



HAL
open science

Better Space-Time-Robustness Trade-Offs for Set Reconciliation

Belazzougui Djamal, Gregory Kucherov, Stefan Walzer

► **To cite this version:**

Belazzougui Djamal, Gregory Kucherov, Stefan Walzer. Better Space-Time-Robustness Trade-Offs for Set Reconciliation. 51st International Colloquium on Automata, Languages, and Programming (ICALP 2024), Jul 2024, Tallinn, Estonia. pp.20:1-20:19, 10.4230/LIPIcs.ICALP.2024.20. hal-04799765

HAL Id: hal-04799765

<https://cnrs.hal.science/hal-04799765v1>

Submitted on 23 Nov 2024

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.



Distributed under a Creative Commons Attribution 4.0 International License

Better Space-Time-Robustness Trade-Offs for Set Reconciliation

Djamal Belazzougui ✉ 

CAPA, DTISI, Centre de Recherche sur l'Information Scientifique et Technique, Algiers, Algeria

Gregory Kucherov ✉ 

LIGM, CNRS, Université Gustave Eiffel, Marne-la-Vallée, France

Stefan Walzer ✉ 

Karlsruhe Institute of Technology, Germany

Abstract

We consider the problem of reconstructing the symmetric difference between similar sets from their representations (sketches) of size linear in the number of differences. Exact solutions to this problem are based on error-correcting coding techniques and suffer from a large decoding time. Existing probabilistic solutions based on Invertible Bloom Lookup Tables (IBLTs) are time-efficient but offer insufficient success guarantees for many applications. Here we propose a tunable trade-off between the two approaches combining the efficiency of IBLTs with exponentially decreasing failure probability. The proof relies on a refined analysis of IBLTs proposed in (Bæk Tejs Houen et al. SOSA 2023) which has an independent interest. We also propose a modification of our algorithm that enables telling apart the elements of each set in the symmetric difference.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms; Theory of computation → Randomness, geometry and discrete structures; Theory of computation → Sketching and sampling; Theory of computation → Error-correcting codes

Keywords and phrases data structures, hashing, set reconciliation, invertible Bloom lookup tables, random hypergraphs, BCH codes

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.20

Category Track A: Algorithms, Complexity and Games

Related Version *Preliminary version:* <https://arxiv.org/abs/2404.09607>

1 Introduction

The problem of *set reconciliation* (or *database reconciliation* [12]) lies at the intersection of several lines of algorithmic research. One intuitive way to define it is in terms of communication protocols: assume that two parties Alice and Bob hold similar sets S and T , respectively, that they seek to reconcile (i.e. identify differences) using low communication overhead. Rather than exchanging representations of entire sets, they exchange their small *sketches* holding enough information to identify differences of S and T provided that their number does not exceed a pre-defined parameter D .

In this paper, we study the version of the problem where the differences $S \Delta T := (S \setminus T) \cup (T \setminus S)$ should be recovered from the sketches of S and T rather than from the sketch of one of them and the entire other set. This opens a way to the setting, relevant to some applications, when only sketches of sets are stored in a database, and for each pair of sets, items proper to one of them can be retrieved from the corresponding sketches. As an example, consider a vast database of highly similar genomic sequences, such as those of SARS-Cov-2 genomes¹, that are stored in the form of sketches, rather than raw sequences, still supporting retrieval of differences between any two of them. Beyond this example, set reconciliation

¹ More than 16M SARS-Cov-2 are available in GISAID database at the time of writing, which typically differ one from another by just a few characters, over about 30,000 characters of length.



problem occurs in many other scenarios in distributed systems, where information needs to be synchronized between computational units. Those include blockchain systems, P2P systems, ad-hoc networks, synchronizing information across devices or data centers, and others.

The set reconciliation problem is also raised in the streaming framework. It is a common and natural requirement that the sketch can be efficiently updated when a new key is added to or deleted from the set. For example, [6] considers the *straggler identification* problem defined on a stream of insertions and deletions of keys modeling a stream of people entering and leaving a building. In this model, the sketch summarizes information about people currently in the building and if their number does not exceed D , is capable of listing them.

1.1 Prior work

Solutions to set reconciliation can be categorized into probabilistic and exact ones (although there exists an interplay between these two classes of algorithms) [25]. Probabilistic algorithms are based on Invertible Bloom Lookup Tables (IBLTs), also known as Invertible Bloom Filters, based on random hash functions. Originally proposed in [6] where they were applied to the problem of *subset reconciliation* (set reconciliation assuming $S \subseteq T$), they (in a slightly different version) were more systematically studied in [12]. Earlier a related data structure, called *k-set data structure*, was proposed in [10]. Set reconciliation by computing a difference of IBLTs (called *difference digest*) was studied in [7]. Work [21] further applies IBLTs to *multi-party* set reconciliation. The idea of “subtracting” IBLTs has also been applied in [9], where the emphasis is to store in an IBLT carefully defined hash values obtained using cryptography techniques. Very recent work [1] studies a simplified version of IBLTs that also applies to set reconciliation. We rely on this construction in the present paper.

An exact solution of the set reconciliation problem uses algebraic techniques, in particular *error-correcting codes* [7, 3, 4]. Work [20] proposes an exact solution based on characteristic polynomials to both subset and set reconciliation problems using $(D + 1) \log U$ bits of space (i.e. essentially the space needed to store D differences). [11] studies the straggler identification problem with multiplicities and proposes an $\mathcal{O}(D \log(Um))$ -bit solution based on polynomials over finite field, where each key occurs at most m times. A somewhat similar solution to subset reconciliation was proposed in [6], based on Newton polynomials. [3] surveys code-theoretic techniques for space recovery many of which apply to set reconciliation as well. Set reconciliation with BCH codes has been implemented in Minsketch software [19], leveraging an efficient syndrome decoding algorithm for BCH codes [5]. [4] proposes a method for sparse recovery based on expander codes; the construction of [4] can be interpreted as a table somewhat similar to IBLT but with a different decoding mechanism. Recent papers [23, 2] apply algebraic techniques to construct specific “hash functions” that guarantee successful decoding for IBLTs. The general downside of exact solutions is a larger decoding time, typically growing at least quadratically in D and relying on finite field arithmetic. Note also that the relationship between codes and set reconciliation is two-way: [22] proposes a construction of codes based on set reconciliation with IBLTs.

1.2 Our contribution: overview

The IBLT solution to set reconciliation is very efficient but provides a poor success guarantee. For example, if we have to compute similarity joins by performing all pairwise reconciliations of the objects in our database, then a significant fraction of comparisons may fail. On the other hand, techniques based on error-correcting codes are exact but have a high cost of

decoding. Here we propose a family of solutions to set reconciliation offering a trade-off between these two approaches, that is having significantly smaller error rate compared to IBLT at the price of an asymptotically vanishing increase of space and time bounds.

Our solution is based on the space-efficient IBLT from [1] complemented by an additional *stash* data structure supporting recovery in case of failure of the main IBLT. The analysis of [1] provides no guarantee in the case when the IBLT decoding fails. Therefore, here we enhance the analysis of [1] and prove a probability bound on the number of missing and extraneous elements reported by the data structure, in the event that the correct decoding fails. This is a key argument of our construction and its proof constitutes the main technical contribution of this paper.

The algorithm we obtain is parametrized and provides a tunable trade-off between the failure probability and space and time consumption. In the table below, we provide an overview of how our solution relates to the existing approaches surveyed in Section 1.1.

■ **Table 1** Comparison of main techniques of set reconciliation to the results of the present work. r is a parameter assumed to verify $r = \min(\mathcal{O}(D/\log^2 U), \mathcal{O}(\log U))$. For our algorithm, decoding time given is expected.

Method	Sketch size (bits)	Insertion time	Decoding time	Failure probability
IBLT [1]	$(c_3 + \epsilon)D \log U$	$\mathcal{O}(1)$	$\mathcal{O}(D)$	$\mathcal{O}(D^{-1})$
IBLT with t -bit hashsum field	$(c_3 + \epsilon)D(\log U + t)$	$\mathcal{O}(1 + t/\log U)$	$\mathcal{O}(D(1 + t/\log U))$	$\min(\mathcal{O}(D^{-1}), \mathcal{O}(D/2^t))$
BCH	$D \log U$	$\mathcal{O}(D \log U)$	$\mathcal{O}(D^2 \log^2 U)$	0
expander code [4]	$\mathcal{O}(D \log^2 U)$	$\mathcal{O}(\log U)$	$\mathcal{O}(D \log U)$	0
this paper	$(c_3 + \epsilon)D \log U + r \log U$	$\mathcal{O}(r \log U)$	$\mathcal{O}(D)$	$2^{-\Omega(r)}$

The paper is organized as follows. In Section 2, we introduce our version of IBLT, which is a slightly modified set sketch from [1]. This data structure is very space-efficient, however its compactness has a price: the decoding process can “go wrong” triggering undesirable *anomalous* steps [1]. In Section 3.2, we provide a refined analysis of anomalies and prove the following fundamental property: if the sketch stores a set S , the decoding produces a set S_{dec} with $|S \Delta S_{dec}| \leq r$, with a failure probability of $2^{-\Omega(r)}$. Section 4 introduces a solution to set reconciliation based on BCH error-correcting codes. Relying on that result, in Section 5, we extend the IBLT with a backup data structure (stash) to obtain an efficient solution to set reconciliation. The algorithm resorts to the stash when the decoding with the main IBLT is not completed, leaving out a small number of keys. The stash is defined using error-correcting codes. In Section 6, we discuss how to extend our algorithm in order to identify the original set of the keys of the set difference.

2 Set reconciliation with Invertible Bloom Lookup Tables

Definition of IBLT. An Invertible Bloom Lookup Table (IBLT) is an array $A[1 : n]$ equipped with k random hash functions $h_1, \dots, h_k : \mathcal{U} \rightarrow [n]$, from the key universe \mathcal{U} to $[n] = \{1, \dots, n\}$. We assume h_i 's to be fully random and denote $h(x) = \{h_1(x), \dots, h_k(x)\}$.

Several variants of IBLTs have been considered, which differ in how entries $A[i]$ are configured. In all of them, $A[i]$ contains a keysum field that holds either bitwise XOR [12, 7, 21, 1] or arithmetic sum [6, 12] of all current keys hashed to i . XOR provides a more elegant and space-efficient option whereas arithmetic sum becomes necessary when multi-sets are considered. In this work, the keysum field is defined with XOR and is set to contain $\log U$ bits ($U = |\mathcal{U}|$).

```

Algorithm initialise:
┌  $A[1 : n] = (0, \dots, 0)$ 

Algorithm toggle( $x$ ):
┌   for  $i \in h(x)$  do
└   ┌  $A[i] \leftarrow A[i] \oplus x$ 

Algorithm merge( $A[1 : n], A'[1 : n]$ ):
┌   for  $i \in [n]$  do
└   ┌  $B[i] \leftarrow A[i] \oplus A'[i]$ 
    └ return  $B$ 

Algorithm looksPure( $i \in [n]$ ):
┌ return  $A[i] \neq 0 \wedge i \in h(A[i])$ 

Algorithm decode:
┌  $S_{\text{dec}} \leftarrow \emptyset$ 
  ┌  $Q \leftarrow \{i \in [n] \mid \text{looksPure}(i)\}$ 
    ┌  $(t, t_{\text{max}}) \leftarrow (0, 2n)$  // time limit  $2n$ 
      ┌ while  $Q \neq \emptyset$  do
        ┌  $Q_{\text{next}} \leftarrow \emptyset$ 
          ┌ for  $i \in Q$  if  $\text{looksPure}(i)$  do
            ┌  $x \leftarrow A[i]$  // detected key  $x$ 
              ┌ toggle( $x$ )
                ┌  $S_{\text{dec}} \leftarrow S_{\text{dec}} \Delta \{x\}$ 
                  ┌  $Q_{\text{next}} \leftarrow Q_{\text{next}} \cup \{i \in h(x) \mid \text{looksPure}(i)\}$ 
                    ┌  $t \leftarrow t + 1$ 
                      ┌ if  $t \geq t_{\text{max}}$  then
                        └ return  $S_{\text{dec}}$ 
                    └  $Q \leftarrow Q_{\text{next}}$ 
                  └ return  $S_{\text{dec}}$ 
                └  $Q \leftarrow Q_{\text{next}}$ 
              └  $Q \leftarrow Q_{\text{next}}$ 
            └  $Q \leftarrow Q_{\text{next}}$ 
          └  $Q \leftarrow Q_{\text{next}}$ 
        └  $Q \leftarrow Q_{\text{next}}$ 
      └  $Q \leftarrow Q_{\text{next}}$ 
    └  $Q \leftarrow Q_{\text{next}}$ 
  └  $Q \leftarrow Q_{\text{next}}$ 
└  $Q \leftarrow Q_{\text{next}}$ 

```

■ **Figure 1** IBLT implementation from [1] with added time limit in decode.

Besides the keysum, $A[i]$ can include a hashsum field that holds a hash sum of the keys in $A[i]$ and/or a counter field that tracks the current number of keys in $A[i]$. These are used for enforcing proper functioning and integrity of the data structure at the price of using additional space. In this work, we use the most compact IBLT configuration with three hash functions ($k = 3$) and $A[i]$ storing the keysum field only. This variant was introduced and analysed in recent paper [1] that we rely on in this work. Later in Section 6 we consider an extension introducing a restricted counter field.

Our IBLT implementation is defined in Figure 1. Initially, all entries $A[i]$ are set to zero with `initialise`. Both insertion and deletion of a key x is done by performing `toggle(x)`.

Decoding the keys stored in an IBLT resembles the *peeling* process of the k -hypergraph, where n IBLT entries correspond to hypergraph vertices and each stored key corresponds to a hyperedge defined as the set of entries the key is hashed to. Peeling a hypergraph iterates the following operation: for any vertex of degree 1, delete the incident hyperedge. If the peeling results in the empty graph with no hyperedges, the input graph is called *peelable*; otherwise the process yields the largest subhypergraph with every vertex of degree at least 2, which is called the *2-core* (hereafter simply *core*). It is known that peelability of random k -hypergraphs ($k \geq 3$) with n vertices and m hyperedges undergoes a phase transition: if $n > c_k m$, then the hypergraph is peelable with high probability (hereafter, *whp*), whereas if $n < c_k m$, the hypergraph is not peelable whp. Here constant c_k is the peelability threshold, in particular $c_3 = 1.22179\dots$ is the smallest of c_k for $k \geq 3$. The following results is shown in [12].

► **Theorem 1.** *Whenever $n > c_k m$, a random k -hypergraph is peelable except with probability $\mathcal{O}(1/n^{k-2})$.*

Set reconciliation using difference IBLT. Consider two sets S and T stored in IBLTs $A_S[1 : n]$ and $A_T[1 : n]$ respectively, using the same hash functions. If S and T have a bounded symmetric difference, that is $|S \Delta T| \leq D$, and $n > c_k D$, then the keys of $S \Delta T$ can

be recovered whp from the *difference IBLT* $A_{ST}[1 : n] = \text{merge}(A_S, A_T)$. It is immediate to see that common keys of S and T are canceled out and A_{ST} stores exactly the keys $S \Delta T$. Thus, the difference IBLT can solve set reconciliation except with probability $\mathcal{O}(1/D)$.

Note that A_{ST} does not allow distinguishing keys of $S \setminus T$ and $T \setminus S$; below in Section 6 we will extend the construction to make this possible. Note also that for the reconciliation to be possible, the size n of IBLTs only depends on the size of the symmetric difference and does not depend on the sizes of S and T .

3 Improved Guarantees for IBLTs

Consider the IBLT implementation as given in Figure 1. It uses only the keysum field in its buckets and was studied, modulo slight changes stated below, in [1] under the name “simple set sketch”. In this section we enhance the analysis given in [1] (reproduced as (i) in Theorem 2 below) by also probabilistically bounding the magnitude of the error of decode in its failure cases (given as (ii)). Note that we restrict our attention to three hash functions ($k = 3$) as this yields the best threshold value.²

► **Theorem 2.** *Let $\varepsilon > 0$ and $n > (c_3 + \varepsilon)|S|$. Let S_{dec} be the set returned by decode.*

- (i) $\Pr[S_{\text{dec}} = S] = 1 - \tilde{\mathcal{O}}(1/n)$.
- (ii) *For any $r = o(n)$ we have $\Pr[|S \Delta S_{\text{dec}}| > r] = 2^{-\Omega(r)}$.*

Changes made to decode. Our version of decode differs from that in [1] in two minor ways. First, we have introduced a *time limit* of $2n$ to the number of peeling steps that are performed and return the current state of S_{dec} when the time limit is reached³. The result from [1] still applies with this change, since the authors proved that decode terminates in $n + \text{poly} \log(n)$ peeling steps with $S_{\text{dec}} = S$ with probability $1 - \mathcal{O}(1/n)$. Second, we insist that Q and Q_{next} are implemented as set data structures, e.g. using hash tables, whereas [1] pointed out that Q and Q_{next} can be implemented as multiset data structures, e.g. using FIFO queues.⁴

It should be clear that our version of decode can be implemented such that it has running time $\mathcal{O}(n)$, both in expectation and with high probability.

3.1 The issue of anomalies

Decode checks if the i th *bucket* of A contains a single key using the function `looksPure(i)` that checks if $h_j(A[i]) = i$ for some $j \in \{1, 2, 3\}$. Importantly, `looksPure(i)` may produce a false positive result when $A[i] = x_1 \oplus x_2 \oplus \dots \oplus x_{\ell-1}$ is the sum of several keys hashing to i and $y = A[i]$ also happens to hash to i , i.e. $h_\sigma(y) = i$ for some $\sigma \in \{1, 2, 3\}$. Such an occurrence $\{x_1, \dots, x_{\ell-1}, y\}$ is called an *anomaly* (of size ℓ) in [1] and causes the key y to be toggled and thus inserted into the data structure. Anomalies corrupt the decoding process and may cause decode to fail or even to run in an infinite loop if no time limit is set [1].

² In [1] any $k \geq 3$ was considered and our analysis here could likewise be extended.

³ In practice a dynamic time limit could be considered. For instance: Abort if a round of decode only toggles keys that have already been toggled in a previous round.

⁴ At no point did the analysis from [1] actually *rely* on Q or Q_{next} being multisets. The motivation was merely that FIFO queues are the more lightweight data structure. We revert this simplification here because an argument from [1] showing that the impact of duplicated elements is negligible no longer applies. In fact, with a FIFO implementation of Q and Q_{next} we could see 2^R copies of the same element in Q in the R th iteration of the while loop with small but non-negligible probability $n^{-\mathcal{O}(1)}$.

Dealing with anomalies is a challenging issue. In [1], it was proved that certain kinds of anomalies and certain interactions of several anomalies do not occur except with probability $\tilde{O}(1/n)$. Since $\mathcal{O}(1/n)$ is the probability for the existence of a non-empty core (see Theorem 1 for $k = 3$), which also causes decoding to fail, anomalies do not increase the failure probability significantly.

In this section we will go further, proving that *even in the unlikely cases where decoding produces an incorrect result*, e.g. due to problematic anomalies, the magnitude of the error is likely to be small.

3.2 Proof Outline

We quantify the effect of anomalies by considering two sets. First, the set F of all *foreign keys*, which are keys from $\mathcal{U} \setminus S$ that are added to the sketch at least once during the decoding process. Second, the set N of all centres of certain anomalies we call *internal anomalies* (defined in Section 3.4). Our proof of Theorem 2 rests on the following arguments.

- In Section 3.3 we bound the size of F . Intuitively, the heuristic `looksPure` is used $\mathcal{O}(n)$ times and fails each time with probability $\mathcal{O}(1/n)$. Each failure can cause one foreign key to be added, suggesting $\mathbb{E}[|F|] = \mathcal{O}(1)$. The precise argument, including a concentration bound for $|F|$, is more subtle since failures do not happen independently.
- In Section 3.4 we import a tail bound on $|N|$ from [1].
- In Section 3.5 we reframe the decoding process as the well-understood *peeling process* on a random hypergraph. Since peeling follows a local rule, the corruption due to anomalies spreads locally as well. This already implies that the majority of keys are correctly decoded if anomalies are rare.
- In Section 3.6 we prove a claim regarding the maximum hyperedge density of any subhypergraph of a given random hypergraph. Its intuitive role is to show that the size of the queue Q in late rounds of decode would always be, in the absence of anomalies, at least linear in the number of remaining keys. This implies that the anomalies only have the potential to stall progress in a significant way if they affect a number of buckets that is linear in the number of remaining keys.
- In Section 3.7 we stitch everything together: The effect of anomalies as quantified by $|F|$ and $|N|$ is probabilistically bounded. Peeling is likely to proceed largely unhindered in its early rounds. In its later rounds it proceeds until the number of remaining keys is linear in $|F| + |N|$.

3.3 Bounding Foreign Keys

Let $y \in \mathcal{U} \setminus S$ be any foreign key. In order for y to be added to the IBLT, y must occur as the sum of several keys in one of the buckets $h_1(y), h_2(y), h_3(y)$. Since $h_1(y), h_2(y), h_3(y)$ are uniformly random in $[n]$ and independent of the initial state of the data structure, decode would have to effectively “guess” a hash of y . We argue that decode is, in this sense, a player in a corresponding guessing game that, even if played optimally, rarely allows a player to guess a lot of hashes correctly.

A guessing game. Let $(X_{j,\sigma})_{j \in \mathbb{N}, \sigma \in [3]}$ be a family of i.i.d. random variables with uniformly random values in $[n]$. Consider a game in which the player does not know the random variables. She may issue queries of the form $(i, j) \in [n] \times \mathbb{N}$ and is told whether or not $i \in \{X_{j,\sigma} \mid \sigma \in [3]\}$. If she gets a positive answer, we say she has *solved* group j and she is told the values $X_{j,1}, X_{j,2}$ and $X_{j,3}$.

► **Lemma 3.** *Assume a player is given $6n$ queries. She solves more than r groups with probability $\mathcal{O}(2^{-r})$.*

Proof. Consider a query $(i, j) \in [n] \times \mathbb{N}$ that is the k th query of the form (\cdot, j) where j is an unsolved group. We distinguish two cases.

Case 1: $k \geq n/2$. We generously grant group j as solved. This affects at most $6n/(n/2) = 12$ groups.

Case 2: $k \leq n/2$. Up to $k - 1$ values are already excluded for the variables $X_{j,1}, X_{j,2}$ and $X_{j,3}$. All other values are equally likely. The probability of solving group j with this query is at most $1 - (1 - \frac{1}{n-k+1})^3 < 1 - (1 - \frac{2}{n})^3 \leq \frac{6}{n}$.

Let Y be the number of successes due to Case 2. A simple coupling yields a random variable Z with $Y \leq Z$ and $Z \sim \text{Bin}(6n, \frac{6}{n})$. We can think of Z as the sum of the outcomes of $6n$ Bernoulli trials. In order for Z to attain a value of at least $\ell = r - 12$ there must exist a set $T \subseteq [6n]$ of size ℓ such that the corresponding Bernoulli trials are all successes. A union bound over all choices of T yields:

$$\begin{aligned} \Pr[Y \geq \ell] &\leq \Pr[Z \geq \ell] \leq \binom{6n}{\ell} \left(\frac{6}{n}\right)^\ell \leq \left(\frac{6ne}{\ell}\right)^\ell \left(\frac{6}{n}\right)^\ell \\ &\leq \left(\frac{36e}{\ell}\right)^\ell = \left(\frac{36e}{r-12}\right)^{r-12} \leq 2^{-r+12}, \end{aligned}$$

where the last step assumes $r \geq 72e + 12$.

Overall, if $r \geq 72e + 12$ then the number of successes is bounded by r , except with probability $\mathcal{O}(2^{-r})$. ◀

How this relates to the foreign keys. Let $F \subseteq \mathcal{U} \setminus S$ be the set of all foreign keys that are added to S_{dec} at least once during the execution of `decode`.

► **Lemma 4.** *For any $r \in \mathbb{N}$ we have $\Pr[|F| > r] = \mathcal{O}(2^{-r})$.*

Proof. Before decoding begins, the family $(h_\sigma(y))_{y \in \mathcal{U} \setminus S, \sigma \in [3]}$ of hashes of foreign keys are independent and uniformly random in $[n]$ just as is required in the guessing game above. We interpret `decode` as a player. Whenever `decode` executes the `looksPure` operation for a bucket i containing $A[i] = y \in \mathcal{U} \setminus S$, then we interpret this as the player issuing the query (y, i) . He learns whether $i \in \{h_1(y), h_2(y), h_3(y)\}$ and, on a positive answer, he also learns the values $h_1(y), h_2(y), h_3(y)$. This is sufficient to continue the execution of `decode`, in particular `decode` amounts to a valid player respecting the access limitation to the relevant random variables.

The crucial observation is that any $y \in F$ is solved by `decode`. Indeed, if $y \in F$ then y was added to S_{dec} at some point during the execution of `decode` and this addition was immediately preceded by a corresponding successful `looksPure` check.

The number q of queries issued by `decode` is by definition the number of calls to `looksPure`. Every time a bucket i is considered to be added to Q or Q_{next} there is one such call, and, if i is added, there can be another call when i is removed from Q . Due to the time limit we have $q \leq 2(n + t_{\text{max}}) \leq 6n$. Therefore Lemma 3 implies our claim. ◀

3.4 Bounding Centres of Internal Anomalies

It is worthwhile to clarify the notion of anomalies by giving a precise definition of the set \mathcal{A} of all anomalies.

$$\mathcal{A} = \left\{ \emptyset \neq A \subseteq \mathcal{U} \mid \bigoplus_{x \in A} x = 0, \bigcap_{x \in A} h(x) \neq \emptyset \right\}$$

If $A = \{x_1, \dots, x_\ell\}$ is an anomaly then the presence of any $\ell - 1$ keys from A can be mistaken as the presence of the missing ℓ th key (because of $x_1 \oplus \dots \oplus x_{\ell-1} = x_\ell$). A value $i \in \bigcap_{x \in A} h(x)$ is a centre of the anomaly A (in principle a single anomaly could have multiple centres). The set \mathcal{A} makes no mention of S and contains many anomalies that are unlikely to become relevant during decoding. Immediately threatening is, however, the set of *native anomalies* \mathcal{A}_{nat} define as

$$\mathcal{A}_{\text{nat}} = \{A \in \mathcal{A} \mid |A \setminus S| \leq 1\}.$$

For a native anomaly A , all but at most 1 key from A are present in the beginning, enough to cause trouble right away. In what follows, a subset of the native anomalies, namely the *internal* anomalies \mathcal{A}_{int} and the set N of its centres play a role. They are defined as

$$\mathcal{A}_{\text{int}} = \{A \in \mathcal{A} \mid A \subseteq S\} \text{ and } N = \bigcup_{A \in \mathcal{A}_{\text{int}}} \bigcap_{x \in A} h(x).$$

► **Lemma 5.** *For any $r \in \mathbb{N}$ we have $\Pr[|N| > r] = \mathcal{O}(2^{-r})$.*

Proof. The authors of [1] show in their Lemma 2.2 that the set $N_{\text{nat}} = \bigcup_{A \in \mathcal{A}_{\text{nat}}} \bigcap_{x \in A} h(x)$ of centres of native anomalies satisfies $\Pr[N' \geq r] \leq \left(\frac{3e^4}{r}\right)^r$. For $r \geq 6e^4$ this is at most 2^{-r} . Since $N_{\text{nat}} \supseteq N$ the claim follows. ◀

3.5 Early Rounds of the Peeling Process

We will use a result by Molloy who analysed cores of random hypergraphs [24]. The hypergraph H arising in our setting has vertex set is $[n]$ and the m hyperedges $\{\{h_\sigma(x) \mid \sigma \in [k]\} \mid x \in S\}$. By the *parallel peeling* process we mean an algorithm on a hypergraph. In each round, all vertices of degree 0 or 1 are determined (simultaneously) and then all of these vertices are removed, including all incident hyperedges. When no vertex of degree 0 or 1 remains, the core of the hypergraph is reached, which may be empty. Note that decode implements this peeling process, except that anomalies may occur.

Early Peeling Without Anomalies. A lemma by Molloy is the following.

► **Lemma 6** ([24, Lemma 3]). *For any $\varepsilon, \delta > 0$, there exists $R \in \mathbb{N}$ such that after peeling a k -uniform hypergraph with $\frac{n}{m} > c_3 + \varepsilon$ for R rounds then at most δn vertices remain in expectation.⁵*

In the following we assume that constants $\varepsilon, \delta > 0$ are given and $R \in \mathbb{N}$ is the corresponding constant from Lemma 6.

By the R -neighbourhood of a vertex v or hyperedge e we mean the subhypergraph of H induced by all vertices that can be reached from v or e by traversing at most R hyperedges.

It is a well-known fact (and a consequence of the Poisson limit theorem) that the degree distribution of H converges to a Poisson distribution with parameter $3/c_3$ for $n \rightarrow \infty$. More generally, the distribution of the R -neighbourhoods of the vertices in H converges (see e.g. [15, 16]). This implies that we can choose a constant $D = D(\varepsilon, \delta, R)$ large enough such the probability that a given vertex has a vertex of degree at least D in its R -neighbourhood is at most δ .

⁵ Molloy's Lemma is slightly stronger in that it claims that at most δn vertices remain with probability $1 - o(1)$. Since we need a much stronger concentration bound, we will redo a corresponding step ourselves.

In this context call a vertex v *good* if

- (a) v is removed by the R -round peeling process
- (b) the R -neighbourhood of v does not contain a vertex of degree at least D .

Since the first condition is met by at least $(1 - \delta)n$ vertices in expectation by Lemma 6 and the second condition by at least $(1 - \delta)n$ vertices in expectation by choice of D , the set G of good vertices has expectation at least $(1 - 2\delta)n$.

We will now use the bounded difference inequality by McDiarmid [18].

► **Lemma 7** (McDiarmid's inequality [18]). *Let \mathcal{X} be some set, $d \in \mathbb{N}$ and $f : \mathcal{X}^n \rightarrow \mathbb{R}$ a function such that changing any one argument of f can affect the function value by at most d . Let X_1, \dots, X_n be i.i.d. random variables with values in \mathcal{X} . Then*

$$\Pr[|f(X_1, \dots, X_n) - \mathbb{E}[f(X_1, \dots, X_n)]| > \varepsilon] \leq \exp(-\varepsilon^2/(nd^2)).$$

This gives us a concentration bound on the number of good vertices.

► **Lemma 8.** *In the above context there exists a constant $\gamma = \gamma(\varepsilon, \delta, R, D) > 0$ such that*

$$\Pr[|G| < (1 - 3\delta)n] \leq \exp(-\gamma n).$$

Proof. Note that whether or not a vertex is good is a *local* property in the sense that it only depends on its R -neighbourhood. This gives us a bounded difference property of $|G|$ as follows.

The set G (and therefore $|G|$) is a function of the $3m$ hash values defining the m hyperedges. Assume we change a single hash value from v to $v' \neq v$. A previously non-good vertex might become good only if it was in the vicinity of v and a previously good vertex might become non-good only if it ends up in the vicinity of v' . More precisely, a vertex w might be affected if prior to the change, v or v' (or both) were reachable from w via a path of length at most R that only traverses vertices of degree at most D (vertices reachable through vertices of higher degree are non-good anyway). To bound the number of such paths, note that in every step we can first choose among $\leq D$ incident hyperedges and then among $3 - 1$ endpoints of the chosen hyperedge, or we can choose to end the path at the current vertex. In particular, changing a single incidence affects the number of good vertices by at most $2(2D + 1)^R$. We can therefore apply McDiarmid's inequality (Lemma 7) to conclude that

$$\begin{aligned} \Pr[n - |G| > 3\delta n] &\leq \Pr[(n - |G|) - \mathbb{E}[n - |G|] > \delta n] \\ &\leq \exp(-(\delta n)^2/(n \cdot (4(2D + 1)^{2R}))) = \exp(-\gamma n) \end{aligned}$$

when defining $\gamma = \delta^2/(4(2D + 1)^{2R})$. ◀

Early Peeling With Anomalies. We now get back to the analysis of decode. Let S_{dec}^R be the state of S_{dec} after R rounds of the while-loop are executed (ignoring the time limit). Let $S^R := S_{\text{dec}}^R \triangle S$ be the set of elements stored in the IBLT at that time and $I^R := \bigcup_{x \in S^R} h(x)$ the set of buckets touched by S^R .

► **Lemma 9.** *For any $\varepsilon > 0$ and $\delta > 0$ there exist constants R and β such that*

$$\Pr[|I^R| > 7\delta n] \leq \mathcal{O}(2^{-\beta n}).$$

Proof. Given ε and δ , let $R = R(\varepsilon, \delta)$, $D = D(\varepsilon, \delta, R)$ and $\gamma = \gamma(\varepsilon, \delta, R, D)$, be the corresponding constants from the discussion of the discussion above. Recall from Lemmas 4, 5, and 8 that F is the set of foreign keys added to S at least once, N is the set of centres of internal anomalies and G is the set of good vertices. We know:

20:10 Better Space-Time-Robustness Trade-Offs for Set Reconciliation

$$\begin{aligned}
|F| &\leq \frac{\delta n}{(2D+1)^R} && \text{except with probability } \mathcal{O}\left(2^{-\frac{\delta n}{(2D+1)^R}}\right) && \text{by Lemma 4} \\
|N| &\leq \frac{\delta n}{(2D+1)^R} && \text{except with probability } \mathcal{O}\left(2^{-\frac{\delta n}{(2D+1)^R}}\right) && \text{by Lemma 5} \\
|G| &\geq (1-3\delta)n && \text{except with probability } \exp(-\gamma n) && \text{by Lemma 8}
\end{aligned}$$

The sum of the error probabilities is $\mathcal{O}(2^{-\beta n})$ where $\beta = \delta/(2D+1)^R + \gamma \log_2(e)$. Thus we may assume that all three stated events occur.

We now upper bound $|I^R|$. We pessimistically assume that each $i \in [n]$ that is not good is in I^R , which accounts for at most $3\delta n$ vertices. Any good $i \in [n]$ would be removed by the proper peeling process, so if $i \in I^R$ then an anomaly must have interfered. Therefore, the R -neighbourhood of i must include the centre of an internal anomaly or a vertex incident to a key from F . In other words, the R -neighbourhood of i must include a vertex from $|N \cup h(F)|$. At most $|N \cup h(F)| \cdot (2D+1)^R$ good vertices can be affected in this way, namely those reachable from a vertex in $N \cup h(F)$ via a path of length at most R traversing only vertices of degree at most D . Taken together the size of I_R is bounded by

$$|I_R| \leq 3\delta n + |N \cup h(F)| \cdot (2D+1)^R \leq 3\delta n + \frac{4\delta}{(2D+1)^R} n \cdot (2D+1)^R \leq 7\delta n. \quad \blacktriangleleft$$

3.6 Late Rounds of the Peeling Process

► **Lemma 10.** *There exists a constant δ such that for any $1 \leq r \leq n$ there is an event E with probability $\mathcal{O}(2^{-r})$ such that, when E occurs, then*

$$\forall I \subseteq [n] \text{ with } 5r \leq |I| \leq \delta n : |S[I]| \leq \frac{3}{5}|I|.$$

Proof. We use a union bound over all possible sizes i of I and all possible sets of size i . Let p_i be the probability that a violating set of size i exists and $j := \frac{3}{5}i$. Multiplying the number $\binom{n}{i}$ of choices for I , the number $\binom{m}{j}$ of choices for $T \subseteq S$ of size j and the probability for $T \subseteq S[I]$ (i.e. the probability that all $x \in T$ satisfy $h(x) \subseteq I$) gives

$$\begin{aligned}
p_i &\leq \binom{n}{i} \binom{m}{j} \left(\frac{i}{n}\right)^{3j} \leq \left(\frac{ne}{i}\right)^i \left(\frac{ne}{j}\right)^j \left(\frac{i}{n}\right)^{3j} \\
&= e^{i+j} \left(\frac{i}{j}\right)^j \left(\frac{i}{n}\right)^{-i+2j} \leq e^{3i} \left(\frac{i}{n}\right)^{i/5} = \left(\frac{e^{15}i}{n}\right)^{i/5}.
\end{aligned}$$

Setting $\delta = 1/(2e^{15})$ we can bound the overall failure probability p as

$$\begin{aligned}
p &\leq \sum_{i=5r}^{\delta n} p_i \leq \sum_{i=5r}^{\delta n} \left(\frac{e^{15}i}{n}\right)^{i/5} \leq \sum_{i=5r}^{\delta n} \left(\frac{e^{15} \cdot \delta n}{n}\right)^{i/5} \leq \sum_{i=5r}^{\delta n} \left(\frac{1}{2}\right)^{i/5} \\
&= \sum_{i=5r}^{\delta n} \left(\left(\frac{1}{2}\right)^{1/5}\right)^i \leq 2^{-r} \cdot \sum_{i=0}^{\infty} \left(\left(\frac{1}{2}\right)^{1/5}\right)^i = 2^{-r} \cdot \Theta(1) = \mathcal{O}(2^{-r}). \quad \blacktriangleleft
\end{aligned}$$

3.7 Proof of Theorem 2

We now assemble the previous Lemmas into a proof of Theorem 2. Note that we only claimed an error probability of $2^{-\Omega(r)}$, not $\mathcal{O}(2^{-r})$. By a change of parameters it therefore suffices that we show $\Pr[|S_{\text{dec}} \Delta S| > 185r] = \mathcal{O}(2^{-r})$.

The parameter $\varepsilon > 0$ and $r = o(n)$ are given. We use the absolute constant $\delta > 0$ from Lemma 10. We apply Lemma 9 for ε and $\delta' = \delta/7$ which gives us $R \in \mathbb{N}$ and $\beta > 0$. For convenience, let us summarise the statements used in the following. Since each holds with probability $1 - \mathcal{O}(2^{-r})$ we need only show that they imply $|S_{\text{dec}} \Delta S| \leq 185r$.

$$|I^R| < \delta n \quad \text{by Lemma 9} \quad (1)$$

$$\forall I \subseteq [n] \text{ with } 5r \leq |I| \leq \delta n : |S[I]| \leq \frac{3}{5}|I| \quad \text{by Lemma 10} \quad (2)$$

$$|N| \leq r \text{ and } |F| \leq r \quad \text{by Lemmas 4 and 5} \quad (3)$$

Recall that *anomalous decoding steps* are those that add a key rather than removing a key.

▷ **Claim 11.** There are at most $|N| + 3|F| \leq 4r$ anomalous decoding steps per round.

Proof. There are at most $3|F|$ buckets in which a key from F might be placed. We may conservatively assume that there is an anomalous decoding step in each round for each of these buckets. Anomalous decoding steps where no foreign key is involved can arise, by definition, only in centres of internal anomalies, hence there can be at most $|N|$ per round. \triangleleft

We are most interested in the first $\hat{\rho}$ rounds of decode where $\hat{\rho} := R + \log_{\frac{20}{19}}(\frac{n}{300r}) = \Theta(\log \frac{n}{r})$. We now make sure that we are referring to well-defined rounds. Execution might end due to $Q = \emptyset$, or due to the time limit. If execution ends prior to round $\hat{\rho}$ due to $Q = \emptyset$, just pretend the loop is executed for a suitable number of additional rounds. Note that an additional round does not perform any steps when $Q = \emptyset$. Let us now deal with the time limit.

▷ **Claim 12.** The limit $t_{\max} = 2n$ on steps does not take effect within the first $\hat{\rho}$ rounds.

Proof. We use a potential argument. There are n keys in the beginning. Each non-anomalous decoding step removes a key and each anomalous decoding step adds a key. Hence, if a denotes the number of anomalous decoding steps in the first $\hat{\rho}$ rounds then the total number of decoding steps is at most $n + 2a$. The previous claim gives $a \leq 4r\hat{\rho}$. We can therefore bound the number of decoding steps in the first $\hat{\rho}$ rounds by

$$n + 2a \leq n + 8r\hat{\rho} = n(1 + \Theta(\frac{r}{n} \log \frac{n}{r})) = n(1 + o(1)) \leq 2n.$$

The last step uses $r = o(n)$ and $\lim_{x \rightarrow 0} x \log \frac{1}{x} = 0$. \triangleleft

We can now safely refer to each round $\rho \in \{1, \dots, \hat{\rho}\}$. Let S^ρ be the set of keys stored in the IBLT after ρ rounds, $I^\rho = h(S^\rho)$ the set of buckets touched by these keys and $S_0^\rho = S^\rho \setminus F$.

▷ **Claim 13.** If $300r \leq |I^\rho| \leq \delta n$ then $|I^{\rho+1}| \leq \frac{19}{20}|I^\rho|$.

Proof. Since $S_0^\rho \subseteq S[I^\rho]$ we know $|S_0^\rho| \leq \frac{3}{5}|I^\rho|$ from Equation (2). The number of incidences due to the keys from S_0^ρ is $3|S_0^\rho| \leq \frac{9}{5}|I^\rho| = \frac{9}{10} \cdot 2|I^\rho|$. Thus at most a $\frac{9}{10}$ -fraction of the buckets in I^ρ have two or more incidences to keys from S_0^ρ . This leaves at least $\frac{1}{10}|I^\rho|$ buckets with at most one incidence to S_0^ρ . If i is such a bucket, we would normally expect $i \notin I^{\rho+1}$.

There are only two exceptions. A foreign key might be stored in i at some point during round $\rho + 1$ or a key from S is added back into bucket i at some point during round $\rho + 1$. Each foreign key and each anomalous decoding step affects at most 3 buckets. Hence, by an earlier claim, at most $3(|F| + 4r) \leq 15r$ buckets are exceptional in this sense. Together we get

$$|I^{\rho+1}| \leq |I^\rho| - \frac{1}{10}|I^\rho| + 15r \leq \frac{9}{10}|I^\rho| + 15r \leq \frac{19}{20}|I^\rho|.$$

where the last step uses $|I^\rho| \geq 300r$. \triangleleft

20:12 Better Space-Time-Robustness Trade-Offs for Set Reconciliation

▷ **Claim 14.** If $|I^\rho| \leq 300r$ then $|I^{\rho+1}| \leq 300r$.

Proof. Assume w.l.o.g. that $|I^\rho| \geq 5r$ and modify the last step of the previous claim:

$$\frac{9}{10}|I^\rho| + 15r \leq \frac{9}{10} \cdot 300r + 15r = 285r < 300r. \quad \triangleleft$$

Let now ρ_{\max} be the last round that is fully executed.

▷ **Claim 15.** $|I^{\rho_{\max}}| \leq 300r$.

Proof. We have $|I^R| < \delta n$ by Equation (1). In subsequent rounds $\rho \in \{R+1, R+2, \dots\}$, the size of I^ρ decreases by at least a factor of $\frac{19}{20}$ until being at most $300r$ by Claim 13. By choice of $\hat{\rho}$ we have $I^\rho \leq 300r$ for some $\rho \leq \hat{\rho}$. By Claim 14 $|I^\rho|$ will subsequently not rise above $300r$ for larger values of ρ . \triangleleft

By the last claim and Equation (2) we have

$$|S_0^{\rho_{\max}}| \leq S[|I^{\rho_{\max}}|] \leq \frac{3}{5} \cdot (\max(5r, |I^{\rho_{\max}}|)) \leq \frac{3}{5} \cdot 300r = 180r.$$

The set $S^{\rho_{\max}}$ of all keys after round ρ_{\max} might additionally include up to $|F| \leq r$ foreign keys. Since an additional round might be started but cut short due to the time limit, another $4r$ keys might be added due to anomalous decoding steps. Overall the set $S\Delta S_{\text{dec}}$ of keys returned in the end has size at most $180r + r + 4r = 185r$. This concludes the proof.

4 Set reconciliation with error-correcting codes

The set reconciliation problem can be modeled by encoding a set as a binary string of length U which is then sent through a channel inflicting up to D errors (bitflips). Recovering the errors is then equivalent to reporting the keys inserted to or deleted from the set. This reduction immediately implies that the problem can be solved by an appropriate application of error-correcting codes, in particular a linear code such as BCH, Reed-Solomon or expander code (see [14]). This reduction, in turn, highlights the relationship of set reconciliation to the problem of *sparse recovery*: reconstructing a sparse vector from a set of linear measurements (see e.g. [4]).

We will use a BCH code over the field $GF(2^w)$ with $U = 2^w - 1$. It is known (see e.g. [17]) that if up to D errors have to be corrected, the binary parity-check matrix \mathcal{H} for BCH has dimensions $Dw \times U$. Let each set S be conceptually represented by a binary vector of size U . Denote $\mathcal{H}(S)$ the image of S by \mathcal{H} . Then for two sets S, T , $\mathcal{H}(S) \oplus \mathcal{H}(T) = \mathcal{H}(S\Delta T)$. That is, if $|S\Delta T| \leq D$, then $\mathcal{H}(S\Delta T)$, called *syndrome*, encodes the differences between S and T and can be reconstructed by a syndrome decoding algorithm. This implies that $\mathcal{H}(S)$ can be defined as a sketch of S of Dw bits so that the differences between two sets can be recovered from the XOR of their sketches, provided that there are at most D of them.

From the algorithmic viewpoint, inserting a key amounts to XORing the sketch with the corresponding column of \mathcal{H} composed of D blocks of w bits (elements of $GF(2^w)$). Computing each block amounts to D multiplications in $GF(2^w)$ ⁶. A multiplication in $GF(2^w)$ can be done with $\mathcal{O}(w \log w)$ bit operations [13], but only with $\mathcal{O}(w)$ operations in the RAM model (see Appendix). Thus, an insertion of a key takes time $\mathcal{O}(Dw) = \mathcal{O}(D \log U)$.

⁶ We refer to [5] for details on how keys of U are represented as elements of $GF(2^w)$.

```

Algorithm insert( $x, \mathcal{S}$ ):
┌ toggle( $x$ ) in  $\mathcal{S}.A$ 
├  $\mathcal{S}.H \leftarrow \mathcal{S}.H \oplus h(x)$ 
└ toggle_BCH( $x, \mathcal{S}.C$ )

Algorithm diff( $\mathcal{S}_1, \mathcal{S}_2$ ):
┌  $\hat{\mathcal{S}}.A \leftarrow \text{merge}(\mathcal{S}_1.A, \mathcal{S}_2.A)$ 
├  $\hat{\mathcal{S}}.H \leftarrow \mathcal{S}_1.H \oplus \mathcal{S}_2.H$ 
├  $\hat{\mathcal{S}}.C \leftarrow \text{merge\_BCH}(\mathcal{S}_1.C, \mathcal{S}_2.C)$ 
└ return  $\hat{\mathcal{S}}$ 

Algorithm report( $\mathcal{S}$ ):
┌  $S_{\text{dec}} \leftarrow \text{decode}(\mathcal{S}.A)$ 
├ for  $x \in S_{\text{dec}}$  do
├ ┌  $\mathcal{S}.H \leftarrow \mathcal{S}.H \oplus h(x)$ 
├ └ if  $\mathcal{S}.H \neq 0$  then
├ └ ┌ for  $x \in S_{\text{dec}}$  do
├ └ └ ┌ toggle_BCH( $x, \mathcal{S}.C$ ) // delete  $x$ 
├ └ └ └ from  $\mathcal{S}.C$ 
├ └ └  $S'_{\text{dec}} \leftarrow \text{decode\_BCH}(\mathcal{S}.C)$ 
└ return  $S_{\text{dec}} \Delta S'_{\text{dec}}$ 

```

■ **Figure 2** Sketch operations. `insert`(x, \mathcal{S}) inserts key x to sketch \mathcal{S} , `diff`($\mathcal{S}_1, \mathcal{S}_2$) computes the difference of sketches \mathcal{S}_1 and \mathcal{S}_2 , `report`(\mathcal{S}) reports keys stored in \mathcal{S} .

Decoding is a more complex operation. The efficient syndrome decoding algorithm of [5] requires $\mathcal{O}(D^2w)$ multiplications in $GF(2^w)$ resulting in $\mathcal{O}(D^2w^2) = \mathcal{O}(D^2 \log^2 U)$ time for decoding. Note that if the number $d < D$ of errors is known, then by properties of syndromes, it is sufficient to decode the first dw bits of the sketch, and the decoding complexity becomes $\mathcal{O}(d^2 \log^2 U)$.

Denote `toggle_BCH`(x, C) and `decode_BCH`(C) the operation of inserting/deleting a key x to/from a BCH sketch C and decoding a BCH sketch C , respectively. Complexity of BCH sketches is summarized in the following theorem.

► **Theorem 16.** *Consider a set S containing no more than D keys from U . Let S be stored in a BCH sketch C of $D \log(U + 1)$ bits of space. Then*

- `toggle_BCH`(x, C) takes $\mathcal{O}(D \log U)$ time,
- `decode_BCH`(C) decodes keys of S with no error in $\mathcal{O}(D^2 \log^2 U)$ time.

Reconciliation of two sets is done by simply XORing their sketches C_1 and C_2 , which can be trivially sped up by packing bits in words. We denote this operation by `merge_BCH`(C_1, C_2). A software implementation of a BCH sketch for set reconciliation is available [19].

5 IBLT with stash

Our goal is to enhance the IBLT in order to solve set reconciliation with a much stronger success guarantee than that provided by IBLT alone ([1], Table 1) while keeping small space and fast decoding time provided by the IBLT approach. Given an upper bound D on the size of set differences, our sketch \mathcal{S} consists of three components depending on a parameter r :

- an IBLT $\mathcal{S}.A$ of size $n = (c_3 + \varepsilon)D$ with three hash functions, as defined in Section 2,
- a control checksum $\mathcal{S}.H$ of r bits,
- a *stash* data structure $\mathcal{S}.C$ defined as a BCH syndrome of $r \log U$ bits (see Section 4).

The control checksum $\mathcal{S}.H$ is used to check if the decoding of the main IBLT $\mathcal{S}.A$ succeeded. Before the decoding, $\mathcal{S}.H = \bigoplus_{x \in S} h(x)$ where $h : \mathcal{U} \rightarrow 2^r$ is a random hash function and S is the currently stored key set.

The stash $\mathcal{S}.C$ is configured so that it can decode up to r keys with no error, as described in Section 4. The entire sketch takes $(c_3 + \varepsilon)D \log U + r(1 + \log U) = (c_3 + \varepsilon)D \log U + o(D)$ bits.

Figure 2 shows Algorithms for inserting a key to a sketch, computing a sketch difference, and reporting keys.

A run of `report`(\mathcal{S}) (Figure 2) can follow two scenarios. If the number of keys in the sketched set is no more than D , the main IBLT recovers those keys with probability $1 - \tilde{O}(1/D)$ (Theorem 2). With probability $\tilde{O}(1/D)$, however, the recovery may fail by “getting stuck”, by resulting in a “falsely empty” IBLT (see [1]), or by being aborted after $2n$ steps. Checking the checksum ensures that an incorrect decoding will be recognized except with probability $\mathcal{O}(1/2^r)$. By Theorem 2, with probability $2^{-\Omega(r)}$, the output S_{dec} differs from S by no more than r missing or foreign keys. The algorithm resorts then to the stash in order to correct the output, i.e. to recover the missing/superfluous keys, under assumption that there are at most r of them. In this case, the stash reports those keys with no error.

The execution time of `report` depends on whether the correction step is triggered or not. If the decoding of the main IBLT is deemed to be successful by the checksum, the time taken will be $\mathcal{O}(D(1 + \frac{r}{w}))$, where w is the wordsize. Here we assume that updating the checksum is done by bit packing and takes $\mathcal{O}(\frac{r}{w})$ time. Otherwise the algorithm will spend additional $\mathcal{O}(r \log U)$ time on performing `toggle_BCH` on each of $\mathcal{O}(D)$ key in S_{dec} , followed by the decoding taking additional $\mathcal{O}(r^2 \log^2 U)$ time. Since the stash is activated with probability $\tilde{O}(1/D)$, the expected time is $\mathcal{O}(D(1 + \frac{r}{w}) + r \log U + \frac{r^2 \log^2 U}{D})$. Assuming $r = \min(\mathcal{O}(D/\log^2 U), \mathcal{O}(\log U))$, the time of `report` is $\mathcal{O}(D)$.

We summarize the properties of `report` in the following theorem.

► **Theorem 17.** *Let \mathcal{S} be a sketch built for a key set S and $r = \min(\mathcal{O}(D/\log^2 U), \mathcal{O}(\log U))$. Then*

- \mathcal{S} takes $(c_3 + \varepsilon)D \log U + r(1 + \log U) = (c_3 + \varepsilon)D \log U + o(D)$ bits,
- inserting an element to \mathcal{S} takes time $\mathcal{O}(r \log U)$,
- when $|S| \leq D$, `report`(\mathcal{S}) correctly recovers S with probability $1 - 2^{-\Omega(r)}$,
- `report`(\mathcal{S}) takes $\mathcal{O}(D)$ expected time.

Note that if $r = \omega(\log D)$, the failure probability for `report` to recover a set of up to D keys is $o(1/D)$, as opposed to $\tilde{O}(1/D)$ for recovery without stash [1].

For two sets S and T represented by their respective sketches \mathcal{S}_S and \mathcal{S}_T , `diff`($\mathcal{S}_S, \mathcal{S}_T$) is the sketch of symmetric difference $S\Delta T$. Therefore, Theorem 17 applies directly to the set reconciliation problem. In particular, if $|S\Delta T| \leq D$, then $S\Delta T$ can be recovered with guarantees stated in Theorem 17.

6 Distinguishing $S \setminus T$ and $T \setminus S$

The algorithm we presented computes symmetric difference $S\Delta T$ but is not capable of distinguishing elements of S and T in the output. However, the latter is desirable for many applications. Here we outline how our sketch and set reconciliation protocol (Figure 2) can be modified in order to make this possible. The general idea is to assign a “sign” to keys depending on whether they come from S or T and to keep track of it in a consistent manner. Formal proofs are left to the full version of the paper.

Modified IBLT. We modify our IBLT (Section 2) as follows. IBLT entries will now be *ternary* strings of $\{0, 1, 2\}^\nu$ seen as elements of group \mathbb{Z}_3^ν (additive group of $GF(3^\nu)$). The first *trit* encodes a *sign* and the other $\nu - 1$ trits encode a checksum. We use an appropriate encoding of \mathcal{U} into strings $\{0, 1, 2\}^{\nu-1}$ and denote by \tilde{x} the encoding of x .

Toggling a key x is replaced by two operations: insertion and deletion. Inserting x into an IBLT is done by adding (in \mathbb{Z}_3^ν) $1\tilde{x}$ to each of the three entries at $h(x)$, and deleting x is done by subtracting $1\tilde{x}$ from each of the three entries at $h(x)$ or, equivalently, adding

$2(-\tilde{x}) \equiv (-1)(-\tilde{x})$ where $-\tilde{x}$ is the inverse of \tilde{x} in $\mathbb{Z}_3^{\nu-1}$. Observe that insertion and deletion of x cancel each other out, and that inserting twice the same key x is equivalent to deleting x and vice versa.

IBLT operations (Figure 1) are modified as follows. Operation `looksPure` checks if the first trit equals 1 or $2 \equiv -1$. If it is 2, the key assumed to sit in this entry is the inverse of the retrieved key value. `toggle(x)` deletes or inserts x from/to each of the three entries at $h(x)$ depending on whether x has been identified as *pure* with the first trit 1 or 2, respectively. Merging two IBLTs `merge(A, A')` now becomes non-commutative and is done by subtracting (in \mathbb{Z}_3^{ν}) A' from A entry-wise.

Consider a sketch resulting from `merge(A, A')`. When reporting a set difference (modified Algorithm `report` in Figure 2) the algorithm now outputs *signed elements*, where elements annotated with 1 (i.e. those with the first trit equal to 1 at the moment of toggling) are interpreted as belonging to the first set, while those annotated with 2 are interpreted as belonging to the second. The symmetric difference operation Δ in both `decode` and `report` is also modified as follows. When computing $S_1 \Delta S_2$, an element is canceled only if it occurs in S_1 and S_2 with opposite signs. If an element occurs in S_1 and S_2 with the same sign, it is reported in the resulting set with the opposite sign.

Observe that the modified structure does not prevent anomalies to occur but only makes them less likely: an anomaly can still occur if the first trit is 1 or 2 but the entry actually contains more than one key. On the other hand, our analysis of Section 3 carries over to the signed case. The difference in S_{dec} compared to the original decoding is due to repeated reportings of the same key with different sign combinations, which does not affect the proof ideas of the main Theorem 2.

Modified control hashsum. Here the hashsum uses arithmetic summation instead of XOR. We assume that overflow is implemented in a consistent way, that is $(x + y) - y = x$ even if $x + y$ results in overflow. Toggling x entails subtracting $h(x)$ from the hashsum if x is annotated with 1, and adding if it is annotated -1 . The hashsum of `diff(S1, S2)` is now defined as $S_1.H - S_2.H$.

Modified BCH sketch. In order to deal with signed elements, we use BCH code over field $GF(3^{\nu-1})$. The BCH sketch of a set is still defined as the syndrome vector under assumption that each set element is encoded by value 1 in the error vector. The `merge_BCH` operation will now subtract the syndromes (in $GF(3)$) instead of XORing them. By linearity of the code, the “positive” and “negative” elements will correspond to positions with 1 and $2 \equiv -1$ respectively in the recovered error vector.

Given a BCH sketch resulting from the difference of sketches of input sets S and T , algorithm `report` will subtract the syndrome corresponding to the set S_{dec} of signed elements decoded by the IBLT. Naturally, positive and negative elements are encoded respectively by values 1 and 2 as well. Again, by linearity, the resulting syndrome will encode exactly the elements S'_{dec} so that $S_{dec} \Delta S'_{dec} = S \Delta T$, where positive (resp. negative) elements are those belonging to S (resp. T) only.

Observe that, in general, S_{dec} will contain a subset of $S \Delta T$ as well as possibly some foreign keys, however our modified definition of Δ will ensure a correct recovery of the original set. As an example, assume $x \in S \setminus T$ and assume x has been reported by the IBLT as positive, followed by another reporting of x (produced by an anomaly) as positive as well. From our definition of Δ , x will become negative in S_{dec} and will be added (rather than subtracted) to the BCH sketch. This will result in reporting x as negative again by the BCH sketch, and, by our definition of Δ , will be eventually correctly reported as positive in the final output.

The syndrome of the BCH code over $GF(3^{\nu-1})$ consists of at most $2D$ elements of $GF(3^{\nu-1})$ i.e. at most $4D \log_3 U$ bits by a straightforward encoding. Using ternary representation of keys and arithmetic operations in $GF(3^{\nu})$ instead of $GF(2^w)$ introduces only a constant factor change in time complexities of both insertion and decoding (see Appendix). In conclusion, time and robustness guarantees of Theorem 17 remain valid.

7 Concluding remarks

The presented solution to set reconciliation uses asymptotically negligible additional space and additional time for decoding compared to the IBLT-only solution, but provides a drastic improvement in robustness. The decoding time of our algorithm is small in expectation, however it becomes substantial in worst case. More precisely, when the BCH correction is activated, the decoding time becomes $O(Dr \text{polylog}(U))$ which can approach $O(D^2)$. We believe however this can be overcome by implementing the stash with expander codes which have much smaller insertion and decoding times. We leave details for future work.

Very recently, we learned about paper [8] that proposes a modified construction of IBLT and applies hash functions of restricted independence, rather than assuming them fully random. As a result, the construction takes a smaller space (compared to the original IBLT with the same error guarantee) and requires less randomness for hash functions. However, space is measured in terms of the number of IBLT buckets rather than in bits, which leads to a much larger multiplicative constant compared to our result. The decoding time of [8] appears to be larger than ours as well (not specified in the paper). On the other hand, an emphasis of [8] is to using hash functions of restricted randomness, the issue that we don't deal with in this work.

References

- 1 Jakob Bæk Tejs Houen, Rasmus Pagh, and Stefan Walzer. Simple set sketching. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 228–241. SIAM, 2023.
- 2 Daniella Bar-Lev, Avi Mizrahi, Tuvi Etzion, Ori Rottenstreich, and Eitan Yaakobi. Coding for IBLTs with listing guarantees. In *IEEE International Symposium on Information Theory, ISIT 2023, Taipei, Taiwan, June 25-30, 2023*, pages 1657–1662. IEEE, 2023. doi:10.1109/ISIT54713.2023.10206563.
- 3 Mahdi Cheraghchi. Coding-theoretic methods for sparse recovery. In *49th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2011, Allerton Park & Retreat Center, Monticello, IL, USA, 28-30 September, 2011*, pages 909–916. IEEE, 2011. doi:10.1109/Allerton.2011.6120263.
- 4 Mahdi Cheraghchi and João Ribeiro. Simple codes and sparse recovery with fast decoding. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 156–160. IEEE, 2019.
- 5 Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008. doi:10.1137/060651380.
- 6 David Eppstein and Michael T Goodrich. Straggler identification in round-trip data streams via Newton's identities and invertible Bloom filters. *IEEE Transactions on Knowledge and Data Engineering*, 23(2):297–306, 2011.
- 7 David Eppstein, Michael T. Goodrich, Frank Uyeda, and George Varghese. What's the difference? Efficient set reconciliation without prior context. In *Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11*, pages 218–229, New York, NY, USA, 2011. Association for Computing Machinery. doi:10.1145/2018436.2018462.

- 8 Nils Fleischhacker, Kasper Green Larsen, Maciej Obremski, and Mark Simkin. Invertible Bloom lookup tables with less memory and randomness. *CoRR*, abs/2306.07583, 2023. [arXiv:2306.07583](https://arxiv.org/abs/2306.07583), [doi:10.48550/arXiv.2306.07583](https://doi.org/10.48550/arXiv.2306.07583).
- 9 Nils Fleischhacker, Kasper Green Larsen, and Mark Simkin. Property-preserving hash functions for Hamming distance from standard assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022*, pages 764–781, Cham, 2022. Springer International Publishing.
- 10 Sumit Ganguly. Counting distinct items over update streams. *Theoretical Computer Science*, 378(3):211–222, 2007. Algorithms and Computation. [doi:10.1016/j.tcs.2007.02.031](https://doi.org/10.1016/j.tcs.2007.02.031).
- 11 Sumit Ganguly and Anirban Majumder. Deterministic k-set structure. *Information Processing Letters*, 109(1):27–31, 2008. [doi:10.1016/j.ip1.2008.08.010](https://doi.org/10.1016/j.ip1.2008.08.010).
- 12 Michael T Goodrich and Michael Mitzenmacher. Invertible Bloom lookup tables. In *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 792–799. IEEE, 2011.
- 13 David Harvey and Joris van der Hoeven. Polynomial multiplication over finite fields in time $\mathcal{O}(n \log n)$. *J. ACM*, 69(2):12:1–12:40, 2022. [doi:10.1145/3505584](https://doi.org/10.1145/3505584).
- 14 Mark G Karpovsky, Lev B Levitin, and Ari Trachtenberg. Data verification and reconciliation with generalized error-control codes. *IEEE Transactions on Information Theory*, 49(7):1788–1793, 2003.
- 15 Jeong Han Kim. Poisson cloning model for random graphs. In *Proc. ICM, Vol. III*, pages 873–898, 2006. URL: <https://www.mathunion.org/fileadmin/ICM/Proceedings/ICM2006.3/ICM2006.3.ocr.pdf>.
- 16 M. Leconte, M. Lelarge, and L. Massoulié. Convergence of multivariate belief propagation, with applications to cuckoo hashing and load balancing. In *Proc. 24th SODA*, pages 35–46, 2013. URL: <http://dl.acm.org/citation.cfm?id=2627817.2627820>.
- 17 F.J. MacWilliams and N.J.A. Sloane. *The Theory of Error-Correcting Codes*. North-holland Publishing Company, 2nd edition, 1978.
- 18 Colin McDiarmid. *On the method of bounded differences*, pages 148–188. London Mathematical Society Lecture Note Series. Cambridge University Press, 1989. [doi:10.1017/CB09781107359949.008](https://doi.org/10.1017/CB09781107359949.008).
- 19 Minisketch: an optimized library for BCH-based set reconciliation. <https://github.com/sipa/minisketch/>. Accessed: 2023-03-28.
- 20 Yaron Minsky, Ari Trachtenberg, and Richard Zippel. Set reconciliation with nearly optimal communication complexity. *IEEE Transactions on Information Theory*, 49(9):2213–2218, 2003. [doi:10.1109/TIT.2003.815784](https://doi.org/10.1109/TIT.2003.815784).
- 21 Michael Mitzenmacher and Rasmus Pagh. Simple multi-party set reconciliation. *Distributed Computing*, 31:441–453, 2018.
- 22 Michael Mitzenmacher and George Varghese. Biff (Bloom filter) codes: Fast error correction for large data sets. In *2012 IEEE International Symposium on Information Theory Proceedings*, pages 483–487. IEEE, 2012.
- 23 Avi Mizrahi, Daniella Bar-Lev, Eitan Yaakobi, and Ori Rottenstreich. Invertible Bloom lookup tables with listing guarantees. *CoRR*, abs/2212.13812, 2022. [arXiv:2212.13812](https://arxiv.org/abs/2212.13812), [doi:10.48550/arXiv.2212.13812](https://doi.org/10.48550/arXiv.2212.13812).
- 24 Michael Molloy. Cores in random hypergraphs and boolean formulas. *Random Structures & Algorithms*, 27(1):124–135, 2005.
- 25 Thomas Morgan. *An exploration of two-party reconciliation problems*. PhD thesis, Harvard University, Cambridge, MA, May 2018. URL: <https://dash.harvard.edu/handle/1/39947174>.

A Appendix

Below we describe how to perform multiplication in $GF(2^w)$ and $GF(3^\nu)$ in RAM model in times $O(w)$ and $O(\nu)$ respectively.

A.1 Multiplications in $GF(2^w)$ in RAM model in $O(w)$ time

The algorithm we describe is folklore, but since we did not find a reference we briefly sketch it here. The multiplication of two elements x and y in $GF(2^w)$ proceeds in two phases. In the first phase, one computes the product $z = x \cdot y$, where \cdot refers to a carry-less multiplication of the binary representation of x and y . Thus, z will be a bitstring of length $2w - 1$. This phase can be easily implemented in $O(w)$ time using $O(w)$ XOR and shift operation (the result is the XOR of copies of y right-shifted by bit-positions of x that are equal to one)⁷. The second phase is the polynomial euclidean division of z by an irreducible polynomial P of degree w . The remainder of the division will be the final result of the multiplication and can also be easily implemented in $O(w)$ time using $O(w)$ XOR and shift operations: successively for decreasing positions i of z starting from position $2w - 1$ check whether the bit at that position is equal to one and if so, do $z \leftarrow (z \oplus (P \ll (i - w)))$.

A.2 Multiplications in $GF(3^\nu)$ in RAM model in $O(\nu)$ time

Likewise, multiplication in $GF(3^\nu)$ can be implemented in $O(\nu)$ time. For that we can use a representation of an element using 2ν bits, where each element from the base field $GF(3)$ is represented using 2 bits. Note that any integer in the range $[0..3^\nu - 1]$ can be converted to this representation in $O(\nu)$ time, by doing successive euclidean divisions by 3. Notice that a division by 3 can be simulated using multiplications and other elementary operations. For a technical reason that will become clear later, we will instead use 4 bits to represent each element of the base field, resulting in a representation that uses 4ν bits. Let G be the function that transforms integers of $[0..3^\nu - 1]$ into binary representation of 4ν bits. Let G^{-1} be the inverse of G . Now assume that given two elements x and y represented as integers, we want to compute $z = x \cdot y$. We will first convert x and y into their binary representations $x' = G(x)$ and $y' = G(y)$, then compute the product $z' = x' \cdot y'$, and finally convert z' back into an integer z using function G^{-1} . We now describe how we do multiplication in the 4ν -bit representation. For that, we successively extract the quadbits (4-bit units) of x' , multiply each of them with y' , and aggregate all of them. The reason we use 4 bits instead of 2 to represent each digit is to accomodate the temporary result of the multiplication of two digits (numbers in $[0..2]$) which will be in the range $[0..4]$ (before doing Modulo 3 operation). Finally, the remainder of euclidean division by an irreducible polynomial can also be done in time $O(\nu)$ using $O(\nu)$ operations. Figure 3 shows the pseudocode of all steps of the algorithm.

⁷ Note that in practice, most modern processors natively support carry-less multiplication in constant time using much less complex hardware implementation than standard multiplication, but is not part of the standard RAM operations.

Algorithm $G(x)$:

```

 $x' \leftarrow 0$ 
 $j \leftarrow 0$ 
for  $i \leftarrow 1$  to  $\nu$  do
   $x' \leftarrow x' + ((x \bmod 3) \ll j)$ 
   $x \leftarrow x/3$ 
   $j \leftarrow j + 4$ 
return  $x'$ 

```

Algorithm $G^{-1}(x')$:

```

 $x \leftarrow 0$ 
 $j \leftarrow 4 \cdot (\nu - 1)$ 
for  $i \leftarrow 1$  to  $\nu$  do
   $x \leftarrow x \cdot 3 + ((x' \gg j) \wedge (0011)_2)$ 
   $j \leftarrow j - 4$ 
return  $x$ 

```

Algorithm $\text{MOD3}(x)$:

```

//reduce each quadbit modulo 3
//works only if quadbit in range
[0..5]
 $A \leftarrow (x + 1) \wedge (0100)_2$ 
 $x \leftarrow x - ((A \gg 2) + (A \gg 1))$ 
return  $x$ 

```

Algorithm $\text{reduce}(x, P)$:

```

// $P$  : irreducible polynomial
//represented in quadbits
 $P' \leftarrow \text{MOD3}(2 \cdot P)$ 
 $j \leftarrow 4 \cdot (2 \cdot \nu - 1)$ 
for  $i \leftarrow 1$  to  $\nu$  do
  if  $(x \gg j) > 0$  then
     $x \leftarrow x + (0011)_2^j$ 
    if  $(x \gg j) = (P \gg (\nu - 1))$ 
      then
         $x \leftarrow x - P$ 
      else
         $x \leftarrow x - P'$ 
     $x \leftarrow \text{MOD3}(x)$ 
   $j \leftarrow j - 4$ 
return  $x$ 

```

Algorithm $\text{mult}(x', y')$:

```

 $z' \leftarrow 0$ 
 $j \leftarrow 4 \cdot (\nu - 1)$ 
for  $i \leftarrow 1$  to  $\nu$  do
   $t \leftarrow (x' \gg j) \wedge (0011)_2$ 
   $t \leftarrow t \cdot y'$ 
   $t \leftarrow \text{MOD3}(t)$ 
   $z' \leftarrow (z' \ll 4) + t$ 
   $z' \leftarrow \text{MOD3}(z')$ 
   $j \leftarrow j - 4$ 
 $z' \leftarrow \text{reduce}(z', P)$ 
return  $z'$ 

```

Algorithm $\text{MULT}(x, y)$:

```

 $x' \leftarrow G(x)$ 
 $y' \leftarrow G(y)$ 
 $z' \leftarrow \text{mult}(x', y')$ 
 $z \leftarrow G^{-1}(z')$ 
return  $z$ 

```

■ **Figure 3** Multiplication algorithm $GF(3^\nu)$ modulo an irreducible polynomial.