before being stored on the server. Thus, from the server's point of view, they are completely opaque. A table T containing (in encrypted form) the length of the list matching each keyword w is also stored on the server.

Given `1C-Alloc`, the details of ClipOSSE are straightforward. A short overview is given in text below. The encrypted database generated by `Setup` is essentially equivalent to starting from an empty database, and populating it by making repeated calls to `Update`, one for each keyword–document pair in the database. For that reason, we focus on `Search` and `Update`. The full specification for `Setup`, `Search`, and `Update` is given as pseudo-code in Algorithm 4.

Search. To retrieve the list of identifiers matching keyword w , ClipOSSE calls `1C-Alloc`(m, w, ℓ) to get the set of bucket indices where the elements matching keyword w have been stored. The client retrieves those buckets from the server, and decrypts them to obtain the desired information.

Update. For simplicity, we focus on the case where a single identifier is added. The case of a set of identifiers can be obtained by repeating the process for each identifier in the set. To add the new item to the list matching keyword w , ClipOSSE calls $\text{1C-Alloc}(m, w, \ell)$ to determine the bucket where the new list item should be inserted. The client retrieves that bucket from the server, decrypts it, adds the new item, reencrypts the bucket, and sends it back to the server. If that bucket was already full, the item is overflowing, in the sense of Section 6.1.

Algorithm 4 Clipped One-Choice OSSE (ClipOSSE)

Global parameters: constants $d, \alpha \in \mathbb{N}^*$

ClipOSSE.KeyGen (1^λ) <ol style="list-style-type: none"> 1: Generate keys K, K_{PRF} for Enc, PRF 2: return $K = (K, K_{\text{PRF}})$ 	ClipOSSE.Search ($K, w, \text{st}; \text{EDB}$) <p><i>Client:</i> (search token)</p> <ol style="list-style-type: none"> 1: send $(w, K_w) = \text{PRF}_{K_{\text{PRF}}}(w)$ <p><i>Server:</i></p> <ol style="list-style-type: none"> 1: $\ell \leftarrow \text{Dec}_{K_w}(T[w])$ 2: $S \leftarrow \text{1C-Alloc.Fetch}(m, w, \ell)$ 3: return $\{B^{\text{Enc}}[i] : i \in S\}$
ClipOSSE.Setup (K, N, DB) <ol style="list-style-type: none"> 1: $m \leftarrow 2^{\lceil \log(N / \log \log N) \rceil}$ 2: $\tau \leftarrow \lceil \alpha \log \log N \rceil$ 3: $B_0, \dots, B_{m-1}, T, \text{EDB}, \text{clip} \leftarrow \emptyset$ 4: for all each $(w, \{e_1, \dots, e_\ell\})$ in DB do 5: $K_w \leftarrow \text{PRF}_{K_{\text{PRF}}}(w)$ 6: $T[w] \leftarrow \text{Enc}_{K_w}(\ell)$ 7: for all t from 1 to ℓ do 8: $C \leftarrow \emptyset$ 9: $i \leftarrow \text{1C-Alloc.Add}(m, w, t - 1)$ 10: if then $B[i] < \tau$ 11: $B[i] \leftarrow B[i] \cup \{e_i\}$ 12: else 13: $C \leftarrow C \cup \{e_i\}$ 14: if $S > 0$ then 15: $\text{clip} \leftarrow \text{clip} \cup (w, \ell, C)$ 16: Let $B^{\text{Enc}}[i] = \text{Enc}_K(B_i)$ for each i 17: return $\text{EDB} = (T, (B^{\text{Enc}}[i])), \text{clip}$ 	ClipOSSE.Update ($K, (w, \{e\}), \text{op}, \text{st}; \text{EDB}$) <p><i>Client:</i> (update token)</p> <ol style="list-style-type: none"> 1: send $(w, K_w = \text{PRF}_{K_{\text{PRF}}}(w))$ <p><i>Server:</i></p> <ol style="list-style-type: none"> 1: $\ell \leftarrow \text{Dec}_{K_w}(T[w])$ 2: $i \leftarrow \text{1C-Alloc.Add}(m, w, \ell)$ 3: send $B^{\text{Enc}}[i]$ <p><i>Client:</i></p> <ol style="list-style-type: none"> 1: $B \leftarrow \text{Dec}_K(B^{\text{Enc}}[i])$ 2: if $B < \tau$ then 3: $\text{clip} \leftarrow \emptyset$ 4: $B \leftarrow B \cup \{e\}$ 5: else 6: $\text{clip} \leftarrow \{e\}$ 7: send $B' = \text{Enc}_K(B)$ <p><i>Server:</i></p> <ol style="list-style-type: none"> 1: $B^{\text{Enc}}[i] \leftarrow B'$ <p><i>Client:</i></p> <ol style="list-style-type: none"> 2: return clip

6.4 The Generic Local Transform

The Generic Local Transform takes as input a page-length-hiding page-efficient SSE scheme PE-SSE. It outputs a local SSE scheme $\text{Local}[\text{PE-SSE}]$.

To realize $\text{Local}[\text{PE-SSE}]$, we use two structures. The first structure is an instance of ClipOSSE, which stores most of the database. The second structure is an array of n_{level} instances of PE-SSE. The i -th instance, denoted PE-SSE_i , has page size 2^i . The PE-SSE_i instances are used to store elements that overflow

from ClipOSSE. In addition, a table T stores (in encrypted form) the length of the list matching keyword w , for each keyword⁵.

Fix a keyword w , matching ℓ elements. Let $\ell' = 2^{\lceil \log \ell \rceil}$ be the smallest power of 2 larger than ℓ . Let $i = \log \ell'$. At any point in time, the elements matching w are stored in two locations: ClipOSSE, and PE-SSE _{i} . Each of these two locations stores part of the elements: ClipOSSE stores the elements that did not overflow, and PE-SSE _{i} stores the overflowing elements. Each element exists in only one of the two locations. Again, for simplicity, we define updates for adding a single identifier per keyword. The case of adding a set of identifiers at once can be deduced by repeating the same process for each identifier in the set.

Algorithm 5 Generic Local Transform (Local[PE-SSE])

Global parameters: constant $d \in \mathbb{N}^*$	
Local[PE-SSE].KeyGen(1^λ)	Local[PE-SSE].Setup(K, N, DB)
1: Generate key K_{PRF} for PRF 2: return $K = (K, K_{\text{PRF}})$	1: $n_{\text{level}} \leftarrow \lceil N / \log^d N \rceil$ 2: for all $(w, S) \in DB$ do 3: $K_w \leftarrow \text{PRF}_{K_{\text{PRF}}}(w)$ 4: $T[w] \leftarrow \text{Enc}_{K_w}(S)$ 5: EDB, clip $\leftarrow \text{ClipOSSE.Setup}(DB)$ 6: for all i from 0 to n_{level} do 7: $DB_i \leftarrow \{(w, C) : (w, \ell, C) \in \text{clip}$ and $2^{i-1} < \ell \leq 2^i\}$ 8: PE-SSE _{i} $\leftarrow \text{PE-SSE.Setup}(\lceil N / \log N \rceil, 2^i, DB_i)$
Local[PE-SSE].Update($K, (w, L); EDB$)	Local[PE-SSE].Search($K, w, st; EDB$)
Client: (update token) 1: send $(w, L, K_w = \text{PRF}_{K_{\text{PRF}}}(w))$ Server: 1: $C \leftarrow \text{ClipOSSE.Update}(w, L)$ 2: $\ell \leftarrow \text{Dec}_{K_w}(T[w])$ 3: $T[w] \leftarrow \text{Enc}_{K_w}(\ell + 1)$ 4: send ℓ Client: 1: $i \leftarrow \lceil \log \ell \rceil$ 2: if $\lceil \log \ell \rceil = \lceil \log(\ell + 1) \rceil$ then 3: PE-SSE _{i} .Update(w, C) 4: else 5: $S \leftarrow$ set of matches in PE-SSE _{i} .Search(w) 6: PE-SSE _{$i+1$} .Update($w, S \cup C$)	Client: (search token) 1: send $(w, K_w = \text{PRF}_{K_{\text{PRF}}}(w))$ Server: 1: $i \leftarrow \lceil \log(\text{Dec}_{K_w}(T[w])) \rceil$ 2: return $\text{ClipOSSE.Search}(w) \cup \text{PE-SSE}_i.\text{Search}(w)$

Search. During a search operation, Local[PE-SSE] queries both structures, and combines their output to retrieve all matching elements.

Update. During an update operation to add element e , Local[PE-SSE] forwards the update query to ClipOSSE, and gets as output $C = \emptyset$ if the element did not overflow, or $C = \{e\}$ if the element did overflow. For now, assume that $\lceil \log \ell \rceil = \lceil \log(\ell + 1) \rceil$, that is, the PE-SSE _{i} instance associated with the list remains the same during the update operation. In that case, PE-SSE _{i} is updated for the set C . (Recall from Section 6.1 that a length-hiding SSE such as PE-SSE accepts sets of elements as input in Update.) The length-hiding property is designed to guarantee that the content of C (including whether it is empty) is not

⁵ The same table exists in ClipOSSE. In an actual implementation, they would be the same table, but using ClipOSSE in black box eases the presentation.

leaked to the server. Now assume $\lceil \log \ell \rceil < \lceil \log(\ell + 1) \rceil$. In that case, the PE-SSE instance associated with the list becomes PE-SSE $_{i+1}$ instead of PE-SSE $_i$. The client retrieves all current overflowing elements from PE-SSE $_i$, adds the content of C , and stores the result in PE-SSE $_{i+1}$.

6.5 Overflow of ClipOSSE

The main technical result in this section regards the number of overflowing items in ClipOSSE.

Theorem 3. *Suppose that ClipOSSE receives as input a database of size N , such that the size of the longest list is $\mathcal{O}(N/\log^d N)$ for some $d \geq 2$. Then for any constant c , there exists a choice of parameters of ClipOSSE such that the number of overflowing items is $\mathcal{O}(N/\log^c N)$.*

The proof of Theorem 3 is intricate. For space reasons, we only give a brief overview here. The full proof is given in Appendix E. A more detailed overview may also be found in Appendix E.2. First, we show that the result holds in the special case where all lists have length $N/\log^d N$. This uses a negative association argument, similar to the proof of [DPP18, Theorem 1]. The core of the proof is to then show that this special case implies the general case. This is done by iteratively merging short lists, while showing that this merging process can only have a limited effect on the number of overflowing elements. At the outcome of the merging process, all lists have length $N/\log^d N$, which reduces the problem to the special case. The main technique for the reduction is a stochastic dominance argument, combined with a convexity argument (similar to the proof of [BBF⁺21, Theorem 5]).

The Generic Local Transform itself uses standard SSE techniques, and its properties follow from previous discussions. We provide a formal statement below.

Theorem 4 (Generic Local Transform). *Let N be an upper bound on the size of database DB. Suppose that PE-SSE is a suitable page-efficient scheme with page efficiency P and storage efficiency S . Then Local[PE-SSE] is a correct and secure SSE scheme with storage efficiency $\mathcal{O}(S)$, locality $\mathcal{O}(1)$, and read efficiency $P + \tilde{\mathcal{O}}(\log \log N)$.*

7 Unconditional Static Local SSE – Overview

We sketch our SSE scheme UncondSSE with constant locality, constant storage efficiency and $\mathcal{O}(\log^\varepsilon N)$ read efficiency for any $\varepsilon > 0$, without condition on the maximal list size. On a high level, we proceed in 3 steps and refer to Appendix B for the detailed construction.

First, we construct LocORAM, a generalization of the ORAM of [DPP18] with constant locality and $\mathcal{O}(n^\varepsilon)$ bandwidth, given that the block size is “large enough”. The construction is straightforward, though some details are technical.

Then, we use LocORAM to outsource the stash of Tethys (from [BBF⁺21]) via trivial binpacking, obtaining the scheme `OramTethys` that has $\mathcal{O}(\log^\varepsilon \lambda)$ page efficiency, constant storage efficiency and constant client storage. Note that `OramTethys` requires $p = \Omega(\lambda)$ (which stems from the minimal block size requirement of LocORAM). Given Tethys and LocORAM, this is straightforward.

Now, we are ready to describe UncondSSE. For this, we follow the high level idea of [DPP18] to handle lists with different schemes depending on the list size. For some $d \in \mathbb{N}$, we split the interval $[1, N]$ of possible list lengths into four different subintervals:

1. For the subinterval $[1, N^{1-1/\log \log \lambda})$, we can simply store the lists using `Local[LayeredSSE]`. Here, the read efficiency is $\tilde{\mathcal{O}}(\log \log N)$.
2. For the subinterval $[N^{1-1/\log \log \lambda}, N/\log^d N)$, the lengths are small enough for the local transformation and large enough for `OramTethys`. Thus, we store the lists using `Local[OramTethys]` with $\mathcal{O}(\log^\varepsilon N)$ read efficiency.
3. We further split the subinterval $[N/\log^d N, N/\log^\varepsilon N)$ into a constant number of subintervals, such that `OramTethys` has $\mathcal{O}(\log^\varepsilon N)$ read efficiency for each subinterval.
4. For the subinterval $[N/\log^\varepsilon N, N]$, we can read the entire database. Thus, we simply encrypt DB and fetch it from the server each query.

Note that subinterval (2) is the bottleneck in [DPP18], and we can use the Generic Local Transform to remove it. The result is formalized in Theorem 5.

Theorem 5 (UncondSSE). *The static scheme UncondSSE is correct and $\mathcal{L}_{\text{len-rev}}$ -adaptively secure. It has constant client storage, $\mathcal{O}(1)$ locality and $\mathcal{O}(\log^\varepsilon N)$ read efficiency for any $\varepsilon > 0$.*

Acknowledgments

The authors would like to thank Raphael Bost for his helpful comments. This work was supported by the ANR JCJC project SaFED.

References

- ABK⁺20. Azar, Y., Broder, A., Karlin, A., Mitzenmacher, M., and Upfal, E. 2020 ACM Paris Kanellakis Theory and Practice Award. <https://awards.acm.org/kanellakis> (2020).
- ABKU94. Azar, Y., Broder, A.Z., Karlin, A.R., and Upfal, E. Balanced allocations. In: Proceedings of the twenty-sixth annual ACM symposium on theory of computing, pp. 593–602 (1994).
- ANSS16. Asharov, G., Naor, M., Segev, G., and Shahaf, I. Searchable symmetric encryption: optimal locality in linear space via two-dimensional balanced allocations. In: D. Wichs and Y. Mansour (eds.), 48th Annual ACM Symposium on Theory of Computing, pp. 1101–1114. ACM Press, Cambridge, MA, USA (Jun. 18–21, 2016).

- ASS18. Asharov, G., Segev, G., and Shahaf, I. Tight tradeoffs in searchable symmetric encryption. In: H. Shacham and A. Boldyreva (eds.), *Advances in Cryptology – CRYPTO 2018, Part I, Lecture Notes in Computer Science*, vol. 10991, pp. 407–436. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug. 19–23, 2018).
- ASS21. Asharov, G., Segev, G., and Shahaf, I. Tight tradeoffs in searchable symmetric encryption. *Journal of Cryptology*, vol. 34(2):(2021), pp. 1–37.
- BBF⁺21. Bossuat, A., Bost, R., Fouque, P.A., Minaud, B., and Reichle, M. SSE and SSD: Page-efficient searchable symmetric encryption. In: T. Malkin and C. Peikert (eds.), *Advances in Cryptology – CRYPTO 2021, Part III, Lecture Notes in Computer Science*, vol. 12827, pp. 157–184. Springer, Heidelberg, Germany, Virtual Event (Aug. 16–20, 2021).
- BCSV06. Berenbrink, P., Czumaj, A., Steger, A., and Vöcking, B. Balanced allocations: The heavily loaded case. *SIAM Journal on Computing*, vol. 35(6):(2006), pp. 1350–1385.
- BFHM08. Berenbrink, P., Friedetzky, T., Hu, Z., and Martin, R. On weighted balls-into-bins games. *Theoretical Computer Science*, vol. 409(3):(2008), pp. 511–520.
- Bos16. Bost, R. $\Sigma\phi\phi\phi$: Forward secure searchable encryption. In: E.R. Weippl, S. Katzenbeisser, C. Kruegel, A.C. Myers, and S. Halevi (eds.), *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pp. 1143–1154. ACM Press, Vienna, Austria (Oct. 24–28, 2016).
- CGKO06. Curtmola, R., Garay, J.A., Kamara, S., and Ostrovsky, R. Searchable symmetric encryption: improved definitions and efficient constructions. In: A. Juels, R.N. Wright, and S. De Capitani di Vimercati (eds.), *ACM CCS 2006: 13th Conference on Computer and Communications Security*, pp. 79–88. ACM Press, Alexandria, Virginia, USA (Oct. 30 – Nov. 3, 2006).
- CGLS17. Chan, T.H.H., Guo, Y., Lin, W.K., and Shi, E. Oblivious hashing revisited, and applications to asymptotically efficient ORAM and OPRAM. In: T. Takagi and T. Peyrin (eds.), *Advances in Cryptology – ASIACRYPT 2017, Part I, Lecture Notes in Computer Science*, vol. 10624, pp. 660–690. Springer, Heidelberg, Germany, Hong Kong, China (Dec. 3–7, 2017).
- CJJ⁺14. Cash, D., Jaeger, J., Jarecki, S., Jutla, C.S., Krawczyk, H., Rosu, M.C., and Steiner, M. Dynamic searchable encryption in very-large databases: Data structures and implementation. In: *ISOC Network and Distributed System Security Symposium – NDSS 2014*. The Internet Society, San Diego, CA, USA (Feb. 23–26, 2014).
- CT14. Cash, D. and Tessaro, S. The locality of searchable symmetric encryption. In: P.Q. Nguyen and E. Oswald (eds.), *Advances in Cryptology – EUROCRYPT 2014, Lecture Notes in Computer Science*, vol. 8441, pp. 351–368. Springer, Heidelberg, Germany, Copenhagen, Denmark (May 11–15, 2014).
- DP17. Demertzis, I. and Papamanthou, C. Fast searchable encryption with tunable locality. In: *Proceedings of the 2017 ACM International Conference on Management of Data*, pp. 1053–1067 (2017).
- DPP18. Demertzis, I., Papadopoulos, D., and Papamanthou, C. Searchable encryption with optimal locality: Achieving sublogarithmic read efficiency. In: H. Shacham and A. Boldyreva (eds.), *Advances in Cryptology – CRYPTO 2018, Part I, Lecture Notes in Computer Science*, vol. 10991, pp. 371–406. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug. 19–23, 2018).

- DR96. Dubhashi, D.P. and Ranjan, D. Balls and bins: A study in negative dependence. BRICS Report Series, vol. 3(25).
- GM11. Goodrich, M.T. and Mitzenmacher, M. Privacy-preserving access of outsourced data via oblivious ram simulation. In: International Colloquium on Automata, Languages, and Programming, pp. 576–587. Springer (2011).
- GMOT11. Goodrich, M.T., Mitzenmacher, M., Ohrimenko, O., and Tamassia, R. Oblivious ram simulation with efficient worst-case access overhead. In: Proceedings of the 3rd ACM workshop on Cloud computing security workshop, pp. 95–100 (2011).
- Goo11. Goodrich, M.T. Data-oblivious external-memory algorithms for the compaction, selection, and sorting of outsourced data. In: Proceedings of the twenty-third annual ACM symposium on Parallelism in algorithms and architectures, pp. 379–388 (2011).
- JK77. Johnson, N.L. and Kotz, S. Urn models and their application; an approach to modern discrete probability theory. New York, NY (USA) Wiley (1977).
- KMW10. Kirsch, A., Mitzenmacher, M., and Wieder, U. More robust hashing: Cuckoo hashing with a stash. *SIAM Journal on Computing*, vol. 39(4):(2010), pp. 1543–1561.
- MM17. Miers, I. and Mohassel, P. IO-DSSE: Scaling dynamic searchable encryption to millions of indexes by improving locality. In: ISOC Network and Distributed System Security Symposium – NDSS 2017. The Internet Society, San Diego, CA, USA (Feb. 26 – Mar. 1, 2017).
- PPYY19. Patel, S., Persiano, G., Yeo, K., and Yung, M. Mitigating leakage in secure cloud-hosted data structures: Volume-hiding for multi-maps via hashing. In: L. Cavallaro, J. Kinder, X. Wang, and J. Katz (eds.), ACM CCS 2019: 26th Conference on Computer and Communications Security, pp. 79–93. ACM Press (Nov. 11–15, 2019).
- PR04. Pagh, R. and Rodler, F.F. Cuckoo hashing. *Journal of Algorithms*, vol. 51(2):(2004), pp. 122–144.
- PSSZ15. Pinkas, B., Schneider, T., Segev, G., and Zohner, M. Phasing: Private set intersection using permutation-based hashing. In: J. Jung and T. Holz (eds.), USENIX Security 2015: 24th USENIX Security Symposium, pp. 515–530. USENIX Association, Washington, DC, USA (Aug. 12–14, 2015).
- PTW10. Peres, Y., Talwar, K., and Wieder, U. The $(1 + \beta)$ -choice process and weighted balls-into-bins. In: Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms, pp. 1613–1619. SIAM (2010).
- RMS01. Richa, A.W., Mitzenmacher, M., and Sitaraman, R. The power of two random choices: A survey of techniques and results. *Combinatorial Optimization*, vol. 9:(2001), pp. 255–304.
- TW07. Talwar, K. and Wieder, U. Balanced allocations: the weighted case. In: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing, pp. 256–265 (2007).
- TW14. Talwar, K. and Wieder, U. Balanced allocations: A simple proof for the heavily loaded case. In: International Colloquium on Automata, Languages, and Programming, pp. 979–990. Springer (2014).
- Vöc03. Vöcking, B. How asymmetry helps load balancing. *Journal of the ACM (JACM)*, vol. 50(4):(2003), pp. 568–589.

A Dynamic Page Efficient SSE

In this section, we introduce the SSE scheme `LayeredSSE` based on L2C. Essentially, we interpret lists L_i of identifiers matching keyword w_i as balls of a certain weight and use L2C to manage the balls in m bins. Let N be the maximal size of the database, p be the page size and H be a hash function mapping into $\{1, \dots, m\}^2$ for $m = \lceil w_{\max}/(\log \log \log \lambda \cdot \log \log w_{\max}) \rceil$ and $w_{\max} = N/p$. Assume for now that $|L_i| \leq p$, *i.e.* each keyword has at most p associated keywords. Let $p \leq N^{1-1/\log \log \lambda}$. (This is needed for the requirement $m \geq \lambda^{1/\log \log \lambda}$ of L2C, see Theorem 1.) For convenience, we adapt the notation of L2C to such lists as follows⁶:

- `L2C.Setup`($\{(w_i, L_i)\}_{i=1}^W, w_{\max}$): We interpret the pair (w_i, L_i) as a ball with identifier w_i and weight $|L_i|/p \in [0, 1]$, where L_i is a list of (at most p) identifiers matching keyword w_i . The bin choices for (w_i, L_i) are given by $\alpha_1, \alpha_2 \leftarrow H(w_i)$. Run the setup defined in Algorithm 1 given these balls.
- `L2C.InsertBall`($(w, L), B_{\alpha_1}, B_{\alpha_2}$): Insert ball (w, L) into either bin B_{α_1} or bin B_{α_2} as in Algorithm 1.
- `L2C.Update`($(w, L), L', B_{\alpha_1}, B_{\alpha_2}$): Update the weight of ball (w, L) to weight $|L \cup L'|/p$ as in Algorithm 1 and add identifiers L' to list L . One of the bins now contains the ball $(w, L \cup L')$. If the new weight lies in a different subinterval, one bin contains a residual ball (w, L) that we consider to not match w anymore.

A.1 LayeredSSE

Here, we describe the dynamic page efficient symmetric searchable encryption scheme `LayeredSSE` based on L2C. For a concise overview, we assume that $\ell_i \leq p$ and ignore delete operations for now. Also, we present a version of the scheme with an update that requires 2 RTTs. Later, we show how to treat arbitrary list sizes, introduce delete operations and show how to obtain updates in 1 RTT. A detailed description of `LayeredSSE` is given in algorithm 6.

`LayeredSSE.KeyGen`(1^λ). Sample encryption key K_{Enc} for `Enc` with the given security parameter λ . Return the client's master secret key $K = K_{\text{Enc}}$.

`LayeredSSE.Setup`(K, N, DB). Receive as input the client's secret key K , an upperbound N on the number of identifiers and the initial database $\text{DB} = (\text{DB}(w_i))_{i=1}^W$. Recall that $\text{DB}(w_i) = (\text{id}_1, \dots, \text{id}_{\ell_i})$ is a list of ℓ_i document identifiers and that $\sum_{i=1}^W \ell_i \leq N$. interpret $(w_i, \text{DB}(w_i))$ as a ball of weight $\ell_i/p \in [0, 1]$

⁶ As Algorithm 1 is kept purely combinatorial, balls technically have no content. We still need to retrieve lists L given the keyword w in this context. Thus, we say that the pair (w, L) is a ball identified by w and scaled weight $|L|/p$. We assume that we can retrieve the list L given w from the bin that contains ball (w, L) . Clearly, this does not change the behaviour of L2C and we can still apply Theorem 1 on the given variant.

and call `L2C.Setup` with maximal weight N/p and balls $(w_i, \text{DB}(w_i))_{i=1}^W$ as input. The two random choices $(\alpha_{i,1}, \alpha_{i,2}) \leftarrow \text{H}(w_i)$ in `L2C.Setup` are drawn by evaluating `H` on w_i . The result are m bins $(B_i)_{i=1}^m$ filled with the balls such that each bin has load at most $c \log \log \log(\lambda) \log \log(N/p)$ (see Theorem 1). Thus, each bin contains at most $p \cdot c \log \log \log(\lambda) \log \log(N/p)$ identifiers as weights are scaled by a factor p . (The constant $c \in \mathbb{N}$ only depends on N but not the output of `L2C.Setup`.) Next, each bin is filled up to maximal size with dummy items. Finally, encrypt the bins $B_i^{\text{enc}} \leftarrow \text{Enc}_{\text{K}_{\text{Enc}}}(B_i)$ and return $\text{EDB} = (B_i^{\text{enc}})_{i=1}^m$.

`LayeredSSE.Search(K, w; EDB)`. The client receives its secret key K and keyword w . She sends w to the server, and in return receives bins $B_{\alpha_1}^{\text{enc}}, B_{\alpha_2}^{\text{enc}}$, where $(\alpha_1, \alpha_2) \leftarrow \text{H}(w)$.

`LayeredSSE.Update(K, (w, L'), add; EDB)`. The client receives its secret key K , keyword w and a list L' of new identifiers matching w . She sends w to the server and again receives bins $B_{\alpha_1}^{\text{enc}}, B_{\alpha_2}^{\text{enc}}$ in return, where $(\alpha_1, \alpha_2) \leftarrow \text{H}(w)$. Next, the client decrypts $B_{\alpha_1}^{\text{enc}}, B_{\alpha_2}^{\text{enc}}$ to $B_{\alpha_1}, B_{\alpha_2}$ and retrieves ball (w, L) from the corresponding bin $B_\alpha \in \{B_{\alpha_1}, B_{\alpha_2}\}$. Then, she calls `L2C.UpdateBall` with old ball (w, L) , new identifiers L' and bins $B_{\alpha_1}, B_{\alpha_2}$ to insert the new identifiers L' into one of the bins. Finally, she reencrypts the bins and sends them to the server. The server then replaces the old bins with the updated bins.

Algorithm 6 LayeredSSE

Global parameters: constant $c \in \mathbb{N}$, page size p	
LayeredSSE.KeyGen (1^λ)	LayeredSSE.Update ($\text{K}, (w, L'), \text{add}; \text{EDB}$)
1: Sample K_{Enc} for <code>Enc</code> with input 1^λ 2: return $K = \text{K}_{\text{Enc}}$	<i>Client:</i> 1: return w
LayeredSSE.Setup (K, N, DB)	<i>Server:</i>
1: Set $\tau \leftarrow p \cdot c \log \log \log(\lambda) \log \log(N/p)$ 2: Sample bins B_1, \dots, B_m via <code>L2C.Setup</code> with input $(\{(w_i, \text{DB}(w_i))_{i=1}^W, N/p\})$ 3: Fill B_1, \dots, B_m up to size τ with zeros 4: Set $B_i^{\text{enc}} \leftarrow \text{Enc}_{\text{K}_{\text{Enc}}}(B_i)$ for $i \in [1, m]$ 5: return $\text{EDB} = (B_1^{\text{enc}}, \dots, B_m^{\text{enc}})$	1: Set $\alpha_1, \alpha_2 \leftarrow \text{H}(w)$ 2: return $B_{\alpha_1}^{\text{enc}}, B_{\alpha_2}^{\text{enc}}$
LayeredSSE.Search ($\text{K}, w; \text{EDB}$)	<i>Client:</i>
<i>Client:</i> 1: return w	<i>Client:</i> 1: Set $B_{\alpha_i} \leftarrow \text{Dec}_{\text{K}_{\text{Enc}}}(B_{\alpha_i}^{\text{enc}})$ for $i \in \{1, 2\}$ 2: Retrieve ball (w, L) from B_α for appropriate $\alpha \in \{\alpha_1, \alpha_2\}$ 3: Run <code>L2C.UpdateBall</code> ($(w, L), L', B_{\alpha_1}, B_{\alpha_2}$) 4: Set $B_{\alpha_i}^{\text{new}} \leftarrow \text{Enc}_{\text{K}_{\text{Enc}}}(B_{\alpha_i})$ for $i \in \{1, 2\}$ 5: return $B_{\alpha_1}^{\text{new}}, B_{\alpha_2}^{\text{new}}$
<i>Server:</i>	<i>Server:</i>
1: Set $\alpha_1, \alpha_2 \leftarrow \text{H}(w)$ 2: return $B_{\alpha_1}^{\text{enc}}, B_{\alpha_2}^{\text{enc}}$	1: Replace $B_{\alpha_i}^{\text{enc}}$ with $B_{\alpha_i}^{\text{new}}$ for $i \in \{1, 2\}$

A.2 Security

The scheme `LayeredSSE` is correct as each keyword has two bins that contain its identifiers associated to it (and these bins are consistently retrieved and updated with `L2C`). If the hash function is modeled as a random oracle, the bin choices are uniformly random and Theorem 1 guarantees that bins do not overflow.

Also, `LayeredSSE` is selectively secure and has standard setup leakage N , such as search and update leakage \mathbf{qp} , where \mathbf{qp} is the query pattern⁷. This can be shown with a simple hybrid argument. We sketch the proof here and refer to Appendix D for more details. For setup, the simulator `Sim` receives N , recomputes m and initializes m empty bins B_1, \dots, B_m of size $p \cdot c \log \log \log(\lambda) \log \log(N/p)$ each. `Sim` then outputs $\text{EDB}' = (\text{Enc}_{K'_{\text{Enc}}}(B_i)_{i=1}^m)$ for some sampled key K'_{Enc} . As `Enc` is IND-CPA secure (and bins do not overflow in the real experiment except with negligible probability), the output EDB' is indistinguishable from the output of `Setup` in the real experiment. For a search query on keyword w , `Sim` checks the query pattern \mathbf{qp} whether w was already queried. If w was not queried before, `Sim` a new uniformly random keyword w' . Otherwise, `Sim` responds with the same keyword w' from the previous query. As we assume that keywords are preprocessed by the client via a PRF, the keywords w and w' are indistinguishable. For an update query on keyword w , the client output in the first flow is the same as in a search query and thus, `Sim` can proceed as in search. For the second flow, `Sim` receives two bins $B_{\alpha_1}, B_{\alpha_2}$ from the adversary, directly reencrypts them and sends them back to the adversary. This behaviour is indistinguishable, as the bins are encrypted and again, bins do not overflow except with negligible probability.

For adaptive security, the adversary can issue search and update queries that depend on previous queries. As Theorem 1 assumes selectively chosen `InsertBall` and `UpdateBall` operations, there is no guarantee that bins do not overflow anymore in the real game. Thus, the adversary can potentially distinguish update queries of the simulated game from real update queries if she manages to overflow a bin in the real game, as she would receive bins with increased size only in latter case. Fortunately, we can just add a check in `Update` whether one of the bins overflows after the `L2C.UpdateBall` operation. In that case, the client reverts the update and send back the (reencrypted) original bins. Now, Theorem 1 still guarantees that bins overflow only with negligible probability after `Setup` and we can show that the simulated game is indistinguishable from the real game as before. (Note that `LayeredSSE` is still correct after this modification, since queries are chosen selectively for correctness.) Note that when the client remarks that a bin overflowed in an `Update` in a real world environment, this is due malicious `Update` operations. The client can adapt his reaction accordingly, whereas the server learns no information about the attack without being notified by the client.

We can show that `LayeredSSE` with the adjustment of `Update` is correct $\mathcal{L}_{\text{len-hid}}$ -adaptively secure. The same simulator `Sim` suffices and we omit the details.

A.3 Extensions

Handling Long Lists. We now adapt `LayeredSSE` to handle arbitrary lists L (with potentially more than p identifiers). We proceed similarly to the static

⁷ This is equivalent to page length hiding leakage $\mathcal{L}_{\text{len-hid}}$, as we only restrict ourselves to lists of size at most p .

scheme Pluto from [BBF⁺21] and extend the ideas to updates. For this, we split L into sublists of size at most p . The (encrypted) full sublists of size p can be stored in a hash table T_{full} on the server and the incomplete sublists are handled by LayeredSSE as before. For search, the client needs to know the number of sublists in order to fetch the right amount from the server. This information is also required for update queries in order to know when to insert another full list into T_{full} . This information can be outsourced in a table T_{len} . Here, the client stores for each keyword w (with ℓ matching identifiers) the number of sublists $T_{\text{len}}[w] = \lceil \ell/p \rceil$ in encrypted format. In the following, we describe the updated Setup, Search and Update of LayeredSSE in more detail.

Setup. For setup, let L_i be a list of ℓ_i identifiers matching keyword w_i and PRF be a secure pseudo-random function mapping to $\{0, 1\}^{\lceil \log(N) \rceil}$. We set $x_i = \lceil \ell_i/p \rceil$. The client splits L_i into sublists $L_{i,1}, \dots, L_{i,x_i-1}$ of size p and sublist L_{i,x_i} of size at most p . She evaluates $m_i \leftarrow \text{PRF}_{\text{K}_{\text{PRF}}}(w_i)$, where K_{PRF} is a key for PRF sampled in KeyGen. The mask m_i is used to encrypt the content of T_{len} . After initializing the table T_{len} with N random entries of size $\log(N)$ bits and T_{full} with N/p (arbitrary) lists of size p , she sets $T_{\text{len}}[w_i] = x_i \oplus m_i$ and $T_{\text{full}}[w \parallel i] = L_{i,j}$ for $j \in [1, x-1]$. Next, she generates $(B_i)_{i=1}^{m_i}$ as before with the incomplete lists L_{i,x_i} except that the bin choices for list L_{i,x_i} are $(\alpha_{i,1}, \alpha_{i,2}) \leftarrow \text{H}(w_i \parallel x_i)$. Note that we need to also hash the counter x_i , as after some updates, the incomplete sublist of w_i might become full and a new incomplete sublist has to be started. When the new incomplete sublist gets inserted with L2C, it is interpreted as a new ball and new bins need to be chosen. Finally, she encrypts the content of T_{full} and returns $\text{EDB} = (T_{\text{len}}, T_{\text{full}}, (B_i^{\text{enc}})_{i=1}^{m_i})$.

Search. For search queries on keyword w , the client outputs mask $m \leftarrow \text{PRF}_{\text{K}_{\text{PRF}}}(w)$ in addition to w . The server uses this mask to decrypt the number of sublists $x \leftarrow T_{\text{len}}[w] \oplus m$, retrieves $x-1$ encrypted sublists $L_i \leftarrow T_{\text{full}}[w \parallel i]$ from the table for $i \in [1, x-1]$ and the two bins $B_{\alpha_1}^{\text{enc}}$ and $B_{\alpha_2}^{\text{enc}}$ via $(\alpha_1, \alpha_2) \leftarrow \text{H}(w \parallel x)$. Finally, the server sends the encrypted bins and sublists to the client. Clearly, the client obtains all matching identifiers after decrypting the received lists and bins.

Update. For update queries on keyword w and list L' of (at most p) new identifiers⁸, the client generates mask m as before and sends (w, m) to the server. The server again decrypts x from T_{len} and sends $B_{\alpha_1}^{\text{enc}}, B_{\alpha_2}^{\text{enc}}$ to the client. In addition, the server already sends the bins $B_{\alpha_3}^{\text{enc}}, B_{\alpha_4}^{\text{enc}}$ for $\alpha_3, \alpha_4 \leftarrow \text{H}(w \parallel x+1)$ to the client (in case the incomplete list overflows). The client now retrieves the old (incomplete) list L of identifiers matching w from the decrypted bins $B_{\alpha_1}, B_{\alpha_2}$. We distinguish two cases:

1. If $L \cup L'$ contains more than p identifiers, the client sets $L^{\text{new}} = L \cup L'$ and marks (w, L) as a residual ball inside $B_{\alpha_1}, B_{\alpha_2}$. Then, she splits L^{new} into two sublists L^{new} with p identifiers and $L^{\leq p}$ of at most p identifiers. The client then inserts list $L^{\leq p}$ into bins $B_{\alpha_3}, B_{\alpha_4}$ via $\text{L2C.InsertBall}((w, L^{\leq p}), B_{\alpha_3}, B_{\alpha_4})$

⁸ For updates with more than p identifiers, the client can perform multiple updates.

and sends the updated (reencrypted) bins $\{B_i^{\text{enc}}\}_{i=1}^4$ such as encrypted list $L^{\text{enc}} = \text{Enc}_{\mathcal{K}_{\text{Enc}}}(L^{\neq p})$ to the server.

2. Otherwise, the client proceeds as before, *i.e.* adds the new identifiers L' to $\text{ball}(w, L)$ via `UpdateBall` and reencrypts the received bins.

Finally, the server replaces the old bins with the reencrypted bins, and if she received an encrypted list L^{new} , she stores the received list in $T_{\text{full}}[w \parallel x + 1] = L^{\text{enc}}$ and updates $T_{\text{len}}[w] = x + 1$.

Leakage profile. Now, search and update queries `LayeredSSE` leak the number of sublists $x = \lceil \ell/p \rceil$ for a given keyword w with ℓ matching identifiers. Further, updates leak when a list was completed (so also the value $\lceil (\ell + |L'|)/p \rceil$). This is exactly the leakage modeled by $\mathcal{L}_{\text{len-hid}}$. As tables T_{len} and T_{full} are encrypted, it is straightforward to adapt the security analysis in Appendix A.2 to the extended scheme with respect to leakage function $\mathcal{L}_{\text{len-hid}}$.

Handling Deletes. We apply the generic solution from [Bos16] to handle deletes. We use two instantiations of `LayeredSSE`, Σ_{add} for added items and Σ_{del} one for deletes. For adding identifiers L' to a keyword w , the client adds list L' to Σ_{add} via $\Sigma_{\text{add}}.\text{Update}(\mathcal{K}, (w, L'), \text{add}; \text{EDB})$. For deleting identifiers L' from a keyword w , the client adds list L' to Σ_{del} via $\Sigma_{\text{del}}.\text{Update}(\mathcal{K}, (w, L'), \text{add}; \text{EDB})$. For a search query, the client fetches the identifiers w from both Σ_{add} and Σ_{del} and removes the set of items L_{del} received from Σ_{del} from the set of items L_{add} received from Σ_{add} , *i.e.* sets $L \leftarrow L_{\text{add}} \setminus L_{\text{del}}$.

Optimized RTT. Search queries of `LayeredSSE` need only 1 RTT, whereas update queries unfortunately require 2 RTTs. We can use “piggybacking” in order to reduce the update RTT to 1 as follows. Instead of sending the second flow of the update query directly to the server, the client stashes the response and waits for the next query (either update or search). On the next query, the client sends the stashed response in addition to the query. The server then finishes the pending update query (by storing the received bins and updating the tables) and responds the query subsequently.

A.4 Efficiency

We now inspect the efficiency of `LayeredSSE`. Let $w_{\text{max}} = N/p$. The server stores $m = \lceil w_{\text{max}}/(\log \log \log \lambda \cdot \log \log w_{\text{max}}) \rceil$ bins of size $\mathcal{O}(p \log \log \log(\lambda) \log \log(w_{\text{max}})) \cdot \mathcal{O}(\lambda)$ each, tables T_{len} with N entries of size $\log(N)$ and T_{full} with N/p entries of size $p \cdot \mathcal{O}(\lambda)$ each. (Recall that a single identifier has size $\mathcal{O}(\lambda)$.) As $N = \text{poly}(\lambda)$, the storage efficiency is $\mathcal{O}(1)$ in total. There is no client stash required⁹. Further, the server looks up 4 bins of capacity $\tilde{\mathcal{O}}(p \log \log(N/p))$ and $x - 1$ encrypted lists

⁹ The version of `LayeredSSE` with 1 RTT updates requires a stash of size $\tilde{\mathcal{O}}(p \log \log(N/p))$ to temporarily store the second flow of the update query until the next query.

of p identifiers from T_{full} for a search query on word w , where x is the number of pages needed to store the document indices matching keyword w in plaintext. Thus, the page efficiency is $\tilde{\mathcal{O}}\left(\log \log \frac{N}{p}\right)$. This further implies that LayeredSSE has $\mathcal{O}(1)$ locality if only lists up to size p are inserted.

Theorem 6 (LayeredSSE). *Let N be an upper bound on the size of database DB and p be the page size. Let $p \leq N^{1-1/\log \log \lambda}$. The scheme LayeredSSE is correct and $\mathcal{L}_{\text{len-hid}}$ -adaptively semantically secure if Enc is IND-CPA secure and H is modeled as a random oracle. It has constant storage efficiency and $\tilde{\mathcal{O}}(\log \log N/p)$ page efficiency. If only lists up to size p are inserted, LayeredSSE has constant locality.*

Proof. Efficiency and security follow from the discussions above. □

B Unconditional Static Local SSE

In this section, we present our unconditional SSE scheme UncondSSE with constant locality, constant storage efficiency and $\mathcal{O}(\log^\varepsilon N)$ read efficiency for any $\varepsilon > 0$. For this, we first present a local ORAM construction in Appendix B.1. Then we construct a static SSE scheme with $\mathcal{O}(\log^\varepsilon \lambda)$ page efficiency in Appendix B.2 that works for large page sizes and has constant client storage. Finally, we use those schemes in order to construct UncondSSE in Appendix B.3

B.1 Local ORAM

Let $c \in \mathbb{N}$ be arbitrary. We now construct an ORAM LocORAM with amortized constant locality and $\mathcal{O}(\beta \cdot n^{1/c} \log^2(n))$ bandwidth, where n is the size of the memory array and $\beta = \Omega(n^{(c-1)/c})$ is the block size. As c is arbitrary, we can instantiate LocORAM with bandwidth overhead $\mathcal{O}(n^\varepsilon)$ for any $\varepsilon > 0$ and constant locality, if the block size is sufficiently large. Our scheme follows the blueprint of the scheme of [DPP18]. The reader may find it helpful to refer to their scheme first.

More Preliminaries. Before detailing the construction, we introduce some additional preliminaries.

Definition 11 (ORAM).

- $\text{Initialize}_{\text{ORAM}}(1^\lambda, M)$: Client take as input security parameter λ such as memory array M of n values $\{(i, v_i)\}_{i=1}^n$ of $\mathcal{O}(\lambda)$ bits each. Outputs client state st and encrypted memory EM .
- $\text{Access}_{\text{ORAM}}(\text{st}, i; \text{EM})$: The client takes as input its state and index i . The server takes as input encrypted memory EM . Outputs value v_i assigned to i and updated state st to the client such as updated encrypted memory EM' to the server.

We require read-only ORAM with zero-failure probability for our construction. We say that an ORAM scheme is correct, if for any access sequence on block i , the retrieved block via $\text{Access}_{\text{ORAM}}$ is (i, v_i) . We say that an ORAM scheme is (adaptively) secure, if for any two (adaptively chosen) access sequences S_1 and S_2 of the same length, their access patterns $A(S_1)$ and $A(S_2)$ are computationally indistinguishable by anyone but the client. We refer to [DPP18] for formal definitions.

Lemma 1 (Local Oblivious Sort [Goo11, GM11, DPP18]). *Given an array X containing n comparable elements, we can sort X with a data-oblivious external-memory protocol Obsort that uses $\mathcal{O}\left(\frac{n}{b} \log^2 \frac{n}{b}\right)$ I/O operations and local memory of $4b$ chunks, where an I/O operation is defined as the read/write of b consecutive chunks of X .*

We set chunk size $b = n^{1/c} \log^2 n$ in our ORAM scheme. This suffices for $\mathcal{O}(1)$ locality. We implicitly assume that the sorted array is reencrypted (under the same encryption key as the input array).

The Scheme. We now describe our construction LocORAM of a read-only ORAM (based on [DPP18]). Let c be a constant. We write β for the ORAM blocksize. Essentially, LocORAM is a hierarchical ORAM with c levels. For n blocks of memory, it has constant locality and a bandwidth of $\mathcal{O}(\beta \cdot n^{1/c} \log^2(n))$ with $\mathcal{O}(\beta \cdot n^{1/c} \log^2(n))$ temporary client storage, if $\beta = \Omega(n^{(c-1)/c})$. Now, we give an overview of the construction. A detailed description is given in Algorithm 7.

$\text{LocORAM.Initialize}_{\text{ORAM}}(1^\lambda, M)$. The client receives memory $M = \{k, v_k\}_{k=1}^n$ with blocksize $|k, v_k| = \beta$. She allocates c arrays A_1, A_2, \dots, A_c with space for n_i blocks each, where $n_1 = n^{1/c}$ and $n_i = n^{i/c} + n^{(i-1)/c}$. Let $\pi_i : [1, n_i] \mapsto [1, n_i]$ be pseudorandom permutations. Initially, the client stores all blocks (k, v_k) in A_c at position $\pi_c(k)$. Later on, blocks will also be stored in other levels. Note that while A_c can hold all blocks, lower levels A_i can only store up to $n^{i/c}$ blocks (and the remaining space is reserved for dummy queries). We would still like to store block (k, v_k) at a pseudorandom position. For this, we initialize tables T_i , $i \in [2, c-1]$, which store a scaled index $T_i[k] \in [1, n^{i/c}]$ for block k (if the block is stored in level i). (Note that we do not require table T_1 for the first level, as the client always retrieves the entire array A_1 each read and thus, no pseudorandom accesses is required for the first level.) The block k is later stored at location $\pi_i(T_i[k])$. Further, the client initializes sets R_i that store the blocks mapped to level A_i . Initially, $R_c = M$ and the other sets are empty. Levels are later rebuilt (with a new pseudorandom permutation) after a certain number of reads such that each item is only accessed once per level before the next rebuild (and thus resemble a random access to A_i). The client keeps track of the number of reads cnt_i at level i after the last rebuild, initialized to 0. The client finally encrypts A_i, R_i and T_i and sends them to the server and stores the encryption key, cnt_i , and π_i in its state.

LocORAM.Access_{ORAM}(st, k; EM). The client receives index k of the block to be retrieved and its state st which she parses as $\{\pi_i\}_{i=2}^c, \{\text{cnt}_i\}_{i=2}^c, \text{K}_{\text{Enc}}$. The server receives the encrypted memory EM which she parses as $(\{A_i^{\text{enc}}\}_{i=1}^c, \{R_i^{\text{enc}}\}_{i=2}^{c-1}, \{\text{T}_i^{\text{enc}}\}_{i=2}^{c-1})$. First, the client increments the counts cnt_i . Next, the client retrieves tables T_i^{enc} and array A_1^{enc} from the server¹⁰. After decryption, the client looks for the first level A_{i^*} in which (k, v_k) is stored. This level is the first for which the table T_i has a non-zero entry $\text{T}_i[k] \neq 0$. She performs a dummy query for all other levels A_i (by accessing a random and unqueried position in A_i using π_i) and retrieves (k, v_k) from level A_{i^*} (either by scanning A_1 if $i^* = 1$ or from A_{i^*} at position $\pi_{i^*}(\text{T}_{i^*}[k])$) with the help of the server. Next, the client writes (k, v_k) to A_1 . Note that later, A_1 will be merged with upper levels and thus, we already prepare $\text{T}_i[k] = \text{cnt}_{i+1}$ and add an encryption of block (k, v_k) to the set R_i of blocks stored in the i -th level. Last, the client rebuilds (some of) the levels if necessary. For this, she takes the highest i^* such that $\text{cnt}_{i^*} > n^{(i^*-1)/c}$. If $i^* \geq 2$, all levels A_i below A_{i^*} are emptied and filled with dummy blocks. Then, she chooses new pseudorandom permutations π_i for $i \in [2, i^*]$ and merges all blocks from the lower levels into A_{i^*} via a local oblivious sort **ObISort** (see Lemma 1). Concretely, the client and server interactively sort R_{i^*} with respect to the new π_{i^*} . The array R_{i^*} is temporarily filled up to n_{i^*} blocks with zeros in order to obtain array A_{i^*} of size n_{i^*} . For the oblivious sort, we choose a chunk size of $n^{1/c} \log^2 n$ (which is sufficient for $\mathcal{O}(1)$ locality). Lastly, the server and client empty the lower levels (and its auxiliary data structures) A_i, R_i, T_i for $i \in [2, i^* - 1]$ and A_1 . (Note that an empty and encrypted array A_i can be constructed iteratively with the same method as R_i .) Finally, the client updates its state, reencrypts the received data structures, and sends them back to the server (who updates EM accordingly).

Theorem 7 (LocORAM). *Let n be the size of the memory array, $\beta = \Omega(n^{\frac{c-1}{c}})$ the blocksize. The scheme LocORAM is correct and secure, if the π_i 's are secure pseudorandom permutations and Enc is an IND-CPA secure encryption scheme. Further, LocORAM has amortized constant locality and bandwidth of $\mathcal{O}(\beta \cdot n^{1/c} \log^2(n))$, and requires $\mathcal{O}(\beta \cdot n^{1/c} \log^2(n))$ temporary client storage (during a rebuild). The client state st has constant size.*

Proof. We show that scheme LocORAM is correct (1), secure (2) and analyze its efficiency (3).

(1) We need to show that for all indices $k \in [1, n]$, the server retrieves the block (k, v_k) via the access protocol when executing an adaptive access sequence after the initialization. For this, we observe that (k, v_k) is either in A_1 , $A_i[\pi_i(\text{T}_i[k])]$ (for the minimal $i \in [2, c-1]$ such that $\text{T}_i[k] \neq 0$) or $A_c[\pi_c(c)]$ (if no such i exists). This is because if (k, v_k) has been accessed in the last $n^{1/c}$ operations, the block will be stored in A_1 . If (k, v_k) has been accessed in the last $n^{i/c}$ operations but not in the last $n^{(i-1)/c}$ operations, it was shuffled into array A_i at position $\text{T}_i[k]$ during a previous rebuild of level A_i . Also, as it has not

¹⁰ Downloading all encrypted T_i incurs a bandwidth of $\mathcal{O}(n^{(c-1)/c} \cdot \log(n))$. As we require $\beta = \Omega(n^{(c-1)/c})$ this cost vanishes in the total access bandwidth.

been accessed recently, we have $T_j[k] = 0$ for $j < i$ as tables below are emptied during a rebuild. Otherwise, it was never accessed before and is located in A_c at initial position $\pi_c(k)$. These values are retrieved by the client and thus, the scheme is correct.

(2) We give an simulator Sim . For initialization, the simulator receives $|M|$, the block size β and the security parameter λ . Sim outputs $\text{LocORAM.Initialize}_{\text{ORAM}}(1^\lambda, M')$ for $M' = \{i, 0\}_{i=1}^n$ where the zeros are of size β . Under the IND-CPA security, this output is indistinguishable from the real game, as the output is encrypted. For simulating an access, Sim retrieves A_1 and the tables T_i from the server. First, Sim increments cnt_i . Then, Sim outputs $r_i \leftarrow \text{Enc}_{\kappa_{\text{Enc}}}(0)$ and random indices index_i that were not yet queried in A_i (since A_i was last emptied). Lastly, Sim and the server rebuild the largest level A_i if $\text{cnt}_i \geq n^{(i-1)/c}$ for some $i \geq 2$. For this, Sim simply checks (and updates) cnt_i accordingly and simulates the oblivious search with the server. It follows by inspection that if π_i are pseudorandom permutations and Enc is IND-CPA secure, the interaction with Sim is indistinguishable from the game.

(3) As c is a constant, the client state is $\mathcal{O}(1)$. Further, the client requires $\mathcal{O}(n^{1/c} \log^2 n)$ blocks of temporary storage for the oblivious sort (see Lemma 1). We now inspect the bandwidth and locality. Over the course of n accesses, the following holds for the i -th access.

- The client reads array A_1 of size $\mathcal{O}(\beta \cdot n^{1/c})$, one block from each other level A_i and tables T_i of size $\mathcal{O}(n^{(c-1)/c} \cdot \log(n)) = \mathcal{O}(\beta \cdot \log n)$, as $\beta = \Omega(n^{(c-1)/c})$. In total, this incurs $\mathcal{O}(\beta \cdot n^{1/c})$ bandwidth and $\mathcal{O}(1)$ locality.
- If $i \bmod n^{(i-1)/c} = 0 \wedge i \geq 2$, the client performs a rebuild of array A_i via an oblivious sort with chunk size $n^{1/c} \log^2 n$. According to Lemma 1 and as A_i contains $\mathcal{O}(n^{i/c})$ blocks, the sort requires $\mathcal{O}(n^{(i-1)/c})$ I/O operations. In total, this incurs $\mathcal{O}(\beta \cdot n^{i/c} \log^2 n)$ bandwidth and $\mathcal{O}(n^{(i-1)/c})$ locality and happens $n^{(c-i+1)/c}$ times during n accesses.

In total, the amortized bandwidth B and locality L are

$$L = \frac{n\mathcal{O}(\beta \cdot n^{1/c}) + \sum_{i=2}^c n^{(c-i+1)/c} \mathcal{O}(\beta \cdot n^{i/c} \log^2 n)}{n} = \mathcal{O}(\beta \cdot n^{1/c} \log^2 n),$$

$$B = \frac{n\mathcal{O}(1) + \sum_{i=2}^c n^{(c-i+1)/c} \mathcal{O}(n^{(i-1)/c})}{n} = \mathcal{O}(1). \quad \square$$

Remark 1 (On Deamortization). Generally, hierarchical ORAMs can be deamortized by continuously reshuffling the layers each operation [GMOT11]. Indeed, our ORAM is an iterated version of [DPP18] which uses this technique for their deamortization. We believe that LocORAM can be deamortized in the same manner but leave the details for future work.

B.2 Tethys without Stash.

Now, we introduce a page-length-hiding static SSE scheme OramTethys that has $\mathcal{O}(\log^\varepsilon \lambda)$ page efficiency, constant storage efficiency and constant client storage. We will later use it in the local transformation.

More Preliminaries. Again, we require some additional preliminaries.

Definition 12 (Binpacking). We define the algorithm `Binpack` that takes at most N keyword-identifier pairs `Stash` and a size p as input. `Binpack` proceeds as follows. Allocate bins $B_1, \dots, B_{2N/p}$ and table T_{pos} . Then, take list L of identifiers matching keyword w and insert the identifiers one-by-one into the bin with the smallest index that is not full yet. Set $T_{\text{pos}}[w] = i$, where i is the smallest index of a bin containing an identifier matching w . Finally, fill the bins up to size p with zeros. Finally, output (M, T_{pos}) .

Clearly, if there are at most p identifiers matching keyword w , the identifiers will all fit into B_i, B_{i+1} for $i \leftarrow T_{\text{pos}}[w]$. Also, note that `Binpack` can always fit all N identifiers into the bins.

Lemma 2 (Tethys [BBF⁺21]). The SSE scheme `Tethys` is correct and $\mathcal{L}_{\text{len-hid}}$ -adaptively secure in the random oracle model (under the assumption that there exists an IND-CPA secure encryption scheme and a secure pseudo-random function). It has a client storage $\omega(\log \lambda) / \log N$ pages, and $\mathcal{O}(1)$ storage and page efficiency.

We call the client storage of `Tethys` its stash. In this work, a stash size of $\mathcal{O}(\log^{1+\delta}(\lambda)) = \omega(\log \lambda) / \log N$ pages is sufficient for our construction, for some arbitrary $\delta > 0$.

The Scheme. We now define the static SSE scheme `OramTethys` with client storage $\mathcal{O}(1)$. Let $p = \Omega(\lambda), c \in \mathbb{N}$ and $\delta > 1$. Essentially, we use `Tethys` (see Lemma 2) and outsource its stash using `LocORAM`. We define the static SSE scheme `OramTethys` for given page size p as follows:

`OramTethys.KeyGen`(1^λ). Simply output $K \leftarrow \text{Tethys.KeyGen}(1^\lambda)$.

`OramTethys.Setup`(K, N, DB). The client generates encrypted database EDB' and stash `Stash` using $(\text{EDB}', \text{Stash}) \leftarrow \text{Tethys.Setup}(K, N, \text{DB})$. The stash contains the remaining keyword-identifier pairs (w_i, id_i) that could not be allocated directly in `Tethys`. We want to outsource `Stash` using `LocORAM` with c levels and blocksize $\beta = p$. The items in `Stash` are not necessarily lists of size p . Thus, we group `Stash` into pages of p identifiers using $(M, T_{\text{pos}}) \leftarrow \text{Binpack}(\text{Stash}, p)$. Let $n = \mathcal{O}(\log^{1+\delta}(\lambda))$. After the binpacking, M consists of n pages with p identifiers each (see Lemma 2) and we can access the identifiers matching keyword w in page $i = T_{\text{pos}}[w]$, as there is at most one such page¹¹. As this binpacking process is not data-oblivious, we can not leak i to the server. Thus, the

¹¹ Without loss of generality, we can assume that there are at most p identifiers per keyword in the stash. For this, we can keep full lists inside a table T_{full} such as in `LayeredSSE`. This version of `Tethys` was already described in [BBF⁺21] (see their scheme `Pluto`). Note further that the binpacking algorithm `Binpack` packs a list of identifiers into at most two consecutive bins. Thus, knowledge of i suffices to fetch bin i and $i + 1$. We assume in the following that the list of identifiers is in at most one bin for simplicity.

client sets $T_{\text{pos}}^{\text{enc}} \leftarrow \text{Enc}_{K_{\text{Enc}}}(T_{\text{pos}})$ (after padding T_{pos} to n entries of $\log n$ bits). Further, she applies $(\text{EM}, \text{st}) \leftarrow \text{LocORAM.Initialize}_{\text{ORAM}}(1^\lambda, M)$. Finally, she outputs $\text{EDB} = (\text{EDB}', \text{EM}, T_{\text{pos}}^{\text{enc}})$ and stores state st locally.

$\text{OramTethys.Search}(K, w; \text{EDB})$. The client initiates protocol $\text{Tethys.Search}(K, w; \text{EDB})$ with the server from which she receives some of the identifiers matching keyword w . Next, the client retrieves $i^{\text{enc}} = T_{\text{pos}}^{\text{enc}}[w]$ from the server and decrypts the index of the page containing the remaining identifiers via $i \leftarrow \text{Dec}_{K_{\text{Enc}}}(i^{\text{enc}})$. The client retrieves this page via $\text{LocORAM.Access}_{\text{ORAM}}(\text{st}, i; \text{EM})$.

Lemma 3 (OramTethys). *The SSE scheme OramTethys is correct and $\mathcal{L}_{\text{len-hid}}$ -adaptively secure. Let $\varepsilon > 0$ and $p = \Omega(\lambda)$. There are constants $\delta > 1$ and $c \in \mathbb{N}$ such that it has constant client storage, constant storage efficiency and $\mathcal{O}(\log^\varepsilon \lambda)$ page efficiency. Further, the scheme has $\mathcal{O}(1)$ locality if each list fits into a constant number of pages.*

Proof. We first show that OramTethys is correct and $\mathcal{L}_{\text{len-hid}}$ -adaptively secure. Then, we analyze the efficiency for arbitrary constants c and δ from which we conclude the existence of ε .

(1) As Tethys is correct, it remains to show that all remaining identifiers are fetched from the stash when searching a keyword w . First, note that we store $n = \mathcal{O}(\log^{1+\delta}(\lambda))$ blocks in LocORAM. The scheme LocORAM is correct if the blocksize p is $\Omega\left(n^{\frac{c-1}{c}}\right) = \Omega\left(\log^{\frac{1+\delta(c-1)}{c}} \lambda\right)$ which holds as δ, c are constant and $p = \Omega(\lambda)$ by assumption. As Binpack packs the stash into bins of size p and the accessed index $i = T_{\text{pos}}[w]$ corresponds to the bin containing the identifiers matching keyword w , the scheme OramTethys is correct.

(2) The security follows directly as Tethys is \mathcal{L} -adaptively secure, LocORAM is adaptively secure (with zero-failure probability) and the fact that T_{pos} is encrypted.

(3) We now analyze the efficiency of OramTethys. As the client state of LocORAM is $\mathcal{O}(1)$ and the instantiation of Tethys only stores its keys on the client side (as the stash is stored on the server), OramTethys only requires constant client storage. The storage efficiency of OramTethys is $\mathcal{O}(1)$ because Tethys has constant storage efficiency, and EM and T_{pos} have size $n = \mathcal{O}(\log^{1+\delta}(\lambda))$ pages and entries respectively. We now inspect the page efficiency. First, note that Tethys has constant page efficiency and the access to T_{pos} requires (at most) one page access. The access of the stash through LocORAM requires bandwidth $\mathcal{O}(\beta \cdot n^{1/c} \log^2(n)) = \mathcal{O}\left(p \cdot \log^{\frac{1+\delta}{c} + \delta'} \lambda\right)$ for some arbitrary $\delta' > 0$. As c, δ and δ' are arbitrary constants, scaling them accordingly yields the desired result. Constant locality follows directly from the fact that LocORAM has constant locality and that Tethys accesses at most $\mathcal{O}(\ell/p)$ pages for a search on keyword w , where ℓ is the length of the list of identifiers matching w . (Recall that we assume that $\mathcal{O}(\ell/p)$ is constant for all lists.) \square

B.3 The Scheme

Finally, we describe our unconditional static SSE scheme **UncondSSE** with $\mathcal{O}(\log^\varepsilon(N))$ locality, for arbitrary $\varepsilon > 0$. We follow the high level idea of [DPP18] to handle lists with different schemes depending on the list size. For $d \in \mathbb{N}$, we split the interval $[1, N]$ of possible list lengths into four different subintervals.

1. For the subinterval $[1, N^{1-1/\log \log \lambda})$, the lengths are sufficiently for **LayeredSSE** and can simply store the lists using **Local[LayeredSSE]**. Here, the read efficiency is $\mathcal{O}(\log \log N)$.
2. For the subinterval $[N^{1-1/\log \log \lambda}, N/\log^d N)$, the lengths are simultaneously small enough for the local transformation and large enough for **OramTethys**. Thus, we store the lists using **Local[OramTethys]** with $\mathcal{O}(\log^\varepsilon N)$ read efficiency.
3. We further split the subinterval $[N/\log^d N, N/\log^\varepsilon N)$ into a constant number of subintervals, such that **OramTethys** has $\mathcal{O}(\log^\varepsilon N)$ read efficiency.
4. For the subinterval $[N/\log^\varepsilon N, N]$, lists are large enough to read the entire database. Thus, we simply encrypt DB and fetch it entirely from the server for these lists.

We now present how to divide the interval $[N/\log^d N, N/\log^\varepsilon N)$ into subintervals in more detail.

Handling the Remaining List Sizes. Note that for lists of size in $[N/\log^\varepsilon N, N)$, we can just store an encrypted copy of the database on the server and retrieve the entire copy for each read. We now sketch how we handle the remaining lists of size in $S = [N/\log^d N, N/\log^\varepsilon N)$ for some arbitrary $d \in \mathbb{N}$ and $\varepsilon \in (0, 1)_{\mathbb{R}}$. For this, we split the interval S into a constant number of subintervals S_i such that the borders of each interval differ by a factor $\log^\varepsilon N$. Concretely, we set

$$S_i = [N/\log^{d-i\varepsilon} N, N/\log^{d-(i+1)\varepsilon} N) \text{ for } i \in [0, \lceil d/\varepsilon \rceil].$$

For each S_i , we store lists of size in S_i via **OramTethys** with page size $p = \max(S_i)$. Note that each list has at most size p . Thus, **OramTethys** has $\mathcal{O}(1)$ locality and $\mathcal{O}(\log^\varepsilon \lambda)$ read efficiency (see Lemma 3). Note that page efficiency directly translates to read efficiency in this case, as each list is at most a factor of $\log^\varepsilon \lambda$ smaller than the page size.

UncondSSE. We now present our static SSE scheme **UncondSSE**. For a given $\varepsilon > 0$, it has unconditionally $\mathcal{O}(\log^\varepsilon N)$ read efficiency, constant locality and constant storage efficiency. Let $d \in \mathbb{N}$ be the parameter of the local transformation chosen appropriately.

UncondSSE.KeyGen(1^λ). Generate key K_1 for **Local[LayeredSSE]**, key K_2 for **Local[OramTethys]**, key K_3 for **OramTethys** and encryption key K_4 for **Enc**. Also, generate key K_{PRF} for pseudorandom function PRF mapping to $\{0, 1\}^{\lceil \log(N) \rceil}$. Output $K = (K_1, K_2, K_3, K_4, K_{\text{PRF}})$.

UncondSSE.Setup(K, N, DB). First, we initialize a table T_{len} that stores the encrypted length $\ell_i \oplus m_i$ of each list $DB(w_i)$ at position $T_{\text{len}}[w]$, where $m_i \leftarrow \text{PRF}_{K_{\text{PRF}}}(w_i)$ is a mask. Then, we pad T_{len} up to size N (with random values of $\lceil \log N \rceil$ bits). We split the interval of possible list lengths $[1, N]$ into four different subintervals and handle each subinterval separately. For each subinterval, we define four databases DB_i containing a subset of the keyword-identifier pairs of the given database DB (chosen with respect to the lists length). We set

$$\begin{aligned} DB_1 &= \left\{ DB(w_i) \mid \ell_i \in \left[1, N^{1-\frac{1}{\log \log N}}\right] \right\}, \\ DB_2 &= \left\{ DB(w_i) \mid \ell_i \in \left[N^{1-\frac{1}{\log \log N}}, N/\log^d N\right] \right\}, \\ DB_3 &= \left\{ DB(w_i) \mid \ell_i \in \left[N/\log^d N, N/\log^\varepsilon N\right] \right\}, \\ DB_4 &= \{DB(w_i) \mid \ell_i \in [N/\log^\varepsilon N, N]\}. \end{aligned}$$

The lists in DB_1 are sufficiently small for **LayeredSSE** and thus, we can apply the local transformation and run $\text{EDB}_1 \leftarrow \text{Local}[\text{LayeredSSE}].\text{Setup}(K_1, N, DB_1)$. Note that we still pad the encrypted database to size $\mathcal{O}(N)$ and not $\mathcal{O}(|DB_1|)$ because we can not reveal the distribution of lists amongst each subinterval. The lists in DB_2 are sufficiently large for **OramTethys**. Consequently, we can set $\text{EDB}_2 \leftarrow \text{Local}[\text{OramTethys}].\text{Setup}(K_2, N, DB_2)$. For DB_3 , we further split the interval $[N/\log^d N, N/\log^\varepsilon N]$ into the constant number of subintervals $S_i = [N/\log^{d-i\varepsilon} N, N/\log^{d-(i+1)\varepsilon} N]_{\mathbb{R}}$ for $i \in [0, \lceil d \cdot \varepsilon^{-1} \rceil]$ as described above. We then set $DB_{3,i} = \{DB(w_i) \mid \ell_i \in S_i\}$ and $\text{EDB}_{3,i} \leftarrow \text{OramTethys}.\text{Setup}(K_3, N, DB_{3,i})$. Finally, set $\text{EDB}_3 = (\text{EDB}_{3,1}, \dots, \text{EDB}_{3, \lceil d \cdot \varepsilon^{-1} \rceil})$. Lastly, lists in DB_4 are large enough that we can scan entire database each read. For this, we pad DB_4 up to size N and set $\text{EDB}_4 \leftarrow \text{Enc}_{K_{\text{Enc}}}(DB_4)$. Outputs $\text{EDB} = (\text{EDB}_1, \text{EDB}_2, \text{EDB}_3, \text{EDB}_4, T_{\text{len}})$.

UncondSSE.Search($K, w; \text{EDB}$). For retrieving the identifiers matching keyword w , the client sends w and $m \leftarrow \text{PRF}_{K_{\text{PRF}}}(w)$ to the server. The server decrypts the length $\ell \leftarrow T_{\text{len}}[w] \oplus m_i$ of the list to be fetched and then checks in which subinterval ℓ lies. We distinguish four cases: (1) If $\ell \in [1, N^{1-\frac{1}{\log \log N}})$, the client retrieves the identifiers from the server via $\text{Local}[\text{LayeredSSE}].\text{Search}(K_1, w; \text{EDB}_1)$. (2) If $\ell \in [N^{1-\frac{1}{\log \log N}}, N/\log^d N)$, the client runs $\text{Local}[\text{OramTethys}].\text{Search}(K_2, w; \text{EDB}_2)$ with the server. (3) If $\ell \in [N/\log^d N, N/\log^\varepsilon N)$, the server sets $i \in [0, \lceil d \cdot \varepsilon^{-1} \rceil]$ such that $\ell \in S_i$. Then, server and client run $\text{OramTethys}.\text{Search}(K_3, w; \text{EDB}_{3,i})$. (4) Otherwise, we have $\ell \geq N/\log^\varepsilon N$ and the server sends the entire encrypted database EDB_4 to the client (from which he fetches the corresponding list).

Theorem 8 (UncondSSE). *The scheme UncondSSE is correct and $\mathcal{L}_{\text{len-rev}}$ -adaptively secure. It has constant client storage, $\mathcal{O}(1)$ locality and $\mathcal{O}(\log^\varepsilon N)$ read efficiency for any $\varepsilon > 0$.*

Proof. Security (and correctness) directly follow from the security of $\text{Local}[\text{LayeredSSE}]$, $\text{Local}[\text{OramTethys}]$ and **OramTethys**. The efficiency properties of **UncondSSE** can

also be derived from the efficiency properties of the used SSE schemes (see discussion above). \square

Remark 2 (On RTT and Deamortization). The scheme UncondSSE uses LocORAM and thus, the efficiency properties are amortized. Also, this introduces a large round trip time for some operations. We note that if LocORAM is deamortized, we can adapt UncondSSE in order to have a constant RTT and deamortized efficiency.

C Analysis of L2C

Here, we proof Theorem 1. First, we introduce some additional preliminaries.

Lemma 4 (Chernoff's Bound). *Suppose that X_1, \dots, X_n are independent random variables taking values in $\{0, 1\}$. Let X denote their sum and let $\mu = \mathbb{E}[X]$ denote the expectancy of X . Then for any $\delta > 0$, it holds that*

$$\Pr[X < (1 - \delta)\mu] \leq e^{-\frac{\delta^2\mu}{2}}$$

In the next lemma, we consider a sequence of ball insertions and deletions of arbitrary length, such that the total number of balls in the bins at any point in time is bounded by $n = h \cdot m$. A ball insertion is a standard 2-choice insertion: pick two bins i.u.r., and insert the ball into the least loaded bin. A deletion removes one previously inserted ball. The sequence of additions and deletions is fixed at the input of the problem.

Lemma 5 (2C). *Let $\delta(m)$ be an arbitrary map such that $1 \leq \delta(m) \leq \log m$ for all $m \geq 1$. At the outcome of the sequence of additions and deletions, the most loaded bin contains $O(h + \delta(m) \log \log m)$ items, except with probability $m^{-\Omega(\delta(m) \log \log m)}$.*

In particular, by setting $\delta = 1$, we get that if $m \geq \lambda$, then the failure probability from the claim is negligible. By setting $\delta = \log \log \log m$, we get that if $m \geq \lambda^{1/\log \log \lambda}$, then the failure probability from the claim is negligible.

Proof. We adapt the proof of [Vöc03], which proves a bound $O(h) + \log \log m$ with probability $m^{-\alpha}$, for an arbitrary constant α . The proof uses *witness trees*. The existence of a bin containing more than $Ch + L$ items implies the existence of a witness tree of height $L + C'$, for some suitable constants C, C' . Thus, in order to bound the probability that a bin contains more than $Ch + L$ items, it suffices to bound the probability that a witness tree of height $L + C'$ exists. In more detail, the proof shows that the probability that a witness tree of height $L + 3$ exists is upper-bounded by

$$m^{-\kappa+1+o(1)} + m^{-\alpha}$$

where κ, α are certain parameters (to be discussed later), with:

$$L \leq \log \log m + \log(1 + \alpha) + \kappa.$$

The proof sets α and κ to be constants. The fact that γ and κ are constant is not essential to the argument, and is only used in two places in the proof.

The first place is the end of Section 2.3, when upper-bounding the probability of activation of a pruned witness tree by $m^{-\kappa+1+o(1)}$. The final step of that upper-bound requires $\alpha \cdot \kappa = m^{o(1)}$, which is obviously true for a pair of constants.

The other, more important place where the choice of having constant α and κ comes into play is in the final derivation. The proof shows that, except with probability at most $m^{-\kappa+1+o(1)} + m^{-\alpha}$, the number of items in the most loaded bin is at most:

$$\begin{aligned} L + O(h) &\leq \log \log m + \log(1 + \alpha) + \kappa + O(h) \\ &= \log \log m + O(1) + O(h) \\ &= \log \log m + O(h). \end{aligned}$$

In that final computation, the fact that α and κ are constant makes it possible to absorb the $\log(1 + \alpha) + \kappa$ term into the $O(h)$ term. The other term is only $\log \log m$, which is optimal. If we set $\alpha = \kappa = \delta(m) \log \log m$ instead, we get:

$$\begin{aligned} L + O(h) &\leq \log \log m + \log(1 + \alpha) + \kappa + O(h) \\ &\leq 3\delta(m) \log \log m + O(h). \end{aligned}$$

In the case $\delta = 1$, this worsens the constant in front of the $\log \log$ term, which is likely why the authors chose α and κ to be constant. (A better constant than 3 is possible, we choose 3 for simplicity.) On the other hand, the probability of failure becomes at most

$$m^{-\kappa+1+o(1)} + m^{-\alpha} = m^{-\Omega(\delta(m) \log \log m)}$$

as claimed. Note that the condition $\alpha \cdot \kappa = m^{o(1)}$ is still fulfilled. \square

The next lemma is a direct application of Markov's inequality.

Lemma 6. *For any random variable $X \in [0, N]_{\mathbb{R}}$ and any $R > 0$ (which may depend on N):*

$$\Pr[X > R] = \text{negl}(\lambda) \quad \text{iff} \quad \mathbb{E}[\max(X - R, 0)] = \text{negl}(\lambda).$$

Lemma 7 (Weighted 1C [BFHM08]). *Let $m \in [0, 1]_{\mathbb{R}}$ be some maximal weight. Let $x = (m)_{i \leq n}$ and $x' = (w'_i)_{i \leq n'}$ be (non-negative) weight vectors. Let $\sum_{i=1}^{n'} w'_i \leq n \cdot m$ and $w_i \leq m$ for all $i \in \{1, \dots, n'\}$ ¹². Let $R \in \mathbb{R}^+$. Then it holds that $\mathbb{E}[\max(X_{\text{mlb}} - R, 0)] \geq \mathbb{E}[\max(X'_{\text{mlb}} - R, 0)]$, where X_{mlb} (X'_{mlb}) are a random variable indicating the load of the most loaded bin after throwing n balls with weights x (n' balls with weights x') uniformly and independently at random into m bins¹³.*

¹² [BFHM08] requires that x majorizes x' . This is implied by our condition on x and x' .

¹³ [BFHM08] shows that $\mathbb{E}[X_{\text{mlb}}] \geq \mathbb{E}[X'_{\text{mlb}}]$. As $f(X) = \max(X - R, 0)$ is convex, their proof can be adapted to our formulation.

In words, for 1C, the load above threshold R of the most loaded bin is higher with balls of weight x than with balls of weight x' .

We are now ready to prove Theorem 1.

Proof. Note that the load of a bin is never decreasing, so it is sufficient to analyze the final load of bins B_1, \dots, B_m . Also, note that we can replace Setup with n InsertBall operations. Thus, we can assume without loss of generality that bins B_1, \dots, B_m are initially empty after L2C.Setup. Also, note that $m^{-\Omega(\delta(\lambda) \log \log w)} = \text{negl}(\lambda)$ under the given requirements (see lemma 5). As H is modeled as a random oracle, we assume that the bin choices α_1, α_2 of ball b are chosen independently and uniformly at random from $[1, m]^2$. We split the proof into three parts:

(1) First, we will modify the sequence S such that we can reduce the analysis to only (sufficiently independent) L2C.InsertBall operations, while only increasing the final bin load by a constant factor.

(2) Second, we analyze the maximal bin load when only considering balls of weight at most $1/\log m$. Here, L2C.InsertBall proceeds exactly as weighted 1C. Since uniform weights of value $1/\log m$ are the worst case for the most loaded bin in 1C, the bound follows from a Chernoff's bound as balls are sufficiently small.

(3) Last, we inspect the maximal bin load considering items in the remaining subintervals $(2^{i-1}/\log m, 2^i/\log m]_{\mathbb{R}}$ for $i \in \{1, \dots, \log \log m\}$. Per interval, L2C.InsertBall behaves like unweighted two-choice (independent of other subintervals) and inherits the $\log \log m$ bin load directly, as balls with different weights differ only by a constant factor per interval. Summing up the maximal bin load per interval will yield the desired result.

Part 1 – Adapting the sequence: We observe that update operations updating the weight inside the same subinterval can be ignored. More concretely, let $\text{op}_i = \text{UpdateBall}$ be some update operation on ball b_i with old weight o_i and new weight w_i . If $o_i, w_i \in (\frac{2^{k-1}}{\log m}, \frac{2^k}{\log m}]_{\mathbb{R}}$ for some k , the operation op_i replaces the old weight o_i of ball b_i with the new weight w_i directly (inside the same bin). Thus, we can simply remove op_i and replace the previous operation $\text{op}_j = (b_i, o_j, o_i) = (b_j, o_j, w_j)$ on the same ball with $\text{op}'_j = (b_i, o_j, w_i)$ directly. Clearly this does not change the final load of the bins. (Note that operations between op_j and op_i make the same choices as the concrete weight inside a subinterval never impacts which bin is chosen.)

Now, let $(\text{op}_i)_{i \in I}$ be all remaining update operations for some fixed ball b_* , so $b_* = b_i$ and $\text{op}_i = \text{UpdateBall}$ for $i \in I$. As we removed consecutive update operations in the same subinterval, operation op_i marks the ball b_* as residual ball and calls $\text{InsertBall}(b_*, w_i, B_{\alpha_{*,1}}, B_{\alpha_{*,2}})$. Let $j = \max(I)$ be the index of the last update operation op_j on b_* and k be minimal such that $w_j \leq 2^k/\log m$. As there are only k subintervals below the last interval $(2^{k-1}/\log m, 2^k/\log(m)]_{\mathbb{R}}$, there are at most k such update operations, *i.e.* $|I| \leq k$, and one insert ball operation. Assume without loss of generality that all $k+1$ operations exist. The residual ball left by the i -th update operation has at most size $2^{i-1}/\log m$ and thus, this UpdateBall

operation can be replaced by an $\text{InsertBall}(b_*, 2^{i-1}/\log m, (B_{\alpha_{*,1}}, B_{\alpha_{*,2}}))$ operation. Thus, for ball b_* with final weight w_j , we have to insert k additional balls in order to replace all update operations on ball b_* with inserts. The total weight of these additional balls is

$$\sum_{i=1}^k 2^{i-1}/\log m \leq 2^k/\log(m) \leq 2w_j,$$

since $w_j \geq 2^{k-1}/\log(m)$. Thus, the total weight is increased at most by a factor 3 per ball.

This way, we can iteratively remove all UpdateBall operations at the cost of a factor 3 in the total weight. The remaining operations are InsertBall operations, where each ball b_i is inserted at most once per subinterval and the bin choices are drawn uniformly and independently random per ball. Clearly, if $\mathcal{O}(3 \log \log w_{\max})$ is an upper bound on the load of the most loaded bin for the modified sequence S' , then $\mathcal{O}(3 \log \log w_{\max})$ is an upper bound for the initial sequence S . In the following, we only consider modified sequences S of n such InsertBall operations.

Part 2 – Light balls: Here, we show that the most loaded bin has load at most $3\delta(\lambda) \log \log w_{\max}$ when only considering balls of at most weight $1/\log m$. Let $w \leq w_{\max}$ be the total weight of all such light balls. Without loss of generality, assume that $w_{\max} = w$. At first, we assume that all such balls have weight exactly $1/\log m$ each. We will then reduce the case with arbitrary weights in $[0, 1/\log m]_{\mathbb{R}}$ to the above.

Since we initially assume all balls have weight $1/\log m$, the number of balls is at most $n' = w \log m$. Let X_i be the random variable that denotes the number of balls in bin B_i . Recall that $m = \frac{w}{\delta(\lambda) \log \log w}$. We observe that InsertBall behaves like 1C in this case and thus, we have $\mathbb{E}[X_i] = n'/m = \delta(\lambda) \log \log w \cdot \log m$. Applying Chernoff's bound (Lemma 10), we get:

$$\Pr[X_i \geq (1 + \gamma) \mathbb{E}[X_i]] \leq \exp\left(-\frac{\gamma^2 \mathbb{E}[X_i]}{2}\right).$$

We insert $\gamma = 2$ in the equation above and receive:

$$\Pr[X_i \geq 3\delta(\lambda) \log \log w \cdot \log m] \leq \exp(-2\delta(\lambda) \log \log w \cdot \log m) = m^{-\Omega(\delta(\lambda) \log \log w)}.$$

A union bound yields that the most loaded bin contains at most $3\delta(\lambda) \log \log w \cdot \log m$ balls with probability at most $m^{-\Omega(\delta(\lambda) \log \log w)} = \text{negl}(\lambda)$. As each ball has size $1/\log m$, the most loaded bin has a maximal load of $3\delta(\lambda) \log \log w$ with overwhelming probability.

Now, we show this bound is preserved when allowing arbitrary weights of at most $1/\log m$. We define the weight vectors $x = (1/\log m)_{i=1}^{w \log m}$ and $x' = (w_i)_{i \in I}$. Let X_{mlb} and X'_{mlb} be the random variable indicating the load of the most loaded bin with weights x and with weights x' respectively. We want to show that $\Pr[X'_{\text{mlb}} > 3\delta(\lambda) \log \log w] = \text{negl}(\lambda)$. Lemma 12 implies that the above holds iff $\mathbb{E}[\max(X'_{\text{mlb}} - 3\delta(\lambda) \log \log w, 0)] = \text{negl}(\lambda)$. This expectancy can be upper

bound by $\mathbb{E}[\max(X_{\text{mlb}} - 3\delta(\lambda) \log \log w, 0)]$ as x and x' fulfill the requirements of Lemma 7. As we showed above that $\Pr[X_{\text{mlb}} > 3\delta(\lambda) \log \log w] = \text{negl}(\lambda)$, we can conclude from another application of Lemma 12.

Part 3 – Heavy balls: So far, we have shown that the most loaded bin has load at most $3\delta(\lambda) \log \log w$ with overwhelming probability, when only considering balls of weight smaller or equal to $1/\log m$, for any (modified) sequence S . We will now show that when considering the remaining balls of weight in $(1/\log m, 1]_{\mathbb{R}}$, a maximal load of $\mathcal{O}(w_{\text{max}}/m + \delta(\lambda) \log \log w) = \mathcal{O}(\delta(\lambda) \log \log w)$ is preserved.

For $i \in [1, \log \log m]$, let n_i be the number of balls in each subinterval $A_i = (2^{i-1}/\log m, 2^i/\log m]_{\mathbb{R}}$. Recall that each ball b_i has two bin choices that are drawn uniformly and independently random at the first insertion. These choices are reutilized across the subintervals A_i , if b_i is inserted in multiple subintervals. But note that per subinterval, b_i is only inserted once. Thus, L2C behaves like unweighted 2C on all balls with weights in A_i (independent from the balls in other subintervals). By Lemma 5, the bin with the highest number of balls (of weights in A_i) contains at most $\mathcal{O}(n_i/m + \delta(\lambda) \log \log m)$ balls with overwhelming probability. (Note that there are at most $w \log m$ balls and that $m^{-\Omega(\delta(\lambda) \log \log w)} = \text{negl}(\lambda)$.)

Each ball has weight at most $\max(A_i) = 2^i/\log m$ and thus, the load of the most loaded bin is at most $2^i/\log m (\mathcal{O}(n_i/m + \delta(\lambda) \log \log m))$ when considering balls with weights in A_i . Summing over all A_i 's, when considering only balls with weights in $(1/\log m, 1]$, the load of the most loaded bin is at most

$$\begin{aligned} & \sum_{i=1}^{\log \log m} \frac{2^i}{\log m} \mathcal{O}(n_i/m + \delta(\lambda) \log \log m) \\ &= \sum_{i=1}^{\log \log m} \mathcal{O} \left(2 \frac{n_i 2^{i-1}}{m \log m} + \sum_{i=1}^{\log \log m} \frac{2^i}{\log m} \mathcal{O}(\delta(\lambda) \log \log m) \right) \\ &\leq \mathcal{O} \left(\frac{w_{\text{max}}}{m} + \delta(\lambda) \log \log w_{\text{max}} \right), \end{aligned}$$

as $m = \mathcal{O}(w_{\text{max}})$ and w_{max} is an upper bound on the total weight. The above holds with overwhelming probability, since the probability that 2C fails is $\text{negl}(\lambda)$, and there are only $\log \log m$ subintervals.

As we showed in the first part that it suffices to look at the modified sequence (with only InsertBall operations), we conclude that the load of the most loaded bin is at most $\mathcal{O}(\log \log w_{\text{max}})$. \square

D Security Analysis of LayeredSSE

Lemma 8 (Correctness). *The scheme LayeredSSE is correct if at most p identifiers are associated to each keyword and \mathbf{H} is modeled as a random oracle.*

Proof. We use L2C to insert (and update) the lists of identifiers $\text{DB}(w)$ of length $\ell \leq p$ into m bins. Each list is interpreted as a ball of weight $\ell/p \in [0, 1]$. Theorem 1 implies that the maximal loaded bin has load at most $c \log \log(\lambda)$

$\log \log(N/p)$ for some appropriate constant $c \in \mathbb{N}$ (for $\delta(\lambda) = \log \log \log(\lambda)$), since the bin choices via H are uniformly and independently random by assumption. That means, it contains at most $p \cdot c \log \log \log(\lambda) \log \log(N/p)$ identifiers (as we scaled weights by a factor p). Consequently, the bins only overflow with negligible probability. Further, it follows from inspection that one of the two bins returned by the search algorithm on input w contains all the identifiers matching keyword w . \square

Lemma 9 (Selective Security). *Let $\mathcal{L}_{\text{Stp}}(\text{DB}, N) = N$, $\mathcal{L}_{\text{Srch}}(w) = \text{qp}$ and $\mathcal{L}_{\text{Updt}}(\text{op}, w, L') = \text{qp}$, where qp is the query pattern and $\text{op} = \text{add}$. Let $\mathcal{L} = (\mathcal{L}_{\text{Stp}}, \mathcal{L}_{\text{Srch}}, \mathcal{L}_{\text{Updt}})$. The scheme LayeredSSE is \mathcal{L} -selectively semantically secure if at most p identifiers are associated to each keyword, Enc is IND-CPA secure and H is modeled as a random oracle. Note that $\mathcal{L} = \mathcal{L}_{\text{len-hid}}$ because we restrict ourselves to lists of size at most p .*

Proof. Let Sim denote the simulator and \mathcal{A} an arbitrary honest-but-curious PPT the adversary.

Initially, Sim receives $\mathcal{L}_{\text{Stp}}(\text{DB}, N) = N$ and a series of search and update requests with input $\mathcal{L}_{\text{Updt}}(\text{op}_i, w_i, L'_i) = \mathcal{L}_{\text{Srch}}(w_i) = \text{qp}$. First, Sim initializes $m = \lceil (N/p) / (\log \log(N/p) \log \log \log(\lambda)) \rceil$ bins B_1, \dots, B_m zeroed out up to size $p \cdot c \log \log \log(\lambda) \log \log(N/p)$, and outputs $\text{EDB}' = (\text{Enc}_{\mathsf{K}'_{\text{Enc}}}(B_1), \dots, \text{Enc}_{\mathsf{K}'_{\text{Enc}}}(B_m))$ for some encryption key K'_{Enc} sampled by Sim . Next, Sim simulates the search and update queries.

For search queries, Sim receives sp . If the query pattern sp indicates that the keyword was already queried, Sim outputs the keyword w' from the previous query. Otherwise, Sim outputs a new uniformly random keyword w' (that has not been queried yet).

For update queries, Sim receives sp . First, Sim proceeds as in search for generating the first output w' . After sending w' to the adversary \mathcal{A} , Sim receives two encrypted bins. Sim simply reencrypts both bins and sends them back to the server.

We now show that the real game is indistinguishable from the ideal game. For this, we define four hybrid games.

- Hybrid 0 is identical to the real game.
- Hybrid 1 is the same as Hybrid 0 except the simulated keywords w' are output. By assumption Hybrid 0 and Hybrid 1 are indistinguishable.
- Hybrid 2 is the same as Hybrid 1 except a flag **FAIL** is raised when a bin overflows (*i.e.* contains more than $p \cdot c \log \log N/p$ identifiers after **Setup** or **Update**). Theorem 1 implies that this happens only with negligible probability. Thus, Hybrid 1 and Hybrid 2 are indistinguishable.
- Hybrid 3 is the same as Hybrid 2 except that the encrypted database EDB is replaced with the simulated EDB' and bins are just reencrypted and sent back to the adversary in the second flow of **Update**. Since Enc is IND-CPA secure (and a flag **FAIL** is only raised with negligible probability), it follows that Hybrid 2 and Hybrid 3 are indistinguishable.

- Hybrid 4 is the same as the ideal experiment. The server’s view in the ideal experiment and in Hybrid 3 are identically distributed, so we conclude inductively that the ideal game and the real game are indistinguishable. \square

E Proof Of ClipOSSE

In this section, we prove Theorem 3. Recall that it is assumed all lists have length at most $N/\log^d \lambda$ for $d \geq 2$. This is a limitation of the result, and it is inherent (lists of length close to $N/\log \lambda$ can create too many overflowing elements, and must be handled separately.)

The proof is divided into two parts. First, we show that the result holds when all lists have size exactly $N/\log^d \lambda$. Second, we show that the result still holds as long as all lists have size at most $N/\log^d \lambda$. The second part is the hard part.

E.1 Proof Part 1: All Lists Have Size $N/\log^d \lambda$.

We recall one of the standard formulations of the Chernoff-Hoeffding bound.

Lemma 10 (Chernoff-Hoeffding). *Let $X = \sum_{i \leq n} X_i$ where the X_i ’s are i.i.d. 0-1 random variables, with $p = \mathbb{E}[X_i]$.*

$$\Pr[X > (p + \varepsilon)n] < e^{-\frac{1}{2}\varepsilon^2 n / (p + \varepsilon)}$$

$$\Pr[X < (p - \varepsilon)n] < e^{-\frac{1}{2}\varepsilon^2 n / p}.$$

The following lemma is a direct corollary.

Lemma 11. *Throw n balls into m bins u.i.r. Let $\mu = n/m$ be the average load of a bin. Then the probability that a given bin contains more than $\gamma\mu$ balls is at most:*

$$e^{-\Theta(\gamma)\mu}.$$

Proof. Use Lemma 10 with $p = 1/m$, $\varepsilon = (\gamma - 1)p$. \square

Let $\tau = \beta \log \log \lambda$ be the threshold at which a bin starts to overflow.

By construction, 1C with all lists of size $N/\log^d \lambda$ is exactly a balls-and-bins game with $n = \log^d \lambda$ balls (each ball is a list) and $m = n/\log \log \lambda$ bins. By Lemma 11, the most loaded bin contains less than $\log \lambda \log \log \lambda$ elements, except with negligible probability. We now want to bound the number of bins that overflow. Using the previous lemma again, the probability that a given bucket overflows is $e^{-\Omega(\beta) \log \log \lambda} = \log^{-\Omega(\beta)} \lambda$.

Let $(X_i)_{i \leq m}$ denote the indicator variables that are equal to 1 iff the i -th bucket overflows, 0 otherwise. The number of overflowing buckets is $X = \sum X_i$. We know that $\mathbb{E}[X_i] = \log^{-\Omega(\beta)} \lambda$. We want to show that X cannot be much higher than $m \log^{-\Omega(\beta)} \lambda$.

For that purpose, we use the notion of negative association. Because that notion is only used briefly to establish that Chernoff-Hoeffding bounds apply,

we do not develop the theory here, and instead refer the reader to [DR96] for an excellent survey on the topic. By [DR96, Proposition 13], the occupancy numbers (vector $(B_{n,m}[i])_{i \leq m}$ where $B_{n,m}[i]$ is the number of balls in the i -th bin) are negatively associated. By [DR96, Proposition 7.2], since $X_i = \mathbf{1}_{B_{n,m}[i] \geq \tau}$, and $x \mapsto \mathbf{1}_{x \geq \tau}$ is non-decreasing, the X_i 's are also negatively associated. By [DR96, Proposition 5], it follows that we can apply Chernoff-Hoeffding bounds to $X = \sum X_i$.

Hence, using Lemma 10 with $p = \varepsilon = \mathbb{E}[X_i] = \log^{-\Omega(\beta)} \lambda$, we get:

$$\Pr \left[X > 2m \log^{-\Omega(\beta)} \lambda \right] = e^{-\frac{1}{4}\varepsilon m} = e^{-\frac{1}{4}m \log^{-\Omega(\beta)} \lambda}.$$

The above quantity is negligible as soon as $d - \Omega(\beta) \geq 2$.

Since at most $2m \log^{-\Omega(\beta)} \lambda$ buckets overflow, and the most loaded bucket contains at most $\log \lambda \log \log \lambda$ items, we get that with overwhelming probability, the number of overflowing balls is less than:

$$\log^d \lambda \log^{1-\Omega(\beta)} \lambda.$$

Since each ball corresponds to a list containing $N/\log^d \lambda$ items, with overwhelming probability the number of overflowing *items* is:

$$N \log^{1-\Omega(\beta)} \lambda$$

so we can make it $O(N/\log^c \lambda)$ for any constant c of our choice by picking β suitably (and then picking d to satisfy the condition $d - \Omega(\beta) \geq 2$ encountered earlier).

E.2 Proof Part 2: General Case

Let L denote an arbitrary multiset of list lengths, with $\max L \leq N/\log^d \lambda$, and $\sum L = N$. Let X_L be (the random variable denoting) the number of overflowing elements after inserting the lists in L .

In Part 1 of the proof, we have seen that $\Pr[X_L > R] = \text{negl}(\lambda)$, for a suitable R , when all lists have size $N/\log^d \lambda$. Our goal is to show a similar result for arbitrary L .

Preliminary groundwork. We are going to work with $\mathbb{E}[\max(X_L - R, 0)]$, rather than $\Pr[X_L > R]$. This is made possible by the following lemma.

Lemma 12. *For any integral random variable $X \in [0, N]$ and any $R \geq 0$:*

$$\Pr[X > R] = \text{negl}(\lambda) \quad \Leftrightarrow \quad \mathbb{E}[\max(X - R, 0)] = \text{negl}(\lambda).$$

Proof. By a classic inequality, for any positive integral random variable Y , $\mathbb{E}[Y] = \sum_{i \geq 0} \Pr[Y > i]$. It follows that $\Pr[Y > 0] \leq \mathbb{E}[Y]$. On the other hand, if $Y \leq N$, we get $\mathbb{E}[Y] \leq N \Pr[Y > 0]$. Hence:

$$\Pr[Y > 0] \leq \mathbb{E}[Y] \leq N \Pr[Y > 0].$$

Since $N = \text{poly}(\lambda)(\lambda)$, it follows that $\Pr[Y > 0]$ is negligible iff $\mathbb{E}[Y]$ is negligible. The lemma is obtained by applying that observation to $Y = \max(X - R, 0)$. \square

The following simple lemma will also be useful.

Lemma 13. *Let X, Y be two random variables defined on the same sample space. Let \mathcal{E} be a set of events that forms a partition of the sample space (i.e. pairwise disjoint events whose union is the whole space). If the conditional expectations satisfy $\mathbb{E}[X : E] \leq \mathbb{E}[Y : E]$ for all $E \in \mathcal{E}$, then $\mathbb{E}[X] \leq \mathbb{E}[Y]$.*

Notation

- Let $B(p)$ denote the Bernoulli distribution with mean p : that is, a sample of $B(p)$ is a 0-1 random variable X such that $\Pr[X = 1] = p$ and $\Pr[X = 0] = 1 - p$.
- Let $\text{Bin}(p, n)$ denote the Binomial distribution with n trials, each with probability p : that is, a sample of $\text{Bin}(p, n)$ is distributed like $\sum_{i \leq n} X_i$, where the X_i 's are i.i.d. sampled from $B(p)$.
- Suppose we throw n balls i.u.r. into m buckets. Let $B_{n,m}[i]$ be the (random variable denoting the) load of the i -th bucket. Let $B_{n,m} = (B_{n,m}[i])_{i \leq m}$ be the vector of the load of buckets. Observe that $B_{n,m}[i]$ is distributed according to $\text{Bin}(1/m, n)$. Also note that the $B_{n,m}[i]$'s are not independent, e.g. they are linked by $\sum B_{n,m}[i] = n$.
- If D is a distribution, and E is an event, then $D[E]$ denotes the distribution D conditioned on the event E .
- If X is a random variable, and E is an event, then $\mathbb{E}[X : E]$ denotes the conditional expectation of X , conditioned on the event E .
- If D is a distribution, $X_1, \dots, X_n \leftarrow D$ denotes that X_1, \dots, X_n are i.i.d. random variables, each distributed according to D .

Recall that we are trying to show that the number of overflowing items is bounded by some $R = O(N/\log^c N)$, except with negligible probability, for a constant c of our choice. In Part 1 of the proof, we have already seen that this is true when all lists are of size $N/\log^d N$, for some suitable constant d . In Part 2, we want to prove the same for an arbitrary multiset L of list sizes, assuming $\max L \leq N/\log^d N$. Recall that $\sum L = N$. Let $\tau = \beta \log \log \lambda$ be the threshold at which buckets are cut off. For a given bucket load vector $b = (b[i])_{i \leq m}$, let $\text{over}(b)$ denote the number of overflowing items: $\text{over}(b) = \sum \max(b[i] - \tau, 0)$.

Fix a multiset L of list sizes, with $\max L \leq N/\log^d N$. Let N_ℓ denote the number of lists of size $\ell = 2^i$. Note $\sum \ell N_\ell = N$. Recall that the number of buckets is $m = N/\log \log \lambda$.

Let $D(L)$ be (the random variable denoting) the load of buckets at the output of algorithm 1C, on input L . By abuse of notation, we still write $D(L)$ for a random variable distributed according to $D(L)$. Our goal is to show $\Pr[\text{over}(D(L)) > R] = \text{negl}(\lambda)$. By Step 1, this is equivalent to $\mathbb{E}[\max(\text{over}(D(L)) - R, 0)] = \text{negl}(\lambda)$. To simplify notation, write $F(L) = \max(\text{over}(D(L)) - R, 0)$.

Proof outline. Starting from L , let $\mu = \min L$ be the smallest list size in L . We will merge all N_μ lists of size μ pairwise into $N_\mu/2$ lists of size 2μ . This increases the size of the smallest list in L from μ to 2μ . We can repeat this process as long as the minimum list size μ is less than the maximum list size $N/\log^d N$. Eventually, all lists have size $N/\log^d N$. At that point, we will be able to apply the result from Part 1 of the proof, which deals precisely with the case that all lists have size $N/\log^d N$. This will show that $\mathbb{E}[F(L_{\text{final}})]$ is negligible for the final list L_{final} , obtained after all merging operations are done. In order to show that $\mathbb{E}[F(L)]$ is negligible for the list L we start from, we will show that if $\mathbb{E}[F(L_{i+1})]$ is negligible for a list L_{i+1} obtained *after* a merging operation, then $\mathbb{E}[F(L_i)]$ is also negligible for the list L_i *before* the merging operation. By induction, this will imply that since $\mathbb{E}[F(L_{\text{final}})]$ is negligible, then $\mathbb{E}[F(L)]$ is also negligible for the original list L .

Thus, it suffices to show that if $\mathbb{E}[F(L)]$ is negligible after merging, then it was negligible before merging. This fact is the core of the proof, and involves several techniques. For now, we outline these techniques at a high level, and will provide more details when each technique is introduced. Let $m_\mu = m/\mu$ be the number of superbuckets of size μ . Let us regard lists of size μ as *balls*, and superbuckets of size μ as *bins*. Inserting the lists of size μ amounts to throwing N_μ balls into m_μ bins i.u.r. After merging, a list of size 2μ is viewed as two connected balls. Each pair of connected balls is thrown i.u.r. into two adjacent bins (where the two adjacent bins correspond to one superbucket of size 2μ). When bins are inserted by pairs in that manner, one feature of the resulting distribution is that the bins with even indices (bins number 0, 2, 4, etc) must contain the same total number of balls as the bins with odd indices (bins number 1, 3, 5, etc). When that property is satisfied, let us say that the bins are *balanced*. To recap: inserting merged lists will always yield balanced bins. On the other hand, if we insert lists before the merging step, there is no particular reason that the resulting bins should be balanced. The first main proof technique is to show the following: if we insert lists before the merging step, and condition the resulting distribution of bin occupancies being *balanced*, then the merging operation can only increase $\mathbb{E}[F(L)]$. This step relies on a convexity argument, and uses a special auxiliary operator \diamond . We leave a detailed discussion of those points for later, and continue to focus on the global outline of the proof.

Insofar as merging can only increase $\mathbb{E}[F(L)]$, we get what we want: if $\mathbb{E}[F(L)]$ is negligible after merging, then it was necessarily negligible before merging. However, to apply that argument, we need bins to be balanced. As mentioned earlier, there is no special reason that inserting N_μ balls into m_μ bins i.u.r. should result in balanced bins. This leads to the next proof technique, which is a stochastic dominance argument. Although the distribution obtained by throwing N_μ balls into m_μ bins i.u.r. is not balanced, we show that it is stochastically dominated by balanced distribution, namely the distribution obtained by throwing $N_\mu + \phi(\mu)$ balls into m_μ bins i.u.r. *conditioned on being balanced*. Here, $\phi(\mu)$ is a carefully chosen small quantity. Intuitively, what happens is that although the original distribution may not be balanced, the difference $2\delta = |n_0 - n_1|$ be-

tween the number n_0 of balls in bins with even indices, and the number n_1 of balls in bins with odd indices, must be less than $\phi(\mu)$ (except with negligible probability). As a consequence, by adding less than 2δ balls, we can “correct” the distribution into a balanced one, at the cost of slightly increasing the total number of balls. Since adding new balls can only increase the output of $\text{over}(\cdot)$, this new transformation has the desired property that if $\mathbb{E}[\max(\text{over}(\cdot) - R, 0)]$ is negligible for the distribution at the output of the transformation, it was necessarily negligible before the transformation. On the other hand, because we add new balls, we need to be mindful that each merging increases the total number of balls in the system. However, we show that the total number of balls remains $O(N)$ throughout, which completes the proof.

Full Proof.

Definition 13. Let $a = (a_i)_{i \leq t}$ and $b = (b_i)_{i \leq t}$ be two vectors in \mathbb{N}^t . Then $a \diamond b$ denotes the following vector in \mathbb{N}^{2t} :

$$a \diamond b = (a_1, b_1, a_2, b_2, \dots, a_t, b_t).$$

The notation \diamond is extended in the usual way to combine two sets of vectors ($A \diamond B = \{a \diamond b : a \in A, b \in B\}$), and two distributions of vectors ($D_1 \diamond D_2 = a \diamond b$ where $a \leftarrow D_1, b \leftarrow D_2$). The point of \diamond is the next lemma, which is essentially a convexity argument.

Lemma 14. Let $a = (a_i)_{i \leq t}$ and $b = (b_i)_{i \leq t}$ be two vectors in \mathbb{N}^t . We have:

$$2F(a \diamond b) \leq F(a \diamond a) + F(b \diamond b).$$

Proof. Let $a' \in \mathbb{N}^t$ be defined by $a'_i = \max(a_i - \tau, 0)$, so that a'_i is the number of overflowing elements in bucket i for vector a . (Recall that τ is the threshold at which buckets are cut off.) Define b' in the same way. Observe that $f : x \mapsto \max(x - R, 0)$ is a convex function, which implies that for all x, y , $f(x/2 + y/2) \leq (f(x) + f(y))/2$. As a consequence:

$$\begin{aligned} 2F(a \diamond b) &= 2 \max\left(\sum a'_i + \sum b'_i - R, 0\right) \\ &= 2f\left(\sum a'_i + \sum b'_i\right) \\ &\leq f\left(2 \sum a'_i\right) + f\left(2 \sum b'_i\right) \\ &= F(a \diamond a) + F(b \diamond b). \quad \square \end{aligned}$$

Given a load vector $b \in \mathbb{N}^m$, let $n_0(b) = \sum_{i: i \bmod 2=0} b_i$ (resp. $n_1(b) = \sum_{i: i \bmod 2=1} b_i$) be the total number of balls in bins with even (resp. odd) index. Let $n(b) = n_0(b) + n_1(b)$ be the total number of balls. Let $\delta(b) = \max(n_0(b), n_1(b)) - \lfloor n(b)/2 \rfloor$.

Recall that $B_{n,m}[\delta = 0]$ denotes the distribution $B_{n,m}$ conditioned on the event $\delta = 0$, that is, the bins with even indices contain the same total number of balls as the bins with odd indices. The proof of the following lemma is immediate.

Lemma 15. *For all even n , m :*

$$B_{n,m}[\delta = 0] = B_{n/2,m/2} \diamond B_{n/2,m/2}.$$

Define:

$$B'_{n,m,d} = B_{n,m}[\max(n_0, n_1) \leq n/2 + d].$$

Lemma 16. *If $d = \Omega(\sqrt{n} \log \lambda)$, then the statistical distance between $B_{n,m}$ and $B'_{n,m,d}$ is negligible.*

Proof. By a simple Chernoff bound, the probability that the condition that defines $B'_{n,m,d}$ is not satisfied in $B_{n,m}$ is negligible. Further, if two distributions are identical conditioned on an event with negligible probability not happening, then their statistical distance is negligible. \square

Given two distributions D_1, D_2 on a set equipped with an order relation \preceq , recall that D_1 is said to be stochastically dominated by D_2 if there exists a coupling (X_1, X_2) of D_1 and D_2 (i.e. a random variable (X_1, X_2) such that the marginal distribution of X_i is D_i) such that $\Pr[X_1 \preceq X_2] = 1$.

Lemma 17. *For all n, m, d , $B'_{n,m,d}$ is stochastically dominated by $B_{n+2d,m}[\delta = 0]$ (with respect to the product order on \mathbb{N}^m).*

Proof. If we sample from $B'_{n,m,d}$, then add $n/2 + 2d - n_0$ (resp. $n/2 + 2d - n_1$) balls uniformly at random into buckets of even (resp. odd) indices, we obtain a sample from $B_{n+2d,m}[\delta = 0]$. Hence, there exists a suitable coupling of the two distributions. \square

Let $\phi(\ell) = \sqrt{N/\ell} \log \lambda$. If L is a multiset of list sizes, let $\mu = \min L$. Define $\text{merge}(L)$ by removing all N_μ instances of μ from L , and adding instead $N_\mu/2 + \phi(\mu)$ instances of size 2μ .

Let $N'_1 = N_1 + \sqrt{N} \log \lambda$. By induction, for i in $\{1, \dots, \log N\}$ and $\ell = 2^i$, define:

$$N'_\ell = N_\ell + \frac{N'_{\ell/2}}{2} + \sqrt{\frac{N}{\ell}} \log \lambda.$$

Lemma 18. *For all $\ell \leq N/\log^2 \lambda$, $N'_\ell = O(N/\ell)$.*

Proof. A straightforward induction gives:

$$\begin{aligned} N'_\ell &= N_\ell + \sum_{i=0}^{\log \ell} \sqrt{\frac{N}{2^i}} \log \lambda \\ &= N_\ell + \sqrt{N} \log \lambda \sum_{i=0}^{\log \ell} 2^{-i/2} \\ &= N_\ell + \sqrt{N} \log \lambda \cdot O\left(2^{-\frac{1}{2} \log \ell}\right) \end{aligned}$$

$$\begin{aligned}
 &= N_\ell + O\left(\sqrt{\frac{N}{\ell}} \log \lambda\right) \\
 &= N_\ell + O\left(\frac{N}{\ell}\right) && \text{because } \log^2 \lambda \leq N/\ell \\
 &= O\left(\frac{N}{\ell}\right). && \square
 \end{aligned}$$

Lemma 19. *Let $\mu = \min L$. Assume N_μ is even. If $\mu < N/\log^2 N$, then:*

$$\mathbb{E}[F(L)] \leq \mathbb{E}[F(\text{merge}(L))] + \text{negl}(\lambda).$$

Proof. In the scope of this proof, μ is set to $\min L$. Let $m_\mu = m/\mu$ be the number of superbuckets of size μ . Say that a superbucket is *flat* iff all the buckets it contains have the same number of items. Say that a vector of occupancies $b \in \mathbb{N}^m$ is *k-flat* if all superbuckets of size k are flat.

By construction of 1C, and that fact that $\mu = \min L$, after inserting lists in L , bucket occupancies are μ -flat. The load of a bucket is entirely determined by the number of items in the superbucket of size μ that contains it. As a consequence, there is never a reason to consider superbuckets of size smaller than μ . For that reason, instead of working with \mathbb{N}^m , where each entry corresponds to the load of a bucket, we will work with \mathbb{N}^{m_μ} , where each entry corresponds the load of a superbucket of size μ , divided by μ (so that an entry is the load of one bucket within the superbucket). To avoid creating confusion about whether a “bucket” or “items” refers to the original occupancy vectors in \mathbb{N}^m , or the ones just introduced in \mathbb{N}^{m_μ} , we reserve the term “bucket”, “superbucket”, and “item” to the former setting, so that the meaning of those terms is unchanged. When working in N^{m_μ} , we use balls-and-bins terminology: m_μ is the number of *bins*, and they are occupied by *balls*. Thus, each *bin* corresponds to a superbucket of size μ , and each *ball* corresponds to a list of μ items.

We now have all the tools to prove Lemma 19. Let $L_\cap = L \cap \text{merge}(L)$ be the lists common to L and $\text{merge}(L)$. Observe that the order lists are inserted by the algorithm does not matter, hence we are free to assume lists in L_\cap are inserted first.

Let $a \in \mathbb{N}^{m_\mu}$ denote an arbitrary load vector obtained after inserting the lists in L_\cap . Recall that lists in L_\cap are multiples of 2μ , so the load vector after inserting the lists is 2μ -flat. It follows that a may be written in the form $a = a' \diamond a'$ for some $a' \in \mathbb{N}^{m_\mu/2}$. Let us denote by E_a the event that the outcome of inserting L_\cap is equal to a .

We want to prove $\mathbb{E}[F(L)] \leq \mathbb{E}[F(\text{merge}(L))] + \text{negl}(\lambda)$. By Lemma 13, it suffices to prove the inequality when conditioning on E_a , for every possible a .

Given E_a , all that remains to do to compute $D(L)$ is to insert N_μ lists of length μ . In consequence, we have that $D(L)$ conditioned on E_a is equal to $a + B_{N_\mu, m_\mu}$. The lemma can then be established as follows.

In the computation, we multiply the output of `over` by μ , to reflect the fact that each ball in N^{m_μ} represents a list of μ items.

$$\mathbb{E}[F(L) : E_a]$$

$$\begin{aligned}
&= \mathbb{E} [\max(\mu \cdot \text{over}(D(L) - R, 0) : E_a)] \\
&= \mathbb{E} [\max(\mu \cdot \text{over}(a + X - R, 0)] \\
&\quad \text{where } X \leftrightarrow B_{N_\mu, m_\mu} \\
&\leq \mathbb{E} [\max(\mu \cdot \text{over}(a + X' - R, 0)] + N \text{negl}(\lambda) \\
&\quad \text{where } X' \leftrightarrow B'_{N_\mu, m_\mu, \phi(\mu)} \qquad \text{by Lemma 16}^{14}. \\
&\leq \mathbb{E} [\max(\mu \cdot \text{over}(a + Y - R, 0)] + \text{negl}(\lambda) \\
&\quad \text{where } Y \leftrightarrow B_{N_\mu + 2\phi(\mu), m_\mu}[\delta = 0] \qquad \text{by Lemma 17} \\
&= \mathbb{E} [\max(\mu \cdot \text{over}(a + Y^1 \diamond Y^2) - R, 0)] + \text{negl}(\lambda) \\
&\quad \text{where } Y^1, Y^2 \leftrightarrow B_{(N_\mu + 2\phi(\mu))/2, m_\mu/2} \qquad \text{by Lemma 15} \\
&\leq \mathbb{E} [\max(\mu \cdot \text{over}(a + Y^1 \diamond Y^1) - R, 0)] / 2 \\
&\quad + \mathbb{E} [\max(\mu \cdot \text{over}(a + Y^2 \diamond Y^2) - R, 0)] / 2 + \text{negl}(\lambda) \qquad \text{by Lemma 14} \\
&= \mathbb{E} [\max(\mu \cdot \text{over}(a + Y^1 \diamond Y^1) - R, 0)] + \text{negl}(\lambda) \\
&= \mathbb{E} [\max(\mu \cdot \text{over}((a' + Y^1) \diamond (a' + Y^1)) - R, 0)] + \text{negl}(\lambda) \\
&= \mathbb{E} [\max(2\mu \cdot \text{over}(a' + Y^1) - R, 0)] + \text{negl}(\lambda) \\
&= \mathbb{E} [F(\text{merge}(L)) : E_a] + \text{negl}(\lambda). \qquad \square
\end{aligned}$$

If we start from an arbitrary L , by applying Lemma 19 and computing $\text{merge}(L)$ as in the statement of the lemma, we strictly increase the minimum size of the list. Eventually, all lists have size $N/\log^d N$, while the total number of items remains $O(N)$ (Lemma 18). Hence, the analysis from Part 1 applies, and we are done.

F Applications of L2C

As mentioned in the introduction, the two-choice process is a staple in the resource allocation literature, with applications that range from job allocation to circuit routing. A survey may be found in [RMS01], which also presents some of the underlying analytical techniques. To further illustrate the range of applications, Azar, Broder, Karlin, Mitzenmacher and Upfal have recently received the 2020 ACM Paris Kenallakis Theory and Practice Award for the discovery and analysis of the two-choice process, highlighting in particular its “extensive applications to practice” [ABK⁺20]. An overview of the two-choice process was given in Section 2. In this section, we explain how our result on L2C fits within that broader context, beyond cryptographic applications.

Let us first recall the two-choice allocation process itself. Suppose that n balls are inserted into n bins as follows: for each ball in succession, two bins are sampled uniformly at random among all bins. The ball is then inserted into whichever bin currently contains fewer items. The central analytical result

¹⁴ The fact that the condition $d = \Omega(\sqrt{n} \log \lambda)$ from Lemma 16 is satisfied follows from Lemma 18.

regarding this process as that, once all balls have been inserted, the most loaded bin contains $\mathcal{O}(\log \log n)$ items with high probability.

Since typical applications include job allocation or memory management, it is natural to consider the *weighted* case, where each ball i is assigned some weight $w_i \in \mathbb{R}$ (or $w_i \in \mathbb{N}$). The weighted variant of the two-choice process inserts each ball into whichever of two uniformly random bins currently contains the least weight in total (rather than the least number of balls, in the unweighted case). One would hope that, for instance, if the weights lie in the real interval $[0, 1]$, then the load of the most loaded bin is $\mathcal{O}(\log \log W)$, where $W = \sum w_i$ is the total weight of the balls.

At STOC 2007, Talwar and Wieder analyzed the weighted two-choice process. They showed that when the weights of the balls are drawn *independently* from a fixed distribution of expectation 1, with some mild smoothness assumptions, the load of the most loaded bin is indeed $\mathcal{O}(\log \log n)$ (which is also $\mathcal{O}(\log \log W)$, if W is defined to be the expectation of the total weight). In fact, they show a much stronger result that the gap between the expected load of each bin and the load of the most loaded bin is bounded by $\mathcal{O}(\log \log m)$ with high probability, even when inserting an unbounded number of balls $m \gg n$ into n bins, inspired by a seminal paper by Berenbrink *et al.* showing the same result in the unweighted case [BCSV06]. A simpler proof of a variant of the result was later given in [TW14], again assuming weights drawn independently from a suitable distribution. To our knowledge, all existing analyses of the weighted two-choice process rely on a distributional assumption of that form¹⁵.

It seems quite natural to want to consider the case that there is no distributional assumption: the sequence of ball weights is instead an *arbitrary* sequence, subject only to an upper bound on the weight of an individual ball. Indeed, in many cases, the weights of the balls (corresponding *e.g.* to the cost of a job in a job allocation application, or the size of an object in a memory management application) may be determined by a client, and need not be drawn from a consistent distribution, and may not be independent of each other. To our knowledge, a distribution-free result of that form has only been shown for the one-choice process (which has a $\mathcal{O}(\log)$ overhead, rather than $\mathcal{O}(\log \log)$) by Berenbrink *et al.* [BFHM08], but the same article argues that their analysis technique cannot extend to the two-choice process.

In that context, our result on L2C (Theorem 1) shows that such a distribution-free result holds for the two-choice process, at the cost of slightly tweaking the process. In a nutshell, in L2C, instead of allocating a given ball to whichever of two uniformly sampled bins currently contains the *lowest weight* in total, we

¹⁵ A partial exception may exist: the result on the so-called $(1 + \beta)$ -choice process studied in [PTW10] is written in the same distributional form, but upon closer inspection, it appears that the core potential function argument could be written without the need of a distributional assumption, as long as the weights are bounded. However, as noted also in [TW14], this technique can only prove a logarithmic bound, rather than the $\mathcal{O}(\log \log)$ bound we are hoping for, so it is insufficient for our purpose.

allocate the ball to whichever of the two bins currently contains the *fewest balls within the same weight range*, where weight ranges are defined by partitioning the set of possible weights into $\mathcal{O}(\log \log n)$ sub-intervals. While this slightly alters the process, the computation determining which of the two bins to choose remains trivial. On the other hand, it allows for a distribution-free upper bound on the load of the most loaded bin.

A notable feature of L2C is that it is built essentially by superposing $\mathcal{O}(\log \log n)$ instances of a standard *unweighted* two-choice process. This makes it possible to reduce its analysis to results on the unweighted process. Although we only use this to show a simple upper bound on the most loaded bin, we conjecture that many deeper results known about the unweighted process (*e.g.* the “memoryless” properties from [BCSV06]), could be shown to extend to the weighted case for L2C in a similar manner.

We conjecture that a result of the same form as Theorem 1 would also hold for the standard weighted two-choice process considered in prior literature. However, such a result has remained elusive so far (some obstacles related to that question are discussed in [BFHM08]). We view this question as a compelling open problem.

Algorithm 7 Local Oblivious RAM (LocORAM)

LocORAM.Initialize_{ORAM}($1^\lambda, M$)

- 1: Parse M as $\{i, v_i\}_{i=1}^n$, where $|i, v_i| = \beta$
- 2: Let $n_1 = n^{\frac{1}{c}}$ and $n_i = n^{\frac{1}{c}} + n^{\frac{i-1}{c}}$ for $i \in [2, c]$
- 3: Let A_i be an empty array of size n_i for $i \in [1, c]$
- 4: Let $\pi_i : [1, n_i] \mapsto [1, n_i]$ be pseudorandom permutation for $i \in [2, c]$
- 5: **for all** $i \in [1, n]$ **do**
- 6: Store (i, v_i) at locations $\pi_c[i]$ in A_c
- 7: Encrypt $A_i^{\text{enc}} \leftarrow \text{Enc}_{K_{\text{Enc}}}(A_i)$ for $i \in [1, c]$
- 8: Let R_i be an empty set (of maximal size $n^{i/c}$ blocks for $i \in [2, c-1]$) and $R_c = M$
- 9: Set $\text{cnt}_i \leftarrow 0$ for $i \in [2, c]$
- 10: Let T_i be an empty hash table of size n for $i \in [2, c-1]$
- 11: Encrypt $R_i^{\text{enc}} \leftarrow \text{Enc}_{K_{\text{Enc}}}(R_i)$ and $T_i^{\text{enc}} \leftarrow \text{Enc}_{K_{\text{Enc}}}(T_i)$ for $i \in [2, c-1]$
- 12: Set $\text{st} = (\{\pi_i\}_{i=2}^c, \{\text{cnt}_i\}_{i=2}^c, K_{\text{Enc}})$
- 13: Set $\text{EM} = (\{A_i^{\text{enc}}\}_{i=1}^c, \{R_i^{\text{enc}}\}_{i=2}^{c-1}, \{T_i^{\text{enc}}\}_{i=2}^{c-1})$
- 14: **return** st, EM

LocORAM.Access_{ORAM}($\text{st}, k; \text{EM}$)

Client:

- 1: Retrieve $(A_1^{\text{enc}}, \{T_i^{\text{enc}}\}_{i=2}^{c-1})$ from the server and decrypt to $(A_1, \{T_i\}_{i=2}^{c-1})$
- 2: Set $\text{fnd} \leftarrow \text{false}$ and $\text{cnt}_i \leftarrow \text{cnt}_i + 1$ for $i \in [2, c]$
- 3: **if** $(k, v_k) \in A_1$ **then**
- 4: $\text{fnd} \leftarrow \text{true}$
- 5: **for all** $i \in [2, c-1]$ **do**
- 6: **if** fnd or $T_i[k] = 0$ **then**
- 7: $\text{index}_i \leftarrow \pi_i(n^{i/c} + \text{cnt}_i)$
- 8: **else**
- 9: $\text{index}_i \leftarrow \pi_i(T_i[k])$
- 10: $\text{fnd} = \text{true}$
- 11: $T_i[k] \leftarrow \text{cnt}_{i+1}$
- 12: $r_i \leftarrow \text{Enc}_{K_{\text{Enc}}}(\text{cnt}_{i+1}, v_k)$
- 13: **if** fnd **then**
- 14: $\text{index}_c \leftarrow \pi_c[n + \text{cnt}_c]$
- 15: **else**
- 16: $\text{index}_c \leftarrow \pi_c[k]$
- 17: $A_1[\text{cnt}_2] \leftarrow (k, v_k)$
- 18: **send** $\{\text{index}_i\}_{i=2}^c$ and $\{r_i\}_{i=2}^{c-1}$

Server:

- 1: Set $R_i \leftarrow R_i \cup r_i$ for $i \in [2, c-1]$
- 2: **send** $\{A_i^{\text{enc}}[\text{index}_i]\}_{i=2}^c$

Client:

- 1: Retrieve block (k, v_k) from either A_1 or (decrypted) $A_i[\text{index}_i]$ for some $i \in [2, c]$
- 2: Choose $i \in (c, \dots, 2)$ maximal such that $\text{cnt}_i > n^{\frac{i-1}{c}}$
- 3: **if** i exists **then**
- 4: Let π_j be a new pseudorandom permutation for $j \in [1, i]$
- 5: Set $\text{cnt}_j \leftarrow 0$ for $j \in [2, i]$
- 6: Server updates A_i with the result of $\text{ObSort}(\pi_i, n_i, n^{1/c} \log^2 n; R_i)$
- 7: Empty A_1 and T_j , such as A_j and R_j on the server, for $j \in [2, i-1]$
- 8: Store updated client state
- 9: **send** reencrypted $(A_1^{\text{enc}}, \{T_i^{\text{enc}}\}_{i=2}^{c-1})$ to server

Server:

- 1: Update the encrypted memory EM accordingly
-