

# A rate-compatible solution to the set reconciliation problem

Francisco Lázaro, Balázs Matuz

Institute of Communications and Navigation of DLR (German Aerospace Center),

Wessling, Germany. Email: {Francisco.LazaroBlasco, Balazs.Matuz}@dlr.de

**Abstract**—We consider a set reconciliation setting in which two parties hold similar sets that they would like to reconcile. In particular, we focus on set reconciliation based on invertible Bloom lookup tables (IBLTs), a probabilistic data structure inspired by Bloom filters. IBLT-based set reconciliation schemes have the advantage of exhibiting low computational complexity, however, the schemes available in the literature are known to be far from optimal in terms of communication complexity (overhead). The inefficiency of IBLT-based set reconciliation can be attributed to two facts. First, it requires an estimate of the cardinality of the difference between the sets, which implies an increase in overhead. Second, to cope with uncertainties in the estimation of the cardinality of the set difference, IBLT schemes in the literature oversize the data structures, thus further increasing the overhead. In this work, we present a novel IBLT-based set reconciliation protocol that does not require estimating the cardinality of the set difference. The proposed scheme relies on what we termed multi-edge-type (MET) IBLTs. The simulation results illustrate that the novel scheme outperforms state-of-the-art IBLT-based approaches to set reconciliation in terms of communication cost, i.e., in terms of the number of bits to be exchanged.

## I. INTRODUCTION

### A. Motivation and Problem Description

Set reconciliation problems arise frequently in distributed systems with multiple copies of data across different nodes. Examples are redundant distributed databases, such as Amazon’s Dynamo [1] or Apache Cassandra [2], in which copies of the same database are stored in multiple data centers (nodes). Whenever a new entry is added to or deleted from the database, the change is propagated to all the nodes in the system. This may, however, fail in some situations, for instance when some of the nodes are unavailable due to a hardware failure, or when communication is impaired by network congestion. The task of restoring the consistency among different copies of the database can be cast as a set reconciliation problem. Set reconciliation protocols find also applications in distributed storage systems and remote file synchronization. For the latter application, the simplest configuration consists of a setup with two copies of a file at different network locations. When one of the copies is updated, one would like to update the other copy, while transmitting as little data as possible to the remote network location. This problem can be converted to a set reconciliation problem by relying on tools from graph theory [3]. Set reconciliation protocols can also be applied in gossip-spreading protocols [4]. Recently, similar techniques have been proposed to address the problem of block propagation in a

cryptocurrency network [5], as well as to synchronize the *mempools*<sup>1</sup> of nodes in the Bitcoin network [6].

In its simplest form, the set reconciliation problem can be described as follows. Two hosts  $A$  and  $B$  possess (possibly) different sets,  $\mathcal{S}_A$  and  $\mathcal{S}_B$ , respectively. The goal is to determine the set difference  $\Delta$  between the two sets, i.e., the set of elements that are present in the union of  $\mathcal{S}_A$  and  $\mathcal{S}_B$ , but not in their intersection,  $\Delta = (\mathcal{S}_A \cup \mathcal{S}_B) \setminus (\mathcal{S}_A \cap \mathcal{S}_B)$ . To reconcile their sets,  $A$  and  $B$  exchange information and perform certain computations. Thereby, both communication complexity (overhead) and computational (time) complexity shall be kept as small as possible. Intuitively, when the cardinality of the set difference  $\delta = |\Delta|$  is large,  $A$  and  $B$  may simply exchange their sets. However, this solution turns out to be highly sub-optimal when  $\delta$  is small. In this case, more sophisticated algorithms can be applied to reduce communication complexity.

A common approach to set reconciliation (and other similar problems) is the use of log files. The prerequisite here is the existence of a so-called prior context between the parties involved in the protocol. In our setting, this translates to the assumption that sets  $\mathcal{S}_A$  and  $\mathcal{S}_B$  were synchronized (i.e., identical) at a previous point in time  $\tau$ . After that, both hosts keep track respectively of all the changes in their sets by using local log files. After  $d$  units of time have elapsed, i.e., at time  $\tau + d$ , the two sets can be reconciled relying on the log files. While being simple, the use of log files becomes inefficient when the network connectivity is bad, the storage medium is unreliable, or when changes to the sets  $\mathcal{S}_A$  and  $\mathcal{S}_B$  happen very frequently. Finally, these schemes do not scale well in multi-party settings (when one aims at synchronizing more than 2 sets), given that they require the existence of prior context between all pairs of nodes. Our focus instead will be on context-free set reconciliation.

### B. Related Works

Fundamental limits for one-way set reconciliation, i.e., when one single communication message is allowed, were studied in [7]. Let the elements in  $\mathcal{S}_A$  and  $\mathcal{S}_B$  be  $\ell$  bit vectors, i.e.,  $\mathcal{S}_A, \mathcal{S}_B \subset \{0, 1\}^\ell$ , and let an upper bound  $t$  to the cardinality (or size) of the set difference  $\delta = |\Delta|$  be known, i.e.,  $\delta \leq t$ . Then, exact set reconciliation requires transmitting  $\sim t\ell$  bits. In contrast, when no upper bound is known, exact set reconciliation requires transmitting the whole set. Given an upper bound  $t$  to the set difference  $\delta$ , set reconciliation can be

<sup>1</sup>A mempool is a database of pending Bitcoin transactions, which have not yet been written into the Blockchain.

carried out relying on a binary code of length  $2^\ell$  which can correct  $t$  errors. This solution is shown to have an optimal communication complexity when using a perfect code [7]. However, it requires working with a code of length  $2^\ell$  which is impractical for the values of  $\ell$  required by most applications.

In [8] a scheme based on characteristic polynomial interpolation (CPI) was proposed which achieves nearly optimal communication complexity. In particular, when an upper bound  $t$  to the size of the set difference  $\delta$  is known, exact set reconciliation is possible with communication complexity  $\sim t\ell$  and time complexity  $\mathcal{O}(t^3)$ . When an upper bound to  $\delta$  is not known, [8] proposes an approximate<sup>2</sup> set reconciliation scheme. In this case, the scheme requires an additional feedback message and is thus no longer a one-way reconciliation protocol. Instead, we have a reconciliation scheme that requires one round of communication. This approximate scheme has communication complexity  $\sim \delta\ell$  and time complexity  $\mathcal{O}(\delta^4)$ .

In [9] a low-complexity approximate set reconciliation scheme based on invertible Bloom lookup tables (IBLTs) was proposed, which also requires one communication round. In brief, an IBLT [10] is a probabilistic data structure that supports the insertion and deletion of set elements, as well as listing of all the elements contained in the IBLT, provided that the number of elements in the IBLT is not too large (a formalization of these concepts follows in Section II). From a data structure point of view, an IBLT can be broken down into cells, which have size  $\sim \ell + \log \ell$ . When an upper bound to the size of the set difference is known, the communication complexity of IBLT-based set reconciliation is  $\sim ct(\ell + \log \ell)$ , where  $c \approx 1.22$ , and the time complexity is  $\mathcal{O}(t(\ell + \log \ell))$  [10].

IBLT-based set reconciliation requires inverting an IBLT, an operation that is identical to erasure decoding [11]. As such, different IBLT designs have been proposed in the literature which are akin to different erasure code designs. The original construction in [9] relies on *regular* invertible Bloom lookup tables (IBLTs) in which each set element is mapped into the same number of IBLT cells. By moving to *irregular* IBLTs it is possible to improve the constant  $c$ . In particular, it was first shown in [12] that  $c = 1.09$  can be achieved by relying on irregular IBLTs with two different degrees. In [13] a density evolution analysis of arbitrary irregular IBLTs was presented which allows computing  $c$  and thus optimizing the irregular IBLT. Furthermore, relying on results for random access protocols [14], it was argued that IBLTs with constants  $c$  arbitrarily close to one can be found.

A shortcoming of IBLT-based set reconciliation [9] is the fact that it requires knowledge of an upper bound  $t$  to the size of the set difference  $\delta$ . In most cases, rather than an upper bound to  $\delta$ , an (imperfect) estimate of  $\delta$  is available, denoted by  $\hat{\delta}$ . Obtaining this estimate usually requires an additional communication round between the hosts, which represents an additional overhead [9]. In practice, the estimated set difference size  $\hat{\delta}$  can substantially differ from the actual

one  $\delta$ . Intuitively, an over-estimation ( $\hat{\delta} > \delta$ ) results in an oversizing of the IBLT and thus in an unnecessarily high communication cost, whereas an under-estimation ( $\hat{\delta} < \delta$ ) results in a high probability of failure of the set reconciliation algorithm. To reduce the probability of a failure, [9] considers an oversizing of the IBLT. Additionally, a larger IBLT can be sent upon a failure, attempting set reconciliation using only this second IBLT. This approach is inefficient since it discards the first exchanged IBLT. This inefficiency was identified in [5], where it was proposed to use the two IBLTs to reconcile the sets, reducing the communication cost.

Both CPI and IBLT-based set reconciliation assume that elements in  $\mathcal{S}_A$  and  $\mathcal{S}_B$  have all equal length, i.e., that they are length- $\ell$  bit vectors. However, in some applications, the set elements might have variable lengths. For example, the set elements might be encoded using strings of length ranging between  $\ell_{\min}$  and  $\ell_{\max}$  bits. In this case, a possible workaround is padding all elements to the maximum length  $\ell_{\max}$ . This yields an increase in the overhead (communication cost). Alternatively, one may rely on other techniques for approximate set reconciliation protocols that can inherently deal with variable-length data. Multiple such approaches exist based on so-called approximate set membership data structures. In [15] a low-complexity multi-round protocol based on the exchange of Bloom filters was proposed. Protocols based on the exchange of counting Bloom filters and Cuckoo filters were proposed in [16] and [17], respectively. In [18] approximate set reconciliation protocols based on the exchange of hash lists, Bloom filters, and Merkle trees were proposed and compared. For a comparison of the different set reconciliation protocols in different practical settings, we refer the reader to [6].

### C. Contribution

In this work, we introduce a novel approach to IBLT-based set reconciliation. In contrast to previous works, our approach does not require an upper bound to (or an estimate of) the size of the set difference  $\delta$ . Our work borrows ideas from fountain codes [19]. In particular, we let the transmitter (say host  $A$ ) send the cells of the IBLT one by one. After receiving some cells, the receiver (host  $B$ ) can attempt to reconcile the sets. If set reconciliation fails, the receiver simply waits for more cells and re-attempts set reconciliation. When set reconciliation is successful, the receiver sends an acknowledgment to the transmitter so that it stops sending cells. This approach spares the overhead necessary for set difference estimation which is considerable, especially when the set difference is small [9]. To fully exploit the advantages of our approach, we introduce a novel IBLT structure which we term multi-edge-type (MET) IBLT. This novel data structure borrows ideas from rate-compatible codes, in particular from MET low-density parity-check (LDPC) codes [20]. The main advantage of the adoption of a MET structure is that it allows increasing the number of IBLT cells on demand while still being able to list all the elements with high probability. This is akin to the feature of rate-compatible codes of operating close to capacity at multiple code rates.

The remaining part of the paper is structured as follows. In Section II we introduce the basics of MET IBLTs. In Section III we develop an analysis to assess the communication

<sup>2</sup>By approximate it is meant that the scheme may fail, so that the parties wrongly believe that their sets are reconciled, but they are actually not. In practice, the scheme can be designed to make this failure probability arbitrarily low, at the cost of an increase in the communication cost.

efficiency of MET IBLTs. We sketch a flexible set reconciliation protocol using MET IBLTs in Section IV followed by a design example in Section V and complementary simulation results. Section VI provides a discussion of the contribution.

## II. MULTI-EDGE-TYPE INVERTIBLE BLOOM LOOKUP TABLES

### A. Definitions

We first introduce a few definitions.

**Definition 1** (Key-value pair). *A key-value pair  $z$  is a data structure composed of two fields, a key  $x$  of length  $\nu$  bits and a value  $y$  of length  $\kappa$  bits, where typically  $\kappa \gg \nu$ . The key  $x$  is obtained as a function of  $y$ ,  $x = g(y)$  where the mapping is many to one.<sup>3</sup>*

The term key-value pair has its origin in the field of database systems. In particular, in a so-called key-value database every database entry is a key-value pair. In this context, the key can be thought of as a short (unique) identifier of an element in the database, whereas the value is the actual data which can be orders of magnitude larger than the key. In our case, every key-value pair will represent a set element, and we will assume that the key associated with a set element is obtained as a hash function of the set element's value.

**Definition 2** (Cell). *A cell  $c$  is a data structure containing two different fields, data and count where:*

- *data = (data.x, data.y) is a bit string of length  $\nu + \kappa$ . The bit strings data.x and data.y contain, respectively, the binary XOR of the keys and values that have been mapped to the cell.*
- *count is an integer. It contains the number of key-value pairs that have been mapped to this cell (see Section II-B).*

**Definition 3** (IBLT). *An invertible Bloom lookup table (IBLT) is a data structure that is used to represent a set of key-value pairs. Formally, it consists of an array of  $m$  cells  $\mathbf{c} = [c_1 \ c_2 \ \dots \ c_m]$ . The choice of the hash functions (Section II-B) determines how key-value pairs are mapped into the IBLT cells.*

IBLTs support multiple operations, such as the insertion and deletion of key-value pairs, as well as the listing of all inserted key-value pairs (with high probability). For simplicity, and for being consistent with the literature, we will sometimes abuse  $\mathbf{c}$  to refer to the IBLT itself, rather than only to its  $m$  cells. Furthermore, we will sometimes say that we ‘transmit an IBLT’ when we actually mean that we transmit the array of cells  $\mathbf{c}$ .

**Definition 4** (MET IBLT). *A multi-edge-type (MET) IBLT is a generalization of an IBLT that assigns different labels or types to key-value pairs and cells.*

The notion of types allows putting specific constraints on the mapping of key-value pairs to cells. From a data structure

point of view, introducing a MET structure allows addition of new groups of cells on demand. Furthermore, by judiciously choosing the hash functions used to map key-value pairs into cells we can improve the performance of the recovery operation of our data structure (see Section III). In Section V we will highlight the advantages of MET IBLTs used for set reconciliation compared to unstructured IBLTs which can be seen as a MET IBLT with one single key-value pair type and one single cell type.

### B. MET IBLT Description

We assign to each cell and key-value pair a label referred to as type which will play a role in how the mapping of key-value pairs into cells is done. Let  $d_c$  be the number of different cell types,  $m_i$ ,  $i \in \{1, 2, \dots, d_c\}$  be the number of cells of type  $i$ , and  $m = \sum_i m_i$  be the length (number of cells) of the MET IBLT. For convenience, we assume that the cells are ordered according to their type, so that cells  $c_1, c_2, \dots, c_{m_1}$  are of type 1, cells  $c_{m_1+1}, c_{m_1+2}, \dots, c_{m_1+m_2}$  are of type 2, and so forth.

Further, let the total number of different key-value pair types be  $d_d$ . The type of a key-value pair  $z$  is determined by applying a hash function  $h_{\mathbf{p}}(z.x)$  to the key  $z.x$ , where  $h_{\mathbf{p}}(z.x) = j$  maps an input  $z.x \in \{0, 1\}^\nu$  to an output  $j \in \{1, 2, \dots, d_d\}$ . This hash function is parameterized by the probability vector  $\mathbf{p} = [p_1 \ p_2 \ \dots \ p_{d_d}]$ . The vectors  $[1 \ 2 \ \dots \ d_d]$  and  $\mathbf{p}$  describe a probability mass function (p.m.f.), the so-called *type distribution*. In particular, under the assumption that the input  $z.x$  is uniformly distributed, the output of  $h_{\mathbf{p}}(z.x)$  samples the type distribution.

The mapping of a key-value pair of type  $j$  to cells of type  $i$  is achieved through the hash function  $H_{\mathbf{m}, \mathbf{d}_j}(\cdot)$ . It is parameterized by  $\mathbf{m} = [m_1 \ m_2 \ \dots \ m_{d_c}]$  and  $\mathbf{d}_j = [d_{1,j} \ d_{2,j} \ \dots \ d_{d_c,j}]$  and is used to sample at random  $d_{i,j}$  cells of type  $i$  into which the key-value pair of type  $j$  is mapped. The function outputs a vector  $\mathbf{g} = H_{\mathbf{m}, \mathbf{d}_j}(z.x)$  of length  $\sum_{i=1}^{d_c} d_{i,j}$ , which contains  $d_{i,j}$  indices of cells of type  $i$  (where the indices are between  $1 + \sum_{k=1}^{i-1} m_k$  and  $\sum_{k=1}^i m_k$ ).

**Example 1** (Hash function  $H_{\mathbf{m}, \mathbf{d}_j}$ ). *Assume  $\mathbf{m} = [2 \ 3 \ 5]$  and  $\mathbf{d}_j = [1 \ 1 \ 2]$ . We have  $m_1 = 2$  cells of type 1,  $m_2 = 3$  cells of type 2 and  $m_3 = 5$  cells of type 3. Hence, indices 1 and 2 are associated with cells of type 1. Indices 3, 4, and 5 are associated with cells of type 2, and indices 6, 7, 8, 9, and 10 are associated with cells of type 3. The hash function returns 4 indices in total, one index associated with a cell of type 1, one index associated with a cell of type 2, and two indices associated with cells of type 3. For a first key  $x_1$ , the hash function  $H_{\mathbf{m}, \mathbf{d}_j}(x_1)$  may return  $H_{\mathbf{m}, \mathbf{d}_j}(x_1) = [1 \ 4 \ 7 \ 9]$ . For a second key  $x_2$ , the hash function  $H_{\mathbf{m}, \mathbf{d}_j}(x_2)$  may return  $H_{\mathbf{m}, \mathbf{d}_j}(x_2) = [2 \ 3 \ 6 \ 7]$ .*

MET IBLTs support different operations that are discussed in more detail in the next sections:

- **Initialize()**. This operation sets the different fields of all the cells in the IBLT to zero (see Algorithm 1).

<sup>3</sup>This means that key collisions are in principle possible. However, by choosing  $\nu$  large enough, one can make the key collision probability small enough for practical applications.

---

**Algorithm 1** Initialization

---

```
procedure INITIALIZE()  
  for  $i = 1, 2, \dots, m$  do  
     $c_i.count = 0$   
     $c_i.data = 0$ 
```

---

**Algorithm 2** Insertion

---

```
procedure INSERT( $z$ )  
   $j \leftarrow h_p(z.x)$   
   $\mathbf{g} \leftarrow H_{m,d_j}(z.x)$   
  for  $k = 1, 2, \dots, \text{length}(\mathbf{g})$  do  
     $c_{g_k}.count = c_{g_k}.count + 1$   
     $c_{g_k}.data = \text{XOR}(c_{g_k}.data, z)$ 
```

---

**Algorithm 3** Deletion

---

```
procedure DELETE( $z$ )  
   $j \leftarrow h_p(z.x)$   
   $\mathbf{g} \leftarrow H_{m,d_j}(z.x)$   
  for  $k = 1, 2, \dots, \text{length}(\mathbf{g})$  do  
     $c_{g_k}.count = c_{g_k}.count - 1$   
     $c_{g_k}.data = \text{XOR}(c_{g_k}.data, z)$ 
```

---

**Algorithm 4** Recovery

---

```
procedure RECOVER()  
  while  $\exists i \in [1, m] | c_i.count = 1$  do  
    add  $z = c_i.data$  to the output list  
    call Delete( $z$ )
```

---

- Insert( $z$ ). The insertion operation *adds* the key-value pair  $z$  to the IBLT following certain mapping rules (see Algorithm 2).
- Delete( $z$ ). The deletion operation *removes* the key-value pair  $z$  from the IBLT (see Algorithm 3).
- Recover(). This operation aims at listing all key-value pairs stored in the IBLT. If this operation provides the full list of key-value pairs in the IBLT, we say it succeeds, else it fails (see Algorithm 4).

### C. Insertion of a set $\mathcal{S}$ into a MET IBLT

We consider now the insertion of a set  $\mathcal{S} = \{z_1, z_2, \dots, z_n\}$  of  $n$  elements (key-value pairs) into a MET IBLT that has been previously initialized to zero by Algorithm 1. The elements of  $\mathcal{S}$  are successively inserted into the IBLT as described by Algorithm 2: first, every element  $z = (x, y)$  is assigned a data node type  $j = h_p(z.x)$ . Next, the hash function  $H_{m,d_j}(z.x)$  is used to obtain a vector  $\mathbf{g}$  that contains the indices of the cells into which  $z$  is to be mapped, i.e., those cells whose *data* field is XOR-ed with  $z$ . The respective *count* fields are increased by one.

### D. Recovery of the set $\mathcal{S}$

The recovery operation is a low-complexity algorithm that aims at inverting the IBLT, i.e., its goal is to recover all the key-value pairs that have been inserted into the IBLT. Algorithm 4 describes the simplest form of the recovery operation. This algorithm consists of seeking cells with *count* field equal to one. Such cells contain a single element  $z$  of

$\mathcal{S}$ , which can be directly obtained from their *data* field. Thus, one can extract the element  $z$  and append it to the output list. After that, one deletes  $z$  from the IBLT by calling Delete( $z$ ), which effectively removes  $z$  from all the cells in which it had been mapped and reduces the count fields by one. This process is repeated until no more cells with *count* equal to one can be found. Recovery succeeds if at the end all  $m$  cells of the IBLT have *count* field equal to zero. If some cells have a non-zero count, recovery fails.

As stated in [21], the recovery operation is an instance of *peeling decoding* [22]. Peeling decoding is best understood considering a graph representation of an IBLT. In particular, we may represent an IBLT with  $m$  cells in which a set  $\mathcal{S}$  of  $n$  elements (key-value pairs) have been inserted as a bipartite (or Tanner) graph  $\mathcal{G} = (\mathcal{Z} \cup \mathcal{C}, \mathcal{E})$  composed of a set of  $n$  data nodes  $\mathcal{Z}$ , a set of  $m$  cell nodes  $\mathcal{C}$  and a set of edges  $\mathcal{E}$ . As the names indicate, data nodes represent key-value pairs and cell nodes represent cells of the IBLT. A data node  $z_i \in \mathcal{Z}$  and a cell node  $c_h \in \mathcal{C}$  are connected by an edge if and only if  $z_i = (x_i, y_i)$  is mapped into cell  $c_h$ . A data node  $z$  and a cell node  $c$  are said to be neighbors if they are connected by an edge. We use the shorthand  $\mathcal{N}(c)$  or  $\mathcal{N}(z)$  to denote the set of all neighbors of a cell node  $c$  or a data node  $z$ . The degree of a node is given by the number of edges connected to the node. Thus, the degree of a cell node equals the *count* field of the cell it represents.

**Example 2** (Graph representation). *Figure 1 shows the graph representation of an IBLT with  $m = 4$  cells in which  $n = 3$  pairs have been inserted. Key-value pair  $z_1$  is mapped into cells  $c_1$  and  $c_4$ , key-value pair  $z_2$  is mapped into cells  $c_1$  and  $c_2$  and key-value pair  $z_3$  is mapped into cells  $c_2$  and  $c_4$ . Thus, we have  $\text{count} = 2$  at cells  $c_1, c_2$ , and  $c_4$ , and  $\text{count} = 0$  in  $c_3$ .*

*In the graph representation of the IBLT, each key-value pair  $z$  is represented by a data node  $z$ , shown as a circle in the figure. Each cell  $c$  is represented by a cell node  $c$ , shown as a square. An edge connects a data node  $z_i$  to a cell node  $c_h$  if and only if  $z_i$  is mapped into cell  $c_h$ . For example,  $z_1$  is connected to  $c_1$  and  $c_4$ . We can also observe that the degree of a cell node  $c$  (the number of edges attached to it) equals the count field of the cell it represents. Thus, cell nodes  $c_1, c_2$ , and  $c_4$  have degree 2, whereas  $c_3$  has degree zero.*

Recovery of  $\mathcal{S}$  can be represented as a peeling process on a bipartite graph. In particular, whenever a cell node  $c$  of degree one is present, its only neighbor  $\mathcal{N}(c) = z$  is determined. The key-value pair  $z$  which is represented by the data node  $z$  is added to the output list. Next, the retrieved key-value pair is deleted from the IBLT. This translates into the removal of all edges attached to its associated data node. This process is repeated until no more cell nodes of degree one are present. At this stage, if all cell nodes are of degree zero, recovery succeeded and all key-value pairs are present in the output list. Otherwise, if some cell nodes of a degree larger than zero are present, recovery fails, and the output list will not contain all key-value pairs. Note that the graph structure is unknown to the decoder, and is revealed successively during the decoding process (see Examples 3 and 4).




Figure 1: Graph representation of an IBLT. The data nodes  $z$  are represented by circles and the cell nodes  $c$  by squares. An edge connects a data node  $z_i$  to a cell node  $c_h$  if and only if  $z_i$  is mapped into cell  $c_h$ . The degree of a cell node  $c$  equals the *count* field of the cell it represents. The operator  $\oplus$  denotes a bit-wise XOR.

**Example 3** (Peeling decoding I). *The different steps of the peeling process are shown in Figure 2. Figure 2a shows the bipartite graph representation of an IBLT before the peeling process starts. We observe that the IBLT has  $m = 5$  cells and it stores  $n = 4$  key-value pairs. However, at this stage, the depicted bipartite graph is unknown to the decoder since it does not have any knowledge about  $S$ . The decoder is only aware of the  $m$  cell nodes. For this reason, the data nodes, as well as the edges, are shown in gray. The graph structure will be revealed successively as the recovery operation progresses, and it will only be completely known if decoding succeeds. Otherwise, a part of the graph will remain hidden. We can see that cell node  $c_3$  has degree 1, and thus its associated IBLT cell  $c_3$  has count 1. The recovery operation retrieves the only key-value pair that has been mapped to cell  $c_3$ , i.e., data node  $z_2$ , which is added to the output list of the recovery operation. Afterwards,  $z_2$  is deleted from the IBLT. In the graph representation, this translates to revealing the only neighbor of cell node  $c_3$ , data node  $z_2$  (now shown in black), and deleting all edges attached to it, as shown in Figure 2b. As a consequence, the degree of  $c_1$  becomes one. In the next step, as shown in Figure 2c, the only neighbor of  $c_1$ ,  $z_1$ , is revealed and all edges attached to it are removed. This reduces the degree from  $c_4$  from 2 to 1. Then, data node  $z_4$  is revealed since it is the only neighbor of  $c_4$ . After all the edges attached to  $z_4$  are removed, as shown in Figure 2d, we have two cell nodes of degree 1, namely  $c_2$  and  $c_5$ , both of which have as only neighbor  $z_3$ . In the last step shown in Figure 2e,  $z_3$  is revealed as the only neighbor of  $c_2$ . Finally, all edges attached to  $z_3$  are erased from the graph. In this example, the recovery operation succeeded, and set  $S$  was completely recovered.*

**Example 4** (Peeling decoding II). *Consider now the application of peeling decoding to the graph in Figure 1. In contrast to Example 3, peeling decoding now fails. It cannot even get started since there are no cells with count 1, i.e., there are no cell nodes of degree 1.*




Figure 2: Peeling process on the graph representation of an IBLT with  $n = 4$  key-value pairs and  $m = 5$  cells. The index  $\iota$  is used to denote the different steps of the recovery process. The data nodes and edges shown in gray are those which, at the current iteration, are still unknown to the decoder. The data nodes shown in black have already been recovered by the decoder at the end of the current iteration  $\iota$ . The arrows parallel to the edges indicate the traversal direction of the peeling process.

### III. ANALYSIS

We consider the case where the number of cells  $m$  grows large and are interested in the probability that a random key-value pair cannot be recovered given that  $n \leq m$  key-value pairs were previously stored in the MET IBLT. Our analysis is based on MET density evolution [20] that provides a general framework to obtain the asymptotic performance of LDPC code ensembles. In the asymptotic regime, MET density evolution allows to determine the *average* probability of a key-value pair being present in the output list of the recovery operation, where the average is taken over an ensemble of bipartite graphs. Even while fixing the IBLT parameters and hash functions, different graphs are obtained by inserting different sets  $S$  into the IBLT. Therefore, we are interested in the average probability of recovery failure.

#### A. MET IBLT description

We associate to each cell node a node type  $i \in \{1, 2, \dots, d_c\}$ . Cell nodes of the same type may have *different*

degrees. We introduce the edge-oriented degree distribution polynomial for cell nodes of type  $i$

$$\rho_i(x) = \sum_{k=1} \rho_{i,k} x^{k-1}$$

where the coefficients  $\rho_{i,k}$  denote the number of edges emanating from cell nodes of type  $i$  and degree  $k$  divided by all edges emanating from cell nodes of type  $i$ . As we will shortly see,  $\rho_i(x)$  is induced by the node degree and type distributions of the data nodes.

Likewise, each data node is associated with a data node type  $j \in \{1, 2, \dots, d_d\}$ . We use a  $d_c \times d_d$  degree matrix  $\mathbf{D} = [d_{i,j}]$  to describe the connectivity between the different types of data and cell nodes, where  $d_{i,j}$  corresponds to the number of edges connecting a type- $j$  data node to a type- $i$  cell node (recall Section II). Considering only edges connected to cell nodes of type  $i$  (i.e., row  $i$  of  $\mathbf{D}$ ), we define the expected data node degree  $\bar{d}_i$  from the perspective of a cell node of type  $i$  as

$$\bar{d}_i = \sum_{j=1}^{d_d} p_j d_{i,j}.$$

Further, considering still row  $i$  of  $\mathbf{D}$  only, we denote by  $\lambda_{i,j}$  the fraction of edges connected to data nodes of type  $j$ ,  $\lambda_{i,j} = \frac{p_j d_{i,j}}{\bar{d}_i}$ .

Hence,  $\boldsymbol{\lambda}_i = [\lambda_{i,j}]$ ,  $j \in \{1, 2, \dots, d_d\}$  gives an edge-oriented data node degree distribution considering only edges connected to type  $i$  cell nodes (and ignoring all other edges).

**Example 5.** Given the degree matrix  $\mathbf{D} = \begin{bmatrix} 2 & 0 & 1 \\ 2 & 3 & 3 \end{bmatrix}$  and the probability vector  $\mathbf{p} = [0.25 \ 0.25 \ 0.5]$ . We find that there are three different data node types and two different cell node types. With probability 1/2 a data node is of type three while with probability 1/4 it is of type one or two, respectively. The average data node degrees from the cell node perspective are  $\bar{d}_1 = 1.00$  and  $\bar{d}_2 = 2.75$ . Further, the data node edge-oriented degree distributions from the cell node perspective are  $\boldsymbol{\lambda}_1 = [0.50 \ 0 \ 0.50]$  and  $\boldsymbol{\lambda}_2 = [0.18 \ 0.27 \ 0.55]$ .

**Example 6.** Assume a MET IBLT with the parameters from Example 5. In addition, let  $n = 4$ ,  $m = 5$ , and  $m_1 = 2$ ,  $m_2 = 3$ . Figure 3 exemplifies a bipartite graph with such parameters.<sup>4</sup>

### B. Load threshold computation

**Definition 5 (Load).** Let the load  $\eta = n/m$ , be the ratio between the number of key-value pairs and cells used to store those, i.e., the length of the IBLT.

Let us consider the regime for which  $n$  and  $m$  go to infinity. We are interested in the maximum load, the so-called load threshold  $\eta^*$  of a MET IBLT ensemble, such that recovery is successful with high probability. In the context of LDPC codes, the threshold of code ensembles with peeling decoding is analyzed via *density evolution* [22], [23] which restates the

<sup>4</sup>The connections are data dependent and are chosen at will in this example, respecting the type constraints.




Figure 3: Bipartite graph example of a MET IBLT with  $n = 4$  data nodes and  $m = 5$  cell nodes of different types.

peeling decoder as an *equivalent* iterative message passing algorithm where nodes pass messages along the edges to their neighbors. The messages exchanged by the nodes can be either an *erasure*, i.e., we do not know the corresponding key-value pair yet, or a *non-erasure*, meaning that key-value pair is known.

Let us denote by  $\bar{w}_i^{(\iota)}$  the (average) probability that the message sent from a cell node of type  $i$  over an edge at the  $\iota$ -th iteration is an erasure and by  $\bar{q}_i^{(\iota)}$  the (average) probability that the message sent from a data node to a cell node of type  $j$  over an edge at the  $\iota$ -th iteration is an erasure. MET density evolution iteratively computes  $\bar{w}_i^{(\iota)}$  and  $\bar{q}_i^{(\iota)}$ .

As a first step in our analysis, we need the edge perspective degree distribution for cell nodes of type  $i$  in the limit when  $m \rightarrow \infty$ . As shown in [24], it is

$$\rho_i(x) = e^{-\frac{\eta}{f_i} \bar{d}_i (1-x)}$$

where  $f_i \triangleq m_i/m$  is the fraction of cell nodes of type  $i$  and  $\eta/f_i = n/m_i$  is the load for the  $m_i$  cell nodes of type  $i$ . The message sent by a cell node of type  $i$  will only be a non-erasure if all other incoming messages are non-erasures, i.e.,

$$\bar{w}_i^{(\iota)} = 1 - \rho_i \left( 1 - \bar{q}_i^{(\iota)} \right) = 1 - e^{-\frac{\eta}{f_i} \bar{d}_i \bar{q}_i^{(\iota)}}. \quad (1)$$

In other words, if all but one key-value pairs mapped into a cell are known, also the remaining key-value pair can be determined (non-erasure message). Let  $q_{i,j}^{(\iota)}$  be the erasure probability for messages sent from a type- $j$  data node to a type- $i$  cell node at the  $\iota$ -th iteration. The message sent by a data node will only be an erasure if all other incoming messages are erasures, i.e.,

$$q_{i,j}^{(\iota)} = \prod_{k=1}^{d_c} \left( \bar{w}_k^{(\iota-1)} \right)^{b_{k,j}} \quad (2)$$

where  $b_{k,j} = d_{k,j}$  if  $k \neq i$  else  $b_{k,j} = \max(0, d_{k,j} - 1)$ . In other words, if a key-value pair has been recovered from any cell, its value can be subtracted from all other cells where it has been previously mapped into. The average erasure probability

at the input of a cell node of type  $i$  at the  $\iota$ -th iteration  $\bar{q}_i^{(\iota)}$  is given by

$$\bar{q}_i^{(\iota)} = \sum_{j=1}^{d_d} \lambda_{i,j} q_{i,j}^{(\iota)}. \quad (3)$$

By initially setting  $\bar{q}_i^{(0)} = 1$  for the 0-th iteration, and iteratively computing (1),(2),(3) it is possible to derive the probability that the messages passed along the edges of the graph are erasures as the number of iterations  $\iota$  grows.

Let us define by  $\gamma_j^{(\iota)}$  the probability that a data node of type  $j$  is erased at the end of the  $\iota$ -th iteration. We have,

$$\gamma_j^{(\iota)} = \prod_{i=1}^{d_c} (\bar{w}_i^{(\iota)})^{d_{i,j}}.$$

The load threshold of a MET IBLT ensemble can be formally defined in terms of  $\gamma_j^{(\iota)}$  as follows.

**Definition 6** (Load threshold). *The load threshold  $\eta^*$  is the largest value of  $\eta$  such that  $\gamma_j^{(\iota)} \rightarrow 0$  for all  $j \in \{1, 2, \dots, d_d\}$  when  $m \rightarrow \infty$  and  $\iota \rightarrow \infty$ .*

Intuitively, in the regime of  $m \rightarrow \infty$ , the probability that a random key-value pair is successfully recovered from the IBLT tends to one, as long as the number  $n$  of key-value pairs previously inserted in the IBLT is below  $\eta^*m$ .

### C. Numerical Examples

We report load thresholds complemented by finite-length simulations of the probability  $P_e$  that a randomly chosen key-value pair is not present in the output list of the recovery operation. For illustration purposes, we consider two different MET IBLT designs with different load thresholds  $\eta^*$ .

**Example 7.** Consider a design  $E_1$  with 3 different data and cell node types,  $d_d = d_c = 3$ , a probability vector  $\mathbf{p} = [0.2 \ 0.2 \ 0.6]$ ,  $\mathbf{m} = [m/3 \ m/3 \ m/3]$ , and a degree matrix

$$\mathbf{D} = \begin{bmatrix} 1 & 2 & 1 \\ 2 & 1 & 1 \\ 1 & 2 & 1 \end{bmatrix}.$$

We obtain  $\eta^* = 0.815$ .

**Example 8.** For design  $E_2$ , we have  $d_d = 4$ ,  $d_c = 2$ ,  $\mathbf{p} = [0.046 \ 0.427 \ 0.398 \ 0.129]$ ,  $\mathbf{m} = [m/2 \ m/2]$ , and

$$\mathbf{D} = \begin{bmatrix} 6 & 3 & 1 & 4 \\ 14 & 0 & 2 & 6 \end{bmatrix}.$$

We obtain  $\eta^* = 0.935$ .

Figure 4 depicts  $P_e$  as a function of  $\eta = n/m$  for different  $m$  considering both MET IBLTs designs  $E_1$  and  $E_2$ . Observe that for increasing  $m$ ,  $P_e$  shows a sharp drop at  $\eta \approx \eta^*$  where the load threshold  $\eta^*$  is indicated by a vertical line. Coarsely speaking, for large  $m$  one can store approximately  $\eta^*m$  key-value pairs in the MET IBLT. The probability that the recovery of a randomly selected key-value pair fails goes to zero. Observe that design  $E_2$  allows storing more key-value pairs than design  $E_1$  in an IBLT of a given size  $m$  (for  $m$




Figure 4: Probability of unsuccessful recovery of a key-value pair,  $P_e$ , as a function of the channel load  $\eta$  for two different MET IBLT designs and different values of  $m$ .

large enough). Note that design  $E_2$  was obtained by running a computer search to improve the load threshold. In Section V we will build on the analysis introduced earlier in this section to find good MET designs suitable for a set reconciliation application.

## IV. SET RECONCILIATION PROTOCOL

In this section, we explain how IBLTs can be used for set reconciliation. We describe a set reconciliation protocol referred to as rate-compatible protocol which relies on the presented MET IBLTs.

### A. Set Reconciliation with IBLTs

Following [9], let us first illustrate how IBLTs can be used for set reconciliation. First, hosts  $A$  and  $B$  insert all the elements of their sets  $\mathcal{S}_A$  and  $\mathcal{S}_B$  into the IBLTs  $\mathbf{c}^{\{A\}}$  and  $\mathbf{c}^{\{B\}}$ , respectively. The mapping rules for both IBLTs  $\mathbf{c}^{\{A\}}$  and  $\mathbf{c}^{\{B\}}$  are the same and are agreed upon beforehand by the hosts. Then, host  $A$  sends  $\mathbf{c}^{\{A\}}$  to host  $B$  that creates a so-called difference IBLT  $\mathbf{c}^{\{\Delta\}}$ . This is done by subtracting all the cells of  $\mathbf{c}^{\{A\}}$  from those of  $\mathbf{c}^{\{B\}}$  applying Algorithm 5 (which simply subtracts the *count* field of the two cells, and xors the *data* field).

We make the following observations regarding the difference IBLT,  $\mathbf{c}^{\{\Delta\}}$ :

- The elements in  $\mathcal{S}_A \cap \mathcal{S}_B$ , the intersection of  $\mathcal{S}_A$  and  $\mathcal{S}_B$ , are not present in the difference IBLT  $\mathbf{c}^{\{\Delta\}}$  (it is as if they had been first inserted and then deleted).
- The elements in the set difference,  $\Delta = \mathcal{S}_A \cup \mathcal{S}_B \setminus \mathcal{S}_A \cap \mathcal{S}_B$  are still present in the difference IBLT. In particular:
  - The elements in  $\mathcal{S}_B \setminus \{\mathcal{S}_A \cap \mathcal{S}_B\}$ , those present in  $\mathcal{S}_B$  but not in  $\mathcal{S}_A$ , have been inserted into  $\mathbf{c}^{\{\Delta\}}$  (and they have not been yet deleted).
  - The elements in  $\mathcal{S}_A \setminus \{\mathcal{S}_A \cap \mathcal{S}_B\}$ , those present in  $\mathcal{S}_A$  but not in  $\mathcal{S}_B$ , have been deleted from  $\mathbf{c}^{\{\Delta\}}$ , without having been inserted before.

- Thus, if we represent the difference IBLT as a bipartite graph, we obtain a graph with  $m$  cell nodes, and  $\delta = |\Delta|$  data nodes (only the elements of the set difference  $\Delta$  are present). Furthermore, the degree and connectivity of the remaining data nodes is preserved compared to the bipartite graph of the IBLTs of host  $A$  and/or  $B$  (see Example 9).

**Example 9 (Difference IBLT).** For ease of illustration, we consider an irregular IBLT with 5 cells in which host  $A$  has the set  $\mathcal{S}_A = \{z_1, z_2, z_3\}$ , whereas host  $B$  has  $\mathcal{S}_B = \{z_3, z_4\}$ , hence the set difference is  $\Delta = \{z_1, z_2, z_4\}$ . Figures 5a and 5b, show the bipartite representation of  $c^{\{A\}}$  and  $c^{\{B\}}$  respectively, as well as the values stored in the cells of the IBLTs. The bipartite graph associated with  $c^{\{\Delta\}}$ , which is obtained by subtracting all cells of  $c^{\{A\}}$  from those of  $c^{\{B\}}$ , is shown in see Figure 5c. Observe that there are no edges connected to the elements in the intersection of the sets,  $\mathcal{S}_A \cap \mathcal{S}_B = \{z_3\}$ .

For illustration purposes, solid edges are attached to elements that are only in  $\mathcal{S}_B$ , i.e., in  $\mathcal{S}_B \setminus \mathcal{S}_A \cap \mathcal{S}_B$ . Whenever a data node is connected to a cell node by a solid edge, the count field in the associated IBLT cell will be increased by one. Dashed edges are attached to elements that are only in  $\mathcal{S}_A$ , i.e., in  $\mathcal{S}_A \setminus \mathcal{S}_A \cap \mathcal{S}_B$ . Whenever a data node is connected to a cell node using a dashed edge, the count field in the associated IBLT cell will be decreased by one. By counting the number of solid and dashed edges we can derive the count field associated with each cell in  $c^{\{\Delta\}}$ . Note however that the count of the cells in  $c^{\{\Delta\}}$  no longer coincides with the degree of the associated cell nodes. For example, observing Figure 5c we can see that in cell  $c_3$  we have  $\text{count} = 0$ , although actually, its associated cell node has degree 2, i.e., we have  $\text{data} = z_2 \oplus z_4$ .

In order to be able to recover the set difference  $\Delta = (\mathcal{S}_A \cup \mathcal{S}_B) \setminus (\mathcal{S}_A \cap \mathcal{S}_B)$  from  $c^{\{\Delta\}}$ , the recovery algorithm must be altered to deal with elements that have been deleted from the IBLT without having been previously inserted. Let us recall that the standard recovery algorithm (Algorithm 4) searches for cells with  $\text{count} = 1$ , because in such cells a single set element has been inserted, and thus it can be directly recovered from the  $\text{data}$  field. However, the modified recovery algorithm (Algorithm 7) needs also to account for elements that have been deleted from the IBLT without having been previously inserted (elements in  $\mathcal{S}_A \setminus \{\mathcal{S}_A \cap \mathcal{S}_B\}$ ). In particular, the modified recovery algorithm has to search for cells where a single element has been *inserted or deleted*. Following [9], we will refer to such cells as *pure*. When a cell is pure, the  $\text{data}$  field of the cell contains a copy of the single element that has been inserted/deleted.

**Example 10 (Purity).** In Figure 5c, we can see that cell  $c_1$  has  $\text{count} = -1$ , but it is not pure, since 3 key-value pairs have been mapped into it,  $z_1, z_2$ , and  $z_4$ . By contrast, cells  $c_2$  and  $c_4$  which have respectively  $\text{count} = -1$  and  $1$  are pure since only one key-value pair has been mapped into them.

An issue consists in detecting pure cells, keeping in mind that the decoder only has access to the IBLT cells and the

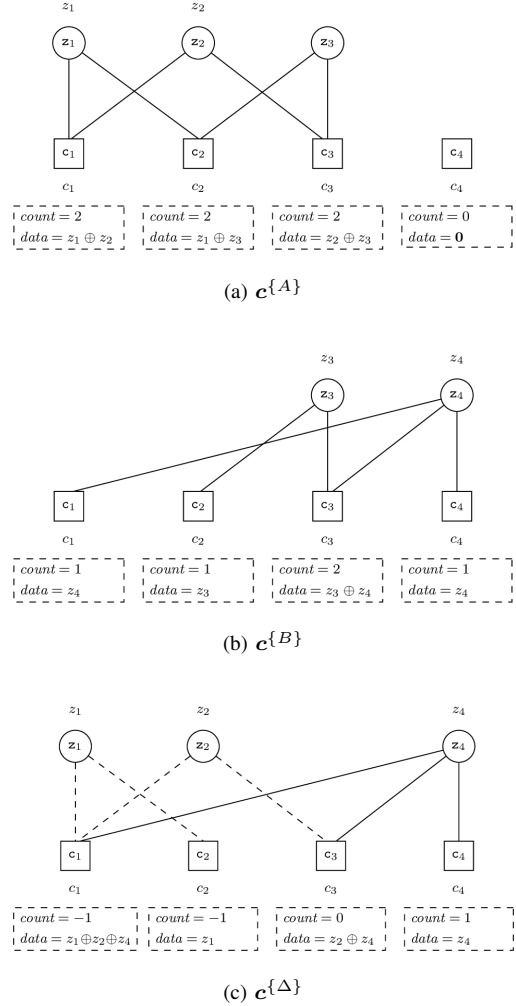


Figure 5: Example of graph representation of the IBLTs generated locally by each of the hosts. The local IBLT of host  $A$  (transmitter),  $c^{\{A\}}$ , is shown in (a), and that of host  $B$ ,  $c^{\{B\}}$ , is shown in (b). The difference IBLT is illustrated in (c).

mapping rules, but the graph structure is *hidden*. It is easy to see how  $\text{count} = \pm 1$  is a necessary but not sufficient condition for a cell to be pure. Thus, altering the recovery algorithm to search for cells with  $\text{count} = \pm 1$  is not sufficient. To distinguish pure cells, one can rely on the fact that the keys are obtained from the values by using a function  $g(\cdot)$ , [9]. Hence, in addition to checking whether  $\text{count} = \pm 1$ , we also check if the key  $z.x$  is valid, i.e., whether  $z.x = g(z.y)$  (see Algorithm 6).<sup>5</sup> Note however that the function  $g(\cdot)$  is a many-to-one mapping. Thus, this purity check may return true for an impure cell (see Section IV-C).

Relying on this probabilistic purity check, it is possible to define a modified recovery algorithm (see Algorithm 7) to recover the set difference from the difference IBLT  $c^{\{\Delta\}}$ . This algorithm searches for pure cells. Whenever a pure cell

<sup>5</sup>In fact the condition  $z.x = g(z.y)$  renders the count field in the cells unnecessary for set reconciliation applications. Still, we keep it in our discussion for the sake of clarity.



---

**Algorithm 5** Cell Subtraction

---

```
procedure  $c^\dagger = \text{SUBTRACT}(c, c')$   
   $c^\dagger.count = c.count - c'.count$   
   $c^\dagger.data = \text{XOR}(c.data, c'.data)$ 
```

---

---

**Algorithm 6** Check Purity

---

```
procedure  $b = \text{ISPURE}(z)$   
  if  $z.x = g(z.y)$  then  
     $b = \text{true}$   
  else  
     $b = \text{false}$ 
```

---

---

**Algorithm 7** Modified Recovery

---

```
procedure  $\text{RECOVER}()$   
  while  $\exists i \in [1, m] \mid c_i.count = \pm 1 \ \& \ \text{IsPure}(c_i.data) = \text{true}$  do  
    add  $z = c_i.data$  to the output list  
    if  $c_i.count = 1$  then  
      call  $\text{Delete}(z)$   
    else  
      call  $\text{Insert}(z)$ 
```

---

is found, the only element mapped in the cell is added to the output list, and then this element is removed from the IBLT. This is achieved by either deleting it from the IBLT in case we have  $count = 1$  or adding it to the IBLT in case we have  $count = -1$ .

**Example 11** (Modified Recovery). *Let us consider modified recovery applied to the difference IBLT shown in Figure 5c. For simplicity, we assume here that the purity check is perfect. Under this assumption, cell  $c_2$ , which has  $count = -1$  can be determined to be pure, allowing to recover  $z_1$ , which is removed from the IBLT by relying on the addition operation. Similarly, cell  $c_4$  is also found to be pure. This allows recovering  $z_4$ , which is removed from the IBLT by calling the delete operation since we had  $count = 1$  in  $c_4$ . Finally,  $z_2$  can be recovered from either  $c_1$  or  $c_3$ .*

Fig. 6a shows the sequence diagram of difference digest, the IBLT based set reconciliation protocol proposed in [9]. For simplicity, we have assumed here that the size of the set difference  $\delta$  is known a priori. As it can be observed, after a request from node B, node A transmits its complete IBLT  $c^{\{A\}}$ . After receiving  $c^{\{A\}}$ , node B computes the difference IBLT. Next, the modified recovery operation is used to attempt reconciling the sets. Finally, node B sends a (negative) acknowledgement to node A to communicate whether the reconciliation was successful or not.

### B. Rate-Compatible Set Reconciliation Protocol

We propose a novel rate-compatible protocol that relies on MET IBLTs. Consider a set reconciliation setting with two hosts, A and B, which possess sets  $\mathcal{S}_A$  and  $\mathcal{S}_B$ , respectively. The hosts communicate with each other through a data network to reconcile their sets, and it is assumed that the protocol parameters, i.e., the parameters that define the IBLTs known to both parties. The protocol follows the following steps:

- 1) Hosts A and B initialize to zero a MET IBLT of length  $m$  cells, with  $d_c$  different cell types and  $d_d$  key-value pair types. Next, each host inserts all the

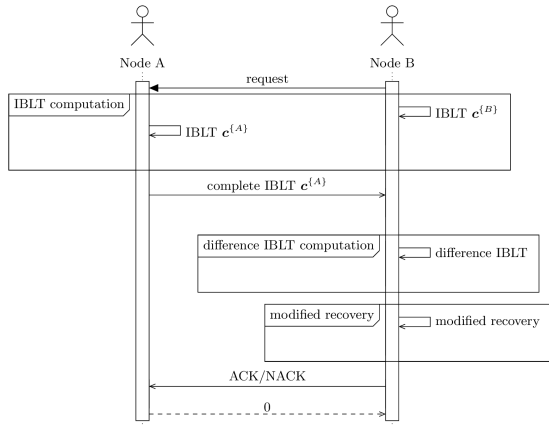
elements of its set into the MET IBLT. The thus obtained IBLTs of A and B will be denoted  $c^{\{A\}}$  and  $c^{\{B\}}$ , respectively. We assume that the cells of  $c^{\{A\}}$  and  $c^{\{B\}}$  are ordered according to their type so that the first  $m_1$  cells correspond to cells of type 1, followed by  $m_2$  cells of type 2, and so on.

- 2) Host A starts sending the cells of its IBLT  $c^{\{A\}}$  one by one, starting with cell  $c_1$ , followed by  $c_2$ , and so on. Host B creates an empty array of IBLT cells  $c^{\{\Delta_0\}}$ .
- 3) When receiving the  $i$ -th cell of  $c^{\{A\}}$ , host B subtracts it from the  $i$ -th cell of its local IBLT  $c^{\{B\}}$  (see Algorithm 5), and it appends the resulting cell to the difference IBLT  $c^{\{\Delta_{i-1}\}}$  to obtain  $c^{\{\Delta_i\}}$ . The cell subtraction operation consists of subtracting the  $count$  fields of the cells and XOR-ing the  $data$  fields.
- 4) Host B runs a modified recovery algorithm on  $c^{\{\Delta_i\}}$  which, if successful, yields at its output the set difference  $\Delta$  between  $\mathcal{S}_A$  and  $\mathcal{S}_B$  (see Algorithm 7).
- 5) If set reconciliation succeeds, host B sends an acknowledgment to host A. Otherwise, host B does not send any acknowledgment, and host A continues sending IBLT cells.
- 6) If recovery is still unsuccessful after all  $m$  cells have been transmitted, a failure is declared.


Fig. 6b shows the sequence diagram of the proposed rate-compatible protocol. In this sequence diagram assume that IBLT cells are not sent one by one, but in groups of  $h$  cells. After having received the first  $h$  cells, node B attempts to reconcile the sets by running the modified recovery operation, which fails. This leads to the transmission of an *implicit NACK*. Here, we say implicit because, in practice, rather than sending a NACK, node A could simply assume that set reconciliation did not yet succeed whenever no acknowledgement is received after a predefined amount of time. Next, Node A sends additional IBLT cells to node B. Node B then re-attempts set reconciliation by running the modified recovery operation, which now succeeds. Finally, Node B sends an acknowledgement to Node A and the protocol ends. We make the following observations. First, when the protocol succeeds, only host B is aware of the set difference. This is a common assumption in set reconciliation protocols. If necessary, host B can then transmit a message to host A to convey the set difference. Second, the protocol will fail if recovery is still unsuccessful after having transmitted all  $m$  cells. However, as illustrated in Section V, one of the advantages of relying on MET IBLTs is that  $m$  may grow on-demand, allowing in principle to generate an unlimited number of cells so that reconciliation is eventually successful.

**Example 12** (Protocol). *Figures 7a and 7b show the graphical representation of the local IBLTs of host A and B ( $c^{\{A\}}$  and  $c^{\{B\}}$ ), respectively. Both bipartite graphs have a total of  $m = 4$  IBLT cell nodes, out of which 2 are of type 1 and 2 of type 2, i.e.,  $m_1 = m_2 = 2$ . Furthermore, we have  $D^T = \begin{bmatrix} 1 & 1 \end{bmatrix}$ , i.e., there is a single data node type. We have  $\mathcal{S}_A = \{z_1, z_2, z_3\}$  and  $\mathcal{S}_B = \{z_2, z_3, z_4\}$ . Thus we have that the set difference is  $\Delta = \{z_1, z_4\}$ .*

*According to the rate-compatible protocol, host A sends its*



(a) Difference digest



(b) Rate-compatible set reconciliation

Figure 6: Sequence diagrams of difference digest and the proposed rate-compatible protocol. For simplicity, in difference digest we assume that the size of the set difference  $\delta$  is known a priori. In the rate-compatible protocol, we assume that cells are sent in packets (or groups) of  $h$  cells.

local IBLT cells one by one to host B. After receiving the  $i$ -th cell, host B subtracts it from the  $i$ -th cell of its local IBLT. The thus obtained difference IBLT is denoted  $c^{\{\Delta_i\}}$ . Thus, for  $h \leq i$ , the  $h$ -th cell of  $c^{\{\Delta_i\}}$  is obtained by subtracting the  $h$ -th cell of  $c^{\{A\}}$  from the  $h$ -th cell of  $c^{\{B\}}$ . Moreover, we have no information about the  $h$ -th cell of  $c^{\{\Delta_i\}}$  for  $h > i$ . These cells are considered erased. After that, host B attempts to recover the set difference by running the modified recovery algorithm, which amounts to a peeling operation on the graph representation of  $c^{\{\Delta_i\}}$ . Figures 7c, 7d, and 7e show the

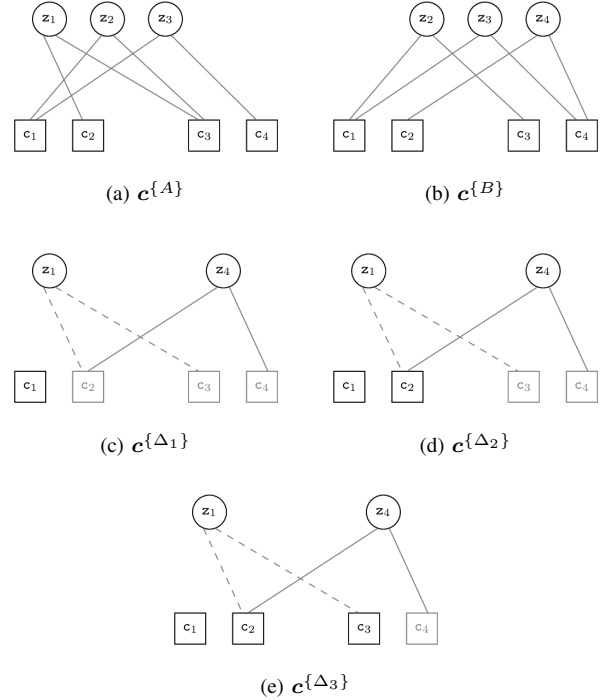


Figure 7: Set reconciliation example. Figures (a) and (b) show the graph representation of  $c^{\{A\}}$  and  $c^{\{B\}}$ . Figures (c), (d), and (e) show the graph representation of  $c^{\{\Delta_i\}}$ ,  $i = 1, 2, 3$ , i.e., after having node B has received 1, 2 and 3 cells respectively. The cell nodes shown in gray represent IBLT cells which have not yet been received and are thus treated as if they were erased (or punctured).

graph representation of  $c^{\{\Delta_i\}}$ , for  $i \in \{1, 2, 3\}$ . Note that the bipartite graph is provided for illustration purposes only. In the beginning, it is hidden to host B and it is successively revealed during the recovery operation. In particular, host B attempts set reconciliation by inverting  $c^{\{\Delta_i\}}$ , i.e., by carrying out the modified recovery operation (Algorithm 7), relying only on the  $i$  cells it has so far received, and treating IBLT cells  $i + 1$  onwards as if they had been erased. For this reason, when representing  $c^{\{\Delta_i\}}$  we mark cell nodes  $c_{i+1}, c_{i+2}, \dots$  in gray.

Observe that reconciliation is not successful after host B has received the first IBLT cell since  $c_1$  is empty in  $c^{\{\Delta\}}$  (and all other cells are still erased). Similarly, after having received 2 cells, reconciliation is still unsuccessful, since  $c_2$  is not pure. Finally, reconciliation is successful after having received 3 cells. In particular, cell  $c_3$  is pure, allowing to recover  $z_1$ , which is then removed from the IBLT. This renders cell  $c_2$  pure, allowing to recover  $z_4$ . Since set reconciliation is successful on  $c^{\{\Delta_3\}}$ , host B sends an acknowledgment to host A, and the reconciliation protocol terminates. Cell  $c_4$  is in this case never transmitted.

The set reconciliation protocol described above has the advantage of exhibiting a low complexity since it relies on the modified recovery operation, which is equivalent to a peeling

operation applied to the graph representation of the difference IBLT. One could argue that this is suboptimal since host  $B$  is not making use of his knowledge of  $\mathcal{S}_B$ . In fact, it is possible to improve the performance of the modified recovery operation (peeling) by relying on the knowledge of  $\mathcal{S}_B$ . For example, whenever the modified recovery operation (peeling) gets stuck because there are no pure cells, one could try to delete different elements of  $\mathcal{S}_B$  with the hope that they are in  $\Delta$  and that their deletion generates some extra pure cells allowing modified recovery (peeling). Such an approach was proposed in [25] and it was shown that it can improve performance at the cost of an increased complexity.

### C. Reconciliation Failures

In the protocol described above, two types of events can lead to undetected failures. The first type of failure is associated with the purity check (see Algorithm 6) returning true for a cell that is not pure. When this happens, an *erroneous* key-value pair will be added to the output list of the recovery algorithm. Furthermore, the removal of this erroneous key-value pair from the IBLT will corrupt a number of cells. With high probability, these corrupted cells will lead the receiver (host  $B$ ) to believe that there are still some key-value pairs that have not been recovered, and the set reconciliation will run until all the  $m$  cells have been transmitted, declaring then a reconciliation failure. Under the assumption that  $g(x)$  returns a random number between 1 and  $2^\kappa$ , we have that the purity check of an impure cell fails with probability  $2^{-\kappa}$ . Assuming a total of  $m$  cells, a set difference size  $\delta$ , and an average data node degree  $\bar{d}$ , the number of different purity checks that have to be carried out on impure cells can be upper bounded as  $m + \delta\bar{d}$ . Thus, we have that the probability of this type of failure can be upper bounded by  $(m + \delta\bar{d})2^{-\kappa}$ . Hence, by choosing  $\kappa$  to be large enough, this probability can be made arbitrarily small.

The other source of failure of the protocol is the fact that host  $B$  may mistakenly assume that all key-value pairs have been recovered, while this is not the case. This happens when, after having received  $i$  cells, the modified recovery (peeling decoding) of  $c^{\{\Delta_i\}}$  leads to an empty IBLT, however, there are one or more key-value pairs that have not yet been recovered. These yet unrecovered key-value pairs are not mapped to any of the first  $i$  cells, i.e., they have no edges to any of the first  $i$  cell nodes. Thus, there is no way for  $B$  to detect the existence of these unrecovered key-value pairs. The best way of dealing with this second type of failure is to never terminate the protocol before having received at least  $m_1$  cells, i.e., all the cells of type 1. Assuming that  $d_{1,j} > 0 \forall j$ , this effectively solves the problem since in this case all key-value pairs will be mapped at least to one of the first  $m_1$  cells.

### D. Protocol Variants

From a practical viewpoint, it might be of interest to change the insertion (encoding) and recovery (decoding) of MET IBLTs in Section II. It is not necessary to generate all  $m$  cells and to recover all  $m$  cells at once. Rather,

encoding and decoding can proceed in rounds similar to rate-compatible coding schemes (see e.g., [26]). We may define an encoding index vector  $[i_0 = 0 \ i_1 \ i_2 \ \dots]$  of positive integers  $0 < i_1 < i_2 \dots$ . In encoding round  $\alpha$  the encoder may generate  $i_\alpha - i_{\alpha-1}$  additional cell nodes. We may for instance choose  $i_\alpha - i_{\alpha-1} = m_\alpha$ . Likewise, we may define a decoding index vector  $[j_0 = 0 \ j_1 \ j_2 \ \dots]$  of non-zero integers  $0 < j_1 < j_2 \dots$ . Once  $j_\alpha$  cells are received, the decoder performs round  $\alpha$  of decoding. If the decoding attempt fails it waits for additional  $j_{\alpha+1} - j_\alpha$  cells and retries again. For instance, we may choose  $j_{\alpha+1} - j_\alpha = 1 \ \forall \alpha$ , i.e., we attempt decoding once receiving a cell.

## V. MET IBLT DESIGN

The rate-compatible protocol in Section IV requires a suitable MET IBLT design that provides good performance for a wide range of sizes of the set difference  $\delta$ , considering that  $\delta$  is not (perfectly) known a priori. A simple solution would be to perform an IBLT design for large  $\delta$  by relying on the results of Section III. However, when  $\delta$  is small, there is no guarantee that this scheme will work with reasonable communication complexity, i.e., that it will be efficient. To illustrate why this is the case, let us consider a  $k$ -regular<sup>6</sup> IBLT of length  $m$ . We focus on the graphical representation of the difference IBLT  $c^{\{\Delta_h\}}$  after  $h$  cells have been exchanged. Let us now consider a key-value pair node associated with the set difference  $\Delta$ . The probability that this node has no edges attached to any of the first  $h$  cell nodes can be approximated as

$$\left(\frac{m-h}{m}\right)^k. \quad (4)$$

Note that the expression in (4) yields a lower bound on the probability of recovery failure  $P_e$ . For fixed  $k$ , e.g.,  $k = 3$ , this probability is high, unless  $h$  is close to  $m$  (see also [27, Section 6] for a more detailed discussion). Thus, we expect the rate-compatible protocol to perform poorly when employing standard IBLTs. This is illustrated by simulation results in Section V-C. The issue can be overcome by MET IBLTs. In particular, we can choose  $m_1$  to be small, and forbid the protocol to terminate before having exchanged  $m_1$  cells. This ensures that every key-value pair in  $\Delta$  is mapped to a cell of type 1.

### A. Design Example

Consider a MET IBLT design in which the number of cells of type  $i$  is  $m_i = 2^{i-1}m_1$ . We use the tools in Section III to predict the asymptotic performance of the MET IBLT. For simplicity, we set  $d_d = 3$ , and consider a degree matrix in the form

$$\mathbf{D}^T = \begin{bmatrix} d_{1,1} & d_{2,1} & d_{3,1} & d_{4,1} & d_{5,1} & \cdots & d_{5,1} \\ d_{1,2} & d_{2,2} & d_{3,2} & d_{4,2} & d_{5,2} & \cdots & d_{5,2} \\ d_{1,3} & d_{2,3} & d_{3,3} & d_{4,3} & d_{5,3} & \cdots & d_{5,3} \end{bmatrix}$$

where the fifth and all successive rows of  $\mathbf{D}$  are equal to the fourth row. This yields a design in which a potentially

<sup>6</sup>in a  $k$ -regular IBLT all key-value pairs and cells are of the same type, and all key-value pairs are mapped exactly on  $k$  cells. Thus, we have  $p = p_1 = 1$ , and  $\mathbf{D} = k$ .

unlimited number of cell types  $d_c$  can be generated.<sup>7</sup> Let us now denote by  $\eta_i^*$  be the load threshold of the MET IBLT when considering only the first  $\sum_{j=1}^i m_j$  cells. The parameters of this design,  $\mathbf{p}$  and  $\mathbf{D}$ , are determined by relying on multi-target optimization. In particular, we choose  $\mathbf{p}$  and  $\mathbf{D}$ , to maximize the minimum among  $\eta_i^*$ ,

$$(\mathbf{p}, \mathbf{D}) = \arg \max_{\mathbf{p}, \mathbf{D}} \left( \min_i \eta_i^* \right).$$

An approximate solution to this optimization problem was obtained by relying on simulated annealing [28], an optimization algorithm to approximate the global optimum of a given function. For the current design example, we set  $i \in \{1, 2, \dots, 8\}$ . Additionally, we introduce the constraint  $d_{i,j} \leq 5 \forall i, j$  in order to limit complexity.<sup>8</sup> The optimization yields

$$\mathbf{p} = [0.1959 \quad 0.1904 \quad 0.6137]$$

and

$$\mathbf{D}^T = \begin{bmatrix} 3 & 1 & 1 & 1 & 1 & \dots & 1 \\ 4 & 4 & 4 & 4 & 5 & \dots & 5 \\ 2 & 1 & 1 & 1 & 1 & \dots & 1 \end{bmatrix}$$

and the load thresholds obtained are

$$\boldsymbol{\eta}^* = [0.7948, 0.7837, 0.7882, 0.8025 \\ 0.8042, 0.7967, 0.7895, 0.7856].$$

Although we only consider the first 8 cell types in our optimization, the resulting MET IBLT still retains a good performance when additional cell types are included by replicating the last row  $\mathbf{D}$ . In particular, we have  $\eta_9^* = 0.7842$ ,  $\eta_{10}^* = 0.7837$  and  $\eta_i^* = 0.7830$  for  $i \geq 11$ .

**Remark 1.** *This multi-target optimization only considers the performance at intermediate points in which all the cells of type  $i$  have been received. In the Appendix, we introduce an extension of the analysis in Section III that allows deriving the load threshold of a MET IBLT when only a fraction of cells of a certain type has been received.*

## B. Simulation Setup

We consider a setup with two hosts,  $A$  and  $B$ , who want to reconcile their sets  $\mathcal{S}_A$  and  $\mathcal{S}_B$  by communicating over a data network without having any knowledge about the size of the set difference  $\delta$ . We apply the set reconciliation protocol of Section IV-B with a MET IBLT from Section V and compare the performance with two other schemes available in the literature, difference digest [9] and characteristic polynomial interpolation (CPI) [8].

Difference digest [9] is an IBLT-based set reconciliation algorithm that relies on the use of regular IBLTs. In contrast to the other two schemes, difference digest requires an estimate of the set difference size, which is computed in an initial

communication round, in which host  $B$  sends to host  $A$  a data structure known as strata estimator which is used to obtain an estimate  $\hat{\delta}$  of the set difference size  $\delta$ . Host  $A$  then calculates a worst-case set difference size  $t = c \cdot \hat{\delta}$ , where  $c$  is chosen so that  $t > \delta$  with high probability (e.g., 0.99). Next, host  $A$  creates an IBLT of size  $m$ , where  $m$  is chosen so that the recovery operation is successful with high probability (e.g., probability 0.99) when  $t$  key-value pairs are inserted in the IBLT. After that, host  $A$  inserts its set  $\mathcal{S}_A$  in the IBLT  $\mathbf{c}^{\{A\}}$  and sends it to host  $B$ . Host  $B$  then also inserts his set  $\mathcal{S}_B$  in an IBLT  $\mathbf{c}^{\{B\}}$  of size  $m$  and then computes the difference IBLT by subtracting the cells of  $\mathbf{c}^{\{B\}}$  and  $\mathbf{c}^{\{A\}}$ , which yields the IBLT  $\mathbf{c}^{\{\Delta\}}$  that also has length  $m$ . Next, host  $B$  attempts to recover the set difference  $\Delta$  by applying a modified recovery operation (see Algorithm 7). If recovery is successful, set reconciliation succeeds, otherwise, we will simply assume that set reconciliation is declared as unsuccessful.<sup>9</sup>

The second scheme, CPI, makes use of the characteristic polynomial of a set. Consider a set  $\mathcal{S} = \{z_1, z_2, \dots, z_n\}$  where the elements  $z_i$  are chosen from  $\mathcal{U} = \{0, 1\}^\ell$ . The characteristic polynomial of set  $\mathcal{S}$  over the finite field  $\mathbb{F}_q$  with  $q > 2^\ell$  is defined as

$$\mathcal{X}_{\mathcal{S}}(Z) = (Z - z_1)(Z - z_2) \dots (Z - z_n)$$

where one relies on an injective mapping of elements of  $\mathcal{U}$  to  $\mathbb{F}_q$ . The key principle behind CPI is that given  $n$  evaluations of  $\mathcal{X}_{\mathcal{S}}(Z)$  it is possible to recover the set  $\mathcal{S}$  since its elements correspond to the zeros of the characteristic polynomial  $\mathcal{X}_{\mathcal{S}}(Z)$ . As we sketch next, CPI can be also operated in a rateless fashion [8] to minimize the communication cost and does not require an estimate of the set difference size. First of all, host  $A$  chooses a list of evaluation points  $r_1, r_2, \dots$  at which it will evaluate the characteristic polynomial of its set.<sup>10</sup> Next, host  $A$  continuously evaluates the characteristic polynomial of his set  $\mathcal{S}_A$ ,  $\mathcal{X}_{\mathcal{S}_A}(Z)$ , at the evaluation points,  $r_1, r_2, \dots$  and sends the evaluations  $\mathcal{X}_{\mathcal{S}_A}(r_1), \mathcal{X}_{\mathcal{S}_A}(r_2), \dots$  to host  $B$ , and it does so until it receives an acknowledgment from host  $B$  which indicates that reconciliation was successful. Thereby, host  $B$  continuously attempts to recover the set difference  $\Delta$ , by evaluating the ratio of two characteristic polynomials  $\mathcal{X}_{\mathcal{S}_A}(r_i)/\mathcal{X}_{\mathcal{S}_B}(r_i)$ . For a set difference size of  $\delta$  this protocol has overall complexity  $\mathcal{O}(d^4)$  [8].

## C. Simulation Results

We compare the performance of the three schemes in terms of communication cost, i.e., in terms of the amount of data, measured in bits, that needs to be exchanged. For illustration purposes, we consider the setup from [9] in which the cardinality of the  $\mathcal{S}_A$  and  $\mathcal{S}_B$  is  $10^5$ , and the set elements are chosen from  $\mathcal{U} = \{0, 1\}^{32}$ . Following [9], we allocate a total of 12 bytes for each IBLT cell: 4 bytes for the *count* field and 8 bytes for the *data* field (4 bytes for the key and

<sup>7</sup>Strictly speaking, the analysis presented in Section III only applies to the case in which the number of cell types  $d_c$  is constant. In practice, this does not hinder us from allowing the number of cell types to grow if needed.

<sup>8</sup>By removing the constraint our optimization algorithm yields only slightly higher thresholds.

<sup>9</sup>In [9] it was actually proposed to reattempt reconciliation by employing IBLTs of length  $2m$ . By allowing difference digest to fail we will slightly penalize the comparison with the two other schemes which continue until reconciliation is successful.

<sup>10</sup>In particular, the evaluations points  $r_1, r_2, \dots$  are chosen among those elements of  $\mathbb{F}_q$  to which no element of  $\mathcal{U}$  is mapped.

4 bytes for the value). For difference digest we use the same configuration as in [9], thus in the initial communication round to estimate the set difference size, we employ a hybrid strata estimator with 7 strata, each with 80 cells per IBLT, and a min-wise estimator with 2160 hashes. Hence, the resulting strata estimator (the data structure used to estimate the set difference size) has a total size of 15 kB. In the second communication round in which the actual set reconciliation takes place, we use a degree-3 regular IBLT. For CPI, we use a configuration in which every evaluation point takes 10 bytes.<sup>11</sup> Finally, for the rate-compatible protocol we use 12 bytes for each IBLT cell, the same as in difference digest, and the MET IBLT parameters reported in Section V. In all cases, we assume that the protocol configuration is known a priori to both reconciliation parties and does not need to be signaled. For example, for the rate-compatible protocol, this assumption translates into the knowledge of the parameters  $p$ ,  $D$ ,  $m$  as well as the hash functions  $H_{m,d_j}$ ,  $h_p$ , and  $g$ .

Figure 8 shows the communication cost, i.e., the total amount of transmitted data needed for set reconciliation for the three schemes. For difference digest, the figure also reports the total data transmission without considering the strata estimator (which takes exactly 15 kB in our example). For the rate-compatible protocol we show two performance curves. The first one uses the MET IBLT design described in Section V setting  $m_1 = 50$ . The second one uses a 3-regular distribution with  $m = 10^4$ .

Observe that the rate-compatible protocol when using a MET IBLT has only a slightly higher communication cost than the more complex CPI scheme, and it outperforms difference digest. This is remarkable since the complexity of our rate-compatible protocol, when relying on the design introduced in Section V, is  $\mathcal{O}(\delta \log(\delta))$ , whereas that of CPI and difference digest (ignoring the strata estimator) are  $\mathcal{O}(\delta^4)$  and  $\mathcal{O}(\delta)$ , respectively. This can be attributed to two different effects. The first one is the presence of an initial communication round to estimate the set difference size in difference digest. This inefficiency dominates the performance for small  $\delta$  since in this case, the size of the strata estimator is much larger than the actual data that has to be exchanged to reconcile the sets. The second one is the rate-compatible nature of the protocol which allows transmitting IBLT cells incrementally.

The performance of the rate-compatible protocol when using a degree 3 regular IBLT is also depicted. Observe that the performance is very poor compared to the MET IBLT design. This is due to the presence of key-value pairs that are not mapped to the first cells exchanged (see discussion in Section V) and hence cannot be recovered upon reception of those cells.

## VI. DISCUSSION

In this paper, we have presented a novel set reconciliation algorithm whose performance can approach that of characteristic polynomial interpolation (CPI) but with lower algorithmic

<sup>11</sup>We operate on a finite field with  $2^{40}$ . Hence, we need 5 bytes to specify the point at which we evaluate the polynomial, and another 5 bytes for the evaluation of the polynomial.

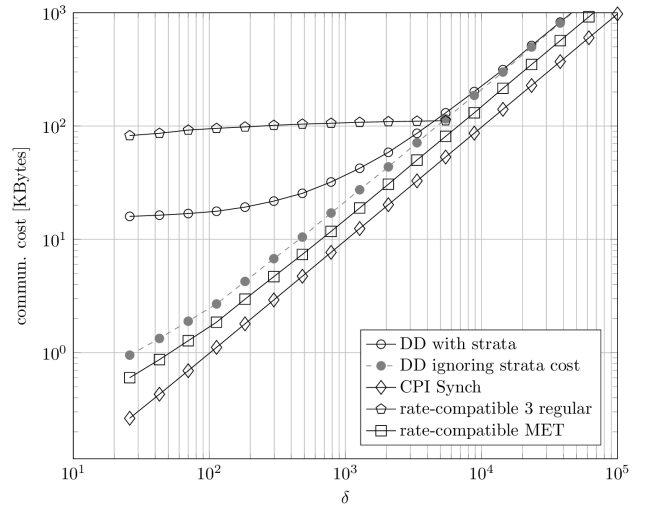


Figure 8: Data transmission required to reconcile sets with  $10^5$  elements as a function of the set difference size  $\delta$ . The solid line with circle markers represents the data transmitted by difference digest [9], including the strata estimator. The gray dashed line with circle markers represents the data transmitted by difference digest ignoring the size of the strata estimator for the computation of the communication cost. The solid line with diamond markers represents the CPI scheme in [8]. Finally, the solid line with square and pentagon markers represent the rate-compatible protocol introduced in this paper using the MET IBLT form Section V-A and a 3-regular IBLT.

complexity. The proposed scheme can be seen as an evolution of difference digest [9] with two major modifications. The first one is the concept (and analysis of) MET IBLTs, a generalization of the IBLT data structure, in which one allows for different types of IBLT cells and key-value pairs. The second modification is the fountain-like protocol, meaning that the number of IBLT cells to be exchanged is not fixed a priori. Instead, one can incrementally transmit as many cells as necessary, and new cells can in principle be generated on the fly.

The novel scheme presents some advantages compared to difference digest. It is not necessary to estimate the cardinality of the set difference which usually requires an additional communication round and introduces a substantial overhead. Likewise, we only have to transmit as many cells as necessary, instead of having to oversize the IBLT to cope with errors in the estimate of the cardinality of the set difference.

The proposed scheme combines the low complexity of difference digest and the efficiency of CPI, and could thus represent an appealing solution to many data synchronization problems in practical distributed systems. A possible application is the reconciliation of large databases. A concrete example is mempool synchronization in the Bitcoin network [6]. Another important application is remote file synchronization [3], [29], [30].

There are several of interesting extensions to this work. An example is the extension of the proposed scheme to a multi-party set reconciliation setting, as proposed for standard IBLTs

in [4]. Another example concerns a broadcasting setting in which a parent node holds a set  $\mathcal{S}_P$  and multiple child nodes hold sets  $\mathcal{S}_{C_1}$ ,  $\mathcal{S}_{C_2}$ , etc. In this setting, the set difference between  $\mathcal{S}_{C_i}$  and  $\mathcal{S}_P$  can be different for different  $i$ . The goal could be for the parent node to convey its set  $\mathcal{S}_P$  efficiently to the children nodes by transmitting IBLT cells over a broadcast channel. By relying on the proposed rate flexible constructions, a child node with a small set difference with respect to the parent node would obtain  $\mathcal{S}_P$  faster than a child node with a larger set difference. This extension could have applications in settings in which a parent node maintains a database and multiple child nodes maintain a local copy of the database.

#### ACKNOWLEDGEMENTS

The authors would like to thank Federico Clazzer for providing some software routines that were used for the Monte Carlo simulations.

The research leading to these results has been carried out under the framework of the project "SiNaKoL". The project started in January 2022 and is led by the Program Coordination Security Research within the German Aerospace Center (DLR), whose support we greatly appreciate.

#### REFERENCES

- [1] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels, "Dynamo: Amazon's highly available key-value store," *ACM SIGOPS operating systems review*, vol. 41, no. 6, pp. 205–220, 2007.
- [2] "Apache Cassandra." [Online]. Available: <https://cassandra.apache.org/>
- [3] S. Agarwal, V. Chauhan, and A. Trachtenberg, "Bandwidth efficient string reconciliation using puzzles," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 11, pp. 1217–1225, 2006.
- [4] M. Mitzenmacher and R. Pagh, "Simple multi-party set reconciliation," *Distributed Comput.*, vol. 31, no. 6, pp. 441–453, 2018. [Online]. Available: <https://doi.org/10.1007/s00446-017-0316-0>
- [5] P. Ozisik, G. Andresen, B. Levine, D. Tapp, G. Bissias, and S. Katkuri, "Graphene: Efficient interactive set reconciliation applied to Blockchain propagation," in *Proc. of Conf. of the ACM Special Interest Group on Data Commun.* Beijing, China: ACM, Aug. 2019, pp. 303–317.
- [6] N. Boškov, A. Trachtenberg, and D. Starobinski, "Gensync: A new framework for benchmarking and optimizing reconciliation of data," *IEEE Trans. Netw. Service Manag.*, 2022.
- [7] M. Karpovsky, L. Levitin, and A. Trachtenberg, "Data verification and reconciliation with generalized error-control codes," *IEEE Trans. Inf. Theory*, vol. 49, no. 7, pp. 1–2, 2003.
- [8] Y. Minsky, A. Trachtenberg, and R. Zippel, "Set reconciliation with nearly optimal communication complexity," *IEEE Trans. Inf. Theory*, vol. 49, no. 9, pp. 2213–2218, 2003.
- [9] D. Eppstein, M. Goodrich, F. Uyeda, and G. Varghese, "What's the difference?: Efficient set reconciliation without prior context," *ACM SIGCOMM Comp. Commun. Review*, vol. 41, no. 4, pp. 218–229, 2011.
- [10] M. Goodrich and M. Mitzenmacher, "Invertible Bloom lookup tables," in *Proc. of 49th Annual Allerton Conf. on Commun., Control, and Comp.* Monticello, IL, USA: IEEE, 2011, pp. 792–799.
- [11] P. Elias, "Error-free coding," *IRE Trans. Inf. Theory*, vol. PGIT-4, pp. 29–37, Sep. 1954.
- [12] M. Rink, "Mixed hypergraphs for linear-time construction of denser hashing-based data structures," in *Proc. of the Int. Conf. on Current Trends in Theory and Practice of Comp. Science.* Springer, 2013, pp. 356–368.
- [13] F. Lázaro and B. Matuz, "Irregular invertible Bloom look-up tables," in *Proc. of the 11th Int. Symp. on Topics in Coding.* IEEE, 2021, pp. 1–5.
- [14] K. Narayanan and H. Pfister, "Iterative collision resolution for slotted ALOHA: An optimal uncoordinated transmission policy," in *Proc. of 7th Int. Symp. on Turbo Codes and Iterative Inf. Processing (ISTC).* Gothenburg, Sweden: IEEE, 2012, pp. 136–139.

- [15] M. Skjegstad and T. Maseng, "Low complexity set reconciliation using Bloom filters," in *Proc. Int. Work. on Foundations of mobile Comp. (FOMC)*, no. 2027. ACM, 2011, pp. 33–41.
- [16] D. Guo and M. Li, "Set reconciliation via counting Bloom filters," *IEEE Trans. Knowledge Data Eng.*, vol. 25, no. 10, pp. 2367–2380, 2013.
- [17] L. Luo, D. Guo, O. Rottenstreich, R. Ma, and X. Luo, "Set Reconciliation with Cuckoo Filters," in *Proc. of the Int. Conf. on Inf. and Knowledge Manag. (CIKM)*. Beijing, China: ACM, Nov. 2019, pp. 2465–2468.
- [18] N. Kruber, "Approximate distributed set reconciliation with defined accuracy," PhD Thesis, Humboldt-Universität zu Berlin, 2019.
- [19] J. Byers, M. Luby, and M. Mitzenmacher, "A digital fountain approach to reliable distribution of bulk data," *IEEE J. Select. Areas Commun.*, vol. 20, no. 8, pp. 1528–1540, Oct. 2002.
- [20] T. Richardson and R. Urbanke, "Multi-edge type LDPC codes," in *Workshop honoring Prof. Bob McEliece on his 60th birthday*. Pasadena, California: California Institute of Technology, May 2002.
- [21] M. Mitzenmacher and G. Varghese, "Biff (bloom filter) codes: Fast error correction for large data sets," in *Proc of 2012 IEEE Int. Symp. on Inf. Theory.* IEEE, 2012, pp. 483–487.
- [22] M. Luby, M. Mitzenmacher, and A. Shokrollahi, "Analysis of random processes via and-or tree evaluation," in *Proc. of the 9-th annual ACM-SIAM Symp. on Discrete Algs.* San Francisco, CAL, USA: ACM, 1998, pp. 364–373.
- [23] T. Richardson, A. Shokrollahi, and R. Urbanke, "Design of capacity-approaching irregular low-density parity-check codes," *IEEE Trans. on Inf. Theory*, vol. 47, no. 2, pp. 619–637, 2001.
- [24] G. Liva, "Graph-based analysis and optimization of contention resolution diversity slotted ALOHA," *IEEE Trans. on Commun.*, vol. 59, no. 2, pp. 477–487, 2011.
- [25] R. Gabrys and A. Coker, "Set reconciliation in two rounds of communication," in *Proc. 18th Int. Command & Control Research & Tech. Symp. (ICCRTS)*. Alexandria, Virginia, USA: IC2I, Jun. 2014.
- [26] J. Ha, J. Kim, and S. McLaughlin, "Rate-compatible puncturing of low-density parity-check codes," *IEEE Trans. Inf. Theory*, vol. 50, no. 11, pp. 2824–2836, 2004.
- [27] A. Shokrollahi, "Theory and applications of Raptor codes," *Mathknow*, pp. 59–89, 2009.
- [28] S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi, "Optimization by simulated annealing," *science*, vol. 220, no. 4598, pp. 671–680, 1983.
- [29] H. Yan, U. Irmak, and T. Suel, "Algorithms for low-latency remote file synchronization," in *Proc of 27th Conf. on Comp. Commun. (INFOCOM)*. IEEE, 2008, pp. 156–160.
- [30] M. Gentili, "Set reconciliation and file synchronization using invertible Bloom lookup tables," Bachelor's Thesis, Harvard College, 2015.

#### APPENDIX

For the rate-compatible protocol, cell nodes are transmitted one by one. The recovery operation can be attempted anytime even if not all cells of a type have been transmitted and hence received. We extend the analysis of Section III to this case. The following analysis is also of interest in case a non-ideal communication channel between transmitter and receiver is assumed which introduces losses of cells (erasures).

Let us denote by  $\epsilon_i$  the probability that a cell node of type  $i$  is not received (erased). For  $m_i \rightarrow \infty$  (by the law of large numbers)  $\epsilon_i$  is the fraction of erased cell nodes of type  $i$ . For the analysis, we assume that these cell nodes are part of the IBLTs bipartite graph, but the outgoing messages (i.e., probability of erasure) on all connected edges are one throughout all iterations. The remaining fraction of  $(1 - \epsilon_i)$  cell nodes of type  $i$  will send out a message given by (1). Overall, the average erasure probability from a cell node of type  $i$  to a data node in case of erasures is modified as follows,

$$\bar{\omega}_i^{(t)} = \epsilon_i + (1 - \epsilon_i) \left( 1 - e^{-\frac{\eta}{f_i} \bar{d}_i \bar{q}_i^{(t)}} \right). \quad (5)$$

Observe that by not transmitting cell nodes of a certain type, the load is increased, i.e.,

$$\eta = \frac{n}{\sum_{i=1}^{d_c} (1 - \epsilon_i) m_i}. \quad (6)$$

Apart from the modification in (5) and the adjustment of the load in (6) the load threshold computation can be carried out as described in Section III.