Batch PIR and Labeled PSI with Oblivious Ciphertext Compression

Alexander Bienstock* Sarvar Patel† Joon Young Seo‡ Kevin Yeo§

Abstract

In this paper, we study two problems: oblivious compression and decompression of ciphertexts. In oblivious compression, a server holds a set of ciphertexts with a subset of encryptions of zeroes whose positions are only known to the client. The goal is for the server to effectively compress the ciphertexts obliviously, while preserving the non-zero plaintexts and without learning the plaintext values. For oblivious decompression, the client, instead, succinctly encodes a sequence of plaintexts such that the server may decode encryptions of all plaintexts value, but the zeroes may be replaced with arbitrary values. We present solutions to both problems that construct lossless compressions only 5% more than the optimal minimum using only additive homomorphism. The crux of both algorithms involve embedding ciphertexts as random linear systems that are efficiently solvable.

Using our compression schemes, we obtain state-of-the-art schemes for batch private information retrieval (PIR) where a client wishes to privately retrieve multiple entries from a server-held database in one query. We show that our compression schemes may be used to reduce communication by up to 30% for batch PIR in both the single- and two-server settings.

Additionally, we study labeled private set intersection (PSI) in the unbalanced setting where one party's set is significantly smaller than the other party's set and each entry has associated data. By utilizing our novel compression algorithm, we present a protocol with 65-88% reduction in communication with comparable computation compared to prior works.

1 Introduction

Protecting user privacy is becoming a core problem in today's society with the continuing growth of cloud-based applications. There are many important cloud services that provide databases of essential information that need to be retrieved by users. In many cases, it is necessary to hide the queried database entry to preserve the user's privacy. This privacy requirement has appeared in many cloud services provided by large organizations including certificate transparency [3], contact discovery [10], device enrollment [7], password leak check [6, 5] and URL blocklists [8].

Private information retrieval (PIR) [27] is an important cryptographic protocol that enables an user to retrieve entries from a public database without revealing the identity of queried entries. PIR has been studied in both the single- and multi-server settings. The main difference is that multi-server PIR requires stronger assumptions of non-colluding servers. In our work, we will study both types of PIR protocols.

*New York University, abienstock@cs.nyu.edu. Part of work done while interning at Google.

†Google, sarvar@google.com.

‡Google, jyseo@google.com.

§Google and Columbia University, kwlyeo@google.com.

	Encoding Size	Encoding Time	Decoding Time
Choi *et al.* [26]	$O(t\lambda)$	$O(n\lambda)$	$O(t\lambda)$
Liu and Tromer [55]	$O(t \log^2 t \log \lambda)$	$O(nt)$	$O(t^3)$
Fleischhacker *et al.* [35]	$O(t)$	$O(n \log n)$	$O(t\sqrt{n})$
Fleischhacker *et al.* [35]	$O(t\lambda)$	$O(n\lambda)$	$O(t\lambda)$
Ours: LSObvCompress	$(1+\epsilon)t$	$O(n\lambda)$	$O(t\lambda)$

Figure 1: Comparison of ciphertext compression for n ciphertexts with t non-zero values for failure probability at most $2^{-\lambda}$. Encoding size is measured in number of ciphertexts.

For many use cases, it is required that users retrieve a batch of multiple entries from the same public database. Some examples include anonymously retrieving encrypted recent messages from a communication system [15], privately fetching relevant advertisements according to user interests [44, 61] and checking validity of multiple certificates [56, 48].

To solve this problem, prior works have studied the notion of batch PIR where the user retrieves a set of t entries in a single query. The naive approach of executing t single-query PIR has the high computational overhead of $O(tn)$ as state-of-the-art single-query PIR schemes still require $O(n)$ server computation linear in the database size. Instead, current batch PIR solutions drastically decrease computation at the cost of increased communication. Angel *et al.* [14] presented a solution that reduce computation to $3n$ that required performing $1.5t$ independent PIR queries on smaller databases. Unfortunately, the number of requests and responses are 50% larger than the naive approach of t single-query PIR executions.

To reduce communication, prior works packed multiple single-query PIR requests into a single ciphertext [14, 13] as well as encoding multiple PIR responses for small database entries into a single ciphertext using vectorization techniques [59]. However, this still requires explicitly encoding $0.5t$ "dummy" requests and responses. Furthermore, response techniques only apply for small entries where ciphertexts can pack multiple entries. In this work, we present compression techniques to avoid encoding non-essential values that are applicable regardless of the database entry size.

While PIR considers the setting of public database, the same problem also occurs for retrieving mutliple entries from private databases with sensitive data where users should not receive information irrelevant to them. This problem has been studied as labeled private set intersection (PSI) or batch symmetric PIR. We will use labeled PSI throughout the rest of our work. Examples use cases with a private database include discovering any contacts using a service [51] and checking all credentials of a user against a database of leaked credentials [13]. In our work, we will also study ways to reduce communication for labeled PSI using ciphertext compression.

1.1 Our Contributions

We identify two compression problems and present efficient schemes for both problems relying only on additive homomorphism of the underlying encryption scheme. This leads to improved batch PIR and labeled PSI constructions (when the encryption scheme used for these constructions is indeed somewhat homomorphic, as required in prior state-of-the-art work).

Oblivious Ciphertext Compression. We study the problem of *oblivious ciphertext compression* where a compressor is given n ciphertexts of which $t < n$ are non-zero. The compressor is unaware of the identity of the t non-zero ciphertexts. The decompressor has the private key for decryption

	Client Storage	Request Overhead	Response Overhead
Baseline	$O(1)$	1x	1x
Cuckoo Hashing [14]	$O(n)$	$\lceil 1.5/r \rceil$x	1.5x
Vectorized [59]	$O(n)$	$\lceil 1.5/r \rceil$x	$\lceil 1.5/d \rceil$x
Keyword [64]	$O(1)$	$\lceil 1.5/r \rceil$x	1.5x
Distributed Point Function (DPF)* [20]	$O(1)$	1.5x	1.5x
Ours: Single-Server	$O(1)$	$\lceil (1+\epsilon)/r \rceil$x	1.5x
Ours: Single-Server	$O(1)$	$\lceil 1.5/r \rceil$x	$\lceil (1+\epsilon)/d \rceil$x
Ours: Two-Server*	$O(1)$	1.5x	$(1+\epsilon)$x

Figure 2: Keyword batch PIR comparisons for retrieving ℓ entries from n-entry database. Request and response overhead is compared to baseline of performing ℓ independent single-query PIR executions. We use r and d to denote the number of requests and plaintext database entries that can fit into a single ciphertext. Asterisks(*) denote two-server PIR protocols.

and knows the location of the t non-zero ciphertexts. The goal is to enable the compressor to construct a succinct encoding that may be correctly decoded by the decompressor.

We present LSObvCompress with encodings of $(1+\epsilon)t$ ciphertexts (where ϵ is a configurable parameter) while requiring only homomorphic addition of ciphertexts. In practice, we achieve ϵ to be as small as 0.05. Note, this is only 5% larger than the optimal compression rate that would consist of only t ciphertexts. Furthermore, oblivious compression requires only $O(n\lambda)$ homomorphic additions and decompression requires only $O(t\lambda)$ plaintext additions such that decompression is successful except with probability $2^{-\lambda}$. Our protocol, LSObvCompress, utilizes novel techniques to compress ciphertexts by encoding them as random linear systems that are efficiently solvable.

LSObvCompress significantly outperforms any prior compression schemes applicable to our setting. In particular, all prior schemes with efficient encoding and decoding produce encodings with $O(t\lambda)$ ciphertexts that is significantly larger than LSObvCompress in practice. Only one previous solution produced encodings with $O(t)$ ciphertexts [35], but the decoding is prohibitively expensive requiring computation of $O(t\sqrt{n})$ discrete logarithms. See Figure 1 for more comparisons. To be fair, we note that prior works study a more challenging version of this problem (see Section 2.1).

Oblivious Ciphertext Decompression. *Oblivious ciphertext decompression* switches the roles of compressor and decompressor. The compressor is given n plaintexts $\mathbf{p} = [p_1, \ldots, p_n]^T$, a subset $I \subset [n]$ of $t < n$ indices and a private encryption key. The goal is to produce a succinct encrypted compression of $\mathbf{p}$. The decompressor must be able to correctly retrieve the ciphertext vector $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]^T$ such that $\tilde{c}_i$ must be an encryption of p_i for all $i \in I$. There are no requirements for any $i \notin I$. The decompressor must decompress obliviously without knowledge of the indices, $I \subset [n]$.

We present LSObvDecompress with nearly identical efficiency as LSObvCompress, with encodings of size $(1+\epsilon)t$ ciphertexts. In practice, we get ϵ to be as small as 0.05 that is only 5% larger than optimal.

For decoding failure probability at most $2^{-\lambda}$, compression requires $O(t\lambda)$ plaintext additions while decompression requires $O(n\lambda)$ homomorphic additions. To our knowledge, no prior works are applicable to this specific problem.

Batch PIR. We apply our compression techniques to obtain state-of-the-art batch PIR schemes

with reduced communication in both the single- and two-server settings. Our techniques work for both small and large entry databases. See Figure 2 for detailed comparisons with prior works.

The cuckoo hashing framework of Angel *et al.* [14] transforms any single-query PIR into a batch PIR protocol. To retrieve ℓ entries, the batch PIR performs 1.5ℓ single-query PIR executions. Recent work by Mughees and Ren [59] used vectorization techniques to pack multiple small entries into a single ciphertext. If d database entries fit into a ciphertext, vectorized batch PIR returns $\lceil 1.5\ell/d \rceil$ ciphertexts. However, this only works for small entries where $d \geq 2$.

In our work, we first present a batch PIR that can reduce response size regardless of database entry sizes. In the cuckoo hashing framework, at most ℓ PIR responses will be encryptions of relevant values while the remaining 0.5ℓ will be encryptions of zero. Using $\mathsf{LSObvCompress}$, we reduce the response size from 1.5ℓ to 1.05ℓ PIR responses. To our knowledge, this is the first response reduction for batch PIR with large database entries. We also show that our techniques are compatible with the vectorization techniques of Mughees and Ren [59]. If d entries fit into a ciphertext, our techniques reduce the response size from $\lceil 1.5\ell/d \rceil$ ciphertexts to $\lceil 1.05\ell/d \rceil$ ciphertexts.

Similar ideas may also be used to reduce the request communication as well. Again, the client only cares about ℓ PIR requests and the remaining 0.5ℓ may be ignored. We leverage $\mathsf{LSObvDecompress}$ to reduce the total request size by only compressing values of the ℓ important requests and, essentially, ignoring the other 0.5ℓ requests. Combined with packing techniques[14] where r single-server PIR requests may fit into a ciphertext, we reduce request sizes from $\lceil 1.5/r \rceil$ to $\lceil 1.05/r \rceil$ ciphertexts. We stil apply apply vectorization [59] to obtained $\lceil 1.5/d \rceil$ response ciphertexts.

Finally, we show similar response reduction may also be obtained in two-server batch PIR protocols by applying $\mathsf{LSObvCompress}$ to prior constructions [20].

Labeled PSI. Next, we show that $\mathsf{LSObvCompress}$ and $\mathsf{LSObvDecompress}$ may also be used to construct improved schemes for labeled PSI. In particular, one can combine our above batch PIR construction, leveraging $\mathsf{LSObvCompress}$ and $\mathsf{LSObvDecompress}$, with any oblivious PRF (OPRF) to obtain a labeled PSI protocol. Our labeled PSI schemes provides a 65-88% reduction in communication with comparable computation over prior solutions [23, 28].

2 Preliminaries

Linear Algebra. We denote $\mathbf{v}$ as column vectors and $\mathbf{v}^T$ as row vectors. We denote the i-th entry of $\mathbf{v}$ by $\mathbf{v}_i$. For two vectors n-length vectors $\mathbf{v}$ and $\mathbf{u}$, we denote the dot product operator as $\mathbf{v} \cdot \mathbf{u} = \sum_{i=1}^{n} \mathbf{v_i} \cdot \mathbf{u_i}$. We define a $n \times m$ matrix using its column vectors as $\mathbf{M} = [\mathbf{v}_1, \ldots, \mathbf{v}_m]$ where the i-th column vector is $\mathbf{v}_i$ of length n. We may also define a matrix using its row vectors as $\mathbf{M} = [\mathbf{v}_1^T, \ldots, \mathbf{v}_n^T]$ where $\mathbf{v}_i^T$ is the i-th row vector of length m. We denote the matrix-vector product $\mathbf{M} \cdot \mathbf{u} = [\mathbf{v}_1 \cdot \mathbf{u}, \ldots, \mathbf{v}_n \cdot \mathbf{u}]$ where $\mathbf{u}$ is a m-length vector. We solve the linear system associated with $n \times m$ matrix $\mathbf{M}$ and n-length vector $\mathbf{u}$ by computing m-length vector $\mathbf{v}$ such that $\mathbf{M} \cdot \mathbf{v} = \mathbf{u}$.

For a vector $\mathbf{v}$ of length n and subset $I = \{i_1, \ldots, i_k\} \subseteq [n]$, we denote by $\mathbf{v}_I = [\mathbf{v}_{i_1}, \ldots, \mathbf{v}_{i_k}]$ containing the entries of $\mathbf{v}$ with indices in I. For $n \times m$ matrix $\mathbf{M} = [\mathbf{v}_1^T, \ldots, \mathbf{v}_n^T]$ and subset $I = \{i_1, \ldots, i_k\} \subseteq [n]$, we denote the sub-matrix consisting of row vectors with indices in I as $\mathbf{M}_{\mathsf{r}(I)} = [\mathbf{v}_{i_1}^T, \ldots, \mathbf{v}_{i_k}^T]$. Similarly, for a $n \times m$ matrix $\mathbf{M} = [\mathbf{v}_1, \ldots, \mathbf{v}_m]$ and subset $I = \{i_1, \ldots, i_k\} \subseteq [m]$, we denote the sub-matrix consisting of column vectors with indices in I as $\mathbf{M}_{\mathsf{c}(I)} = [\mathbf{v}_{i_1}, \ldots, \mathbf{v}_{i_k}]$.

Homomorphic Encryption. Throughout our work, we will define ciphertexts using $\tilde{c}$. A vector of ciphertexts will be defined as $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]$.

In our work, we will mainly consider lattice-based somewhat homomorphic encryption (SHE) where parameters are chosen to support a limited number of homomorphic operations, as used in prior state-of-the-art constructions of batch PIR and labeled PSI [60, 58, 23, 24, 28]. Our compression protocols only use additive hommorphism of these schemes, where noise grows additively. We refer to Appendix A for more details on SHE and recent PIR schemes using SHE [60, 58].

2.1 Oblivious Ciphertext Compression

We define the notion of an *oblivious ciphertext compression* scheme. For this primitive, we only assume additive homomorphism (ciphertext-ciphertext addition). The problem consists of two parties: a compressor and a decompressor. The compressor is given n ciphertexts, $\tilde{\mathbf{c}} = [\tilde{c}_1, ..., \tilde{c}_n]$, to be compressed. Both the compressor and the decompressor know the number of non-zero plaintext entries t. In addition, the decompressor has the private decryption key and the indices of the t non-zero entries, $I \subset [n]$. If $i \in I$, then $\tilde{c}_i$ is an encryption of a non-zero entry. The compressor's job is to produce a succinct encoding of the input ciphertexts with knowledge of only t. The encoding is consumed by the decompressor to recover the original t non-zero plaintext entries. We formally define oblivious ciphertext compression below.

Definition 1 (Oblivious Ciphertext Compression)**.** *Let $\mathbf{p} = [p_1, ..., p_n] \in \mathbb{F}^n$ be a vector of n plaintexts with at most t non-zero entries. Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Eval}, \mathsf{Dec})$ be an additive homomorphic encryption scheme, and let $\tilde{\mathbf{c}} = [\tilde{c}_1, ..., \tilde{c}_n]$ where $\tilde{c}_i = \mathcal{E}.\mathsf{Enc}(\mathbf{pk}_\mathcal{E}, p_i)$ for each $i \in [n]$. An oblivious ciphertext compression scheme consists of a pair of algorithms ($\mathsf{ObvCompress}, \mathsf{Decompress}$) satisfying:*

- $\hat{\mathbf{c}} \leftarrow \mathsf{ObvCompress}(\mathbf{pk}_\mathcal{E}, \tilde{\mathbf{c}}, t; R)$*: Oblivious compression takes in a public key $\mathbf{pk}_\mathcal{E}$, n ciphertexts $\tilde{\mathbf{c}} = [\tilde{c}_1, ..., \tilde{c}_n]$, the number of non-zero plaintext entries t, and randomness R. It outputs compressed ciphertexts $\hat{\mathbf{c}}$.*
- $\mathbf{p} \leftarrow \mathsf{Decompress}(\mathbf{sk}_\mathcal{E}, \hat{\mathbf{c}}, I; R)$*: Decompression takes in a secret key $\mathbf{sk}_\mathcal{E}$, compressed ciphertexts $\hat{\mathbf{c}}$, the non-zero plaintext entry indices $I \subset [n]$ ($|I| \leq t$) of p, and randomness R. It outputs the non-zero plaintext values $\{i, p_i\}_{i \in I}$.*

Let $\gamma = \gamma(\lambda)$ be the bit length of all n ciphertexts produced by the homomorphic encryption scheme $\mathcal{E}$. An oblivious ciphertext compression is δ-compressing if the bit length of $\hat{\mathbf{c}}$ is at most $\delta \cdot \gamma \cdot |\tilde{\mathbf{c}}|$. The failure probability is at most ϵ if, for each plaintext vector $\mathbf{p} = [p_1, ..., p_n]$ and associated ciphertexts $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]$ with at most t non-zero values,

$$\Pr[\mathsf{Decompress}(\mathbf{sk}_\mathcal{E}, \hat{\mathbf{c}}, I) \neq \{i, p_i\}_{i \in I}] \leq \epsilon$$

where $\hat{\mathbf{c}} \leftarrow \mathsf{ObvCompress}(\mathbf{pk}_\mathcal{E}, \tilde{\mathbf{c}}, t)$.

Comparison with Prior Work. Liu and Tromer [55] implicitly studied oblivious ciphertext compression, without explicitly defining the primitive. Fleischhacker *et al.* [35] considered another variant closer to our compression problem that was also implicitly studied in [55]. where the decompressor is not given the identity of the non-zero plaintext indices, $I \subset [n]$. Therefore, this is a harder setting than our compression problem. It is not surprising that the resulting compression rates or decoding efficiency are significantly worse than our constructions (see Figure 1). To our knowledge, our specific variant of compression has not been explicitly studied previously.

2.2 Oblivious Ciphertext Decompression

Next, we define *oblivious ciphertext decompression* that switches the compressor and decompressor roles. The compressor is given the plaintext vector, $\mathbf{p} = [p_1, \ldots, p_n]$ and a subset of t indices, $I \subset [n]$ with $|I| = t$ to produce a succinct encoding $\hat{\mathbf{c}}$. The decompressor is given $\hat{\mathbf{c}}$ and must produce the ciphertext vector $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]^T$ such that each $\tilde{c}_i$ is an encryption of p_i for all $i \in I$. No correctness is required for $i \notin I$. In other words, $\tilde{c}_i$ needs to be an encryption of p_i only when $i \in I$. However, the decompressor must obliviously decode without any knowledge of the relevant indices, I. In fact, the compressed ciphertexts $\hat{\mathbf{c}}$ must not reveal any information about neither the underlying plaintext values $\mathbf{p} = [p_1, \ldots, p_n]^T$ nor the relevant indices I. To our knowledge, no prior works have studied this setting.

Definition 2 (Oblivious Ciphertext Decompression)**.** *Let $p = [p_1, ..., p_n]^T \in \mathbb{F}^n$ be a vector of n plaintexts and $I \subset [n]$ be a subset of $t < n$ indices. Let $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Eval}, \mathsf{Dec})$ be an additive homomorphic encryption scheme. A oblivious ciphertext decompression scheme consists of a pair of algorithms ($\mathsf{Compress}, \mathsf{ObvDecompress}$), where:*

- $\hat{\mathbf{c}} \leftarrow \mathsf{Compress}(\mathbf{sk}_{\mathcal{E}}, p, I; R)$*: The compression algorithm takes in a secret homomorphic encryption key $\mathbf{sk}_{\mathcal{E}}$, a vector of n plaintexts $p = [p_1, ..., p_n]^T$, a subset of t indices $I \subset [n]$ and randomness R. Then, it outputs the compressed ciphertexts $\hat{\mathbf{c}}$.*
- $\mathbf{p} \leftarrow \mathsf{ObvDecompress}(\mathbf{pk}_{\mathcal{E}}, \hat{\mathbf{c}}, n; R)$*: The decompression algorithm takes in a public homomorphic encryption key $\mathbf{pk}_{\mathcal{E}}$, compressed ciphertexts $\hat{\mathbf{c}}$, the number of total plaintexts n, and randomness R. Then, it outputs the ciphertext vector $\tilde{c} = [\tilde{c}_1, \ldots, \tilde{c}_n]^T$.*

Let $\gamma = \gamma(\lambda)$ be the bit length of all n ciphertexts produced by the homomorphic encryption scheme $\mathcal{E}$. A oblivious ciphertext decompression is δ-compressing if the bit length of $\hat{\mathbf{c}}$ is at most $\delta \cdot \gamma \cdot |\tilde{c}|$. The failure probability is at most ϵ if, for each plaintext vector $p = [p_1, ..., p_n]^T$ and subset $I \subset [n]$ of size t, the following holds:

$$\Pr[\exists i \in I \mid \mathsf{Dec}(\mathbf{sk}_{\mathcal{E}}, \tilde{c}_i) \neq p_i] \leq \epsilon$$

where $\hat{\mathbf{c}} \leftarrow \mathsf{ObvCompress}(\mathbf{sk}_{\mathcal{E}}, p, I)$ and $[\tilde{c}_1, \ldots, \tilde{c}_n]^T \leftarrow \mathsf{Decompress}(\mathbf{pk}_{\mathcal{E}}, \hat{\mathbf{c}})$. We note that there are no correctness requirements for ciphertexts $\tilde{c}_i$ such that $i \notin I$.

The scheme is computationally oblivious if, for all pairs of plaintext vectors $p = [p_1, \ldots, p_n]^T$ and $p' = [p'_1, \ldots, p'_n]^T$ and pairs of index sets $I, I' \subset [n]$ of size t, a computationally adversary cannot distinguish between the following:

- $\hat{\mathbf{c}} \leftarrow \mathsf{Compress}(\mathbf{sk}_{\mathcal{E}}, p, I)$
- $\hat{\mathbf{c}}' \leftarrow \mathsf{Compress}(\mathbf{sk}_{\mathcal{E}}, p', I')$.

2.3 Batch PIR and Labeled PSI

Batch (Keyword) PIR. In batch keyword PIR, the client holds a batch of ℓ keys, $\{q_1, \ldots, q_\ell\}$, and the server holds a public database $D \in (\mathcal{K} \times \mathcal{V})^n$ of n key-value pairs with n distinct keys, $\{(k_1, v_1), \ldots, (k_n, v_n)\}$. The client wishes to retrieve the database entries $\{D[q_1], \ldots, D[q_\ell]\}$ from the server. For any $q \in \mathcal{K}$, $D[q]$ denotes the value associated with key q. If $q = k_i$, then $D[q] = v_i$. Otherwise, $D[q] = \perp$. The following two properties must hold:

- *Correctness*: If the protocol is executed correctly, the client recovers $\{D[q_1], \ldots, D[q_\ell]\}$ as desired.
- *Query Privacy*: The server learns no information about the batch query, $\{q_1, \ldots, q_\ell\}$.

One can obtain the definition of single-query PIR if the batch query contains only a single index, $\ell = 1$. Furthermore, one can obtain non-keyword PIR if we restrict the database's key universe to be $\mathcal{K} = [n]$. Throughout our work, we will consider keyword PIR unless otherwise specified.

(Unbalanced) Labeled PSI. In labeled PSI, the receiver and sender hold sets X and Y respectively. The sender also holds a database of associated labels $\{L_y \mid y \in Y\}$. The goal is for the receiver to receive labels that appear in the intersection, $\{(z, L_z) \mid z \in X \cap Y\}$. The following properties must hold:

- *Correctness*: If the protocol is executed correctly, the receiver recovers $\{(z, L_z) \mid z \in X \cap Y\}$ as desired.
- *Receiver (Query) Privacy*: The sender learns no information about the receivers's set X beyond its size $|X|$.
- *Sender (Database) Privacy*: The receiver learns no information about the sender's set Y except for the desired output and its size $|Y|$.

In the unbalanced setting, the receiver's set X is typically much smaller than the sender's set Y, $|X| \ll |Y|$. Note, labeled PSI is similar to batch keyword PIR with the main difference being the additional sender (database) privacy guarantee.

3 Oblivious Ciphertext Compression

In this section, we present our oblivious ciphertext compression scheme, $\mathsf{LSObvCompress}$, based on linear systems. We start with a simpler scheme before presenting our main construction.

3.1 First Attempt: Balls-into-Bins

In this section, we start with a construction which leverages the balls-into-bins random process. Given m bins and n balls, each of the n balls are thrown into one of the m bins uniformly at random. In the context of ciphertext compression, bins correspond to compressed ciphertexts and balls correspond to input non-zero ciphertexts. Throwing a ball into a bin corresponds to homomorphically adding an input ciphertext to one of the compressed ciphertexts. Decompression works by re-simuluating the ball throws for non-zero ciphertexts and decrypting the values at relevant bins. The main observation is that adding a zero-encrypting ciphertext can be thought of as "skipping" the ball throw, as its addition doesn't change the value of the underlying plaintext. Conceptually, the algorithm fails if any of the bins contains more than one ball. We describe the algorithm below.

We suppose that both parties share a hash function H. Upon receiving the input ciphertexts $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]^T$ and the number of non-zero plaintext entries t, the compression algorithm first initializes a vector of $m \geq t$ zero ciphertexts $\hat{\mathbf{c}} = [\hat{c}_1, ..., \hat{c}_m]^T$, where $\hat{c}_i = \mathcal{E}.\mathsf{Enc}(\mathbf{pk}_{\mathcal{E}}, 0)$. Then, for each input ciphertext $\tilde{c}_i$, the algorithm executes the following two operations. First, compute index

$j = H(i) \in [m]$ where H is a random function with range $[m]$. Next, homomorphically add $\tilde{c}_i$ to $\hat{c}_j$, that is, $\hat{c}_j = \mathcal{E}.\mathsf{Eval}(\mathbf{pk}_{\mathcal{E}}, +, [\tilde{c}_i, \hat{c}_j])$. Finally, the algorithm outputs the resulting vector $\hat{\mathbf{c}}$.

The decompression algorithm receives the compression $\hat{\mathbf{c}} = [\hat{c}_1, \ldots, \hat{c}_m]^T$ and non-zero plaintext entry indices I. For every non-zero ciphertext index $i \in I$, the algorithm computes $j = H(i)$ and sets $p_i = \mathcal{E}.\mathsf{Dec}(\mathbf{sk}_{\mathcal{E}}, \hat{c}_j)$. Finally, the algorithm outputs all non-zero plaintext values, $\{i, p_i\}_{i \in I}$.

Note this algorithm can recover the original plaintext vector as long as the hash outputs $H(i)$ are all distinct for every $i \in I$. However, the probability of collision is high unless $m = \Omega(t^2)$ (due to the birthday problem) that is a quadratic blowup with respect to t. Ideally, we would like m to be not much larger than t to obtain an efficient compression rate.

Reformulating as a Linear System. We generalize the aforementioned scheme as constructing and solving a system of linear equations. More specifically, the compression algorithm is responsible for constructing a linear system that the decompression algorithm attempts to solve to recover the original plaintext vector. While this viewpoint seems rather unnecessarily complex, it will serve as an important basis to our main construction. We outline the reformulated algorithm below.

For each $i \in [n]$, the compression algorithm constructs a column vector $\mathbf{v}_i \in \mathbb{F}^m$ where only the $H(i)$-th element is set to 1 and the rest are set to 0. Let $\mathbf{M} = [\mathbf{v}_1, ..., \mathbf{v}_n] \in \mathbb{F}^{m \times n}$ be a matrix. Note that both parties know matrix $\mathbf{M}$ as they share hash function H. The compression algorithm computes and outputs the matrix-vector multiplication $\hat{\mathbf{c}} = \mathbf{M} \cdot \tilde{\mathbf{c}}$.

The decompression algorithm takes in the vector $\hat{\mathbf{c}}$ and produces its decryption $\hat{\mathbf{p}}$. Next, we reconstruct the matrix $\mathbf{M}$ using the random function H. Let $I = \{i_1, ..., i_t\}$ be the set of non-zero plaintext entry indices, and let $\mathbf{M}_{\mathsf{c}(I)} = [\mathbf{v}_{i_1}, ..., \mathbf{v}_{i_t}] \in \mathbb{F}^{m \times t}$ be a sub-matrix of $\mathbf{M}$ consisting of all column vectors whose indices appear in I. Similarly, let $\hat{\mathbf{p}}_I = [\hat{p}_{i_1}, \ldots, \hat{p}_{i_t}]$ for entries of $\hat{\mathbf{p}}$ in I. The algorithm solves the linear system associated with $\mathbf{M}_{\mathsf{c}(I)}$ and $\hat{\mathbf{p}}$ to compute $\mathbf{p}_I$ satisfying $\mathbf{M}_{\mathsf{c}(I)} \cdot \mathbf{p}_I = \hat{\mathbf{p}}_I$ to recover the non-zero $p_{i_j} = (\mathbf{p}_I)_j$ for each $j \in [t]$.

We note that the decompression algorithm can correctly recover the plaintext vector if and only if the linear system $\mathbf{M}_{\mathsf{c}(I)} \cdot \mathbf{p}_I = \hat{\mathbf{p}}_I$ has a unique solution (that is, $\mathbf{M}_{\mathsf{c}(I)}$ has full column rank). For our choice of $\mathbf{M}$, this precisely happens when all hash outputs $H(i)$ are distinct for every $i \in I$.

3.2 Second Attempt: Random Matrices

Recall that in the first attempt, the generated matrix $\mathbf{M}$ consists of random column vectors with Hamming weight exactly one corresponding to the balls-into-bins process. This forced us to set the number of rows and the encoding size to $m = \Omega(t^2)$ to avoid collisions. Taking a closer look, we notice that the way we generate the column vectors are unnecessarily restrictive. Indeed, for our scheme to succeed, we only require the $\mathbf{M}_{\mathsf{c}(I)}$ to have a unique solution. There is no need to restrict rows to Hamming weight one vectors.

This crucial observation leads to the following approach. Instead of sampling random column vectors with Hamming weight 1, we instead sample column vectors uniformly at random from $\{0,1\}^m$. To do this, we can imagine the shared hash function $H : [n] \to \{0,1\}^m$ outputs random binary column vectors of length m. Then, the shared matrix is $\mathbf{M} = [H(1), \ldots, H(n)]$. This way, the generated column vectors will be linearly independent with high probability even when m is small. The rest of the algorithm stays identical.

Failure Probability and Compression Rate. The algorithm's failure probability and compression rate will be parameterized by ϵ and t. Let $m = (1 + \epsilon)t$ be the number of rows. Even when ϵ is very small, the generated $m \times t$ matrix $\mathbf{M}_{\mathsf{c}(I)}$ has a unique solution except with negligible

probability. For example, setting $m = t + \lambda$ with very small $\epsilon = \lambda/t$, the system has full rank with probability $1 - 2^{-\lambda-1}$ (see [37]). The compression rate is almost optimal as the encoding contains $t + \lambda$ ciphertexts that is only λ more than the optimal minimum.

Running Time. Let $m = (1 + \epsilon)t$. We start by analyzing the compression time. Generating a random column vector $\in \{0,1\}^m$ takes $O(m)$ time, so the entire matrix generation takes $O(mn)$ time during compression. Computing the matrix-vector product takes $m \cdot n$ homomorphic ciphertext additions. The compression algorithm performs $O(m \cdot t)$ ciphertext-ciphertext additions. For decompression, we note that solving the linear system associated to $\mathbf{M}_{\mathsf{c}(I)}$ requires $O(m \cdot t^2)$ time using Gaussian elimination.

Comparison to the First Attempt. While the new algorithm can give us very high compression rate, it is computationally very inefficient. Compression requires $O(mt)$ time and decompression requires $O(mt^2)$ time using Gaussian elimination. In practice, this may not be so problematic when $t << n$, but as t grows, the scheme is computationally expensive. Ideally, we would like compression to be close to linear in the number of ciphertexts, n, and decompression to be close to linear in t. In contrast, the first attempt has horrible compression rate of $m = O(t^2)$, but is computationally more efficient. The compression algorithm requires only $O(n)$ time. Furthermore, decompression only used $O(t^2)$ time.

This raises the following question: is it possible to get the best of both worlds - an algorithm that achieves high compression rate but is also practically efficient? We show that this is possible in the next subsection.

3.3 LSObvCompress: Random Band Matrices

In prior attempts, we generated random matrices uniformly at random from $\{0,1\}^{m \times n}$. This allowed the associated random linear systems to be uniquely solvable with high probability even when $m = (1 + \epsilon)t$ was very small. However, solving this linear system is very inefficient, which made the previous scheme impractical for larger t. This is not too surprising, because the generated matrix is very dense. The expected number of non-zero matrix entries is $mn/2$. This suggests that the algorithm for solving the linear system must also have at least $O(mn)$ running time as well.

Looking closely, we again realize that we never needed the generated matrices to be sampled uniformly at random from $\{0,1\}^{m \times n}$. That is, as long as the associated linear system is uniquely solvable with high probability, the distribution itself is irrelevant to the security of the scheme. Therefore, we only require a matrix generation algorithm that generates a "small" linear system that is uniquely and efficiently solvable. For LSObvCompress, we consider random matrices that satisfy these two properties.

Random Band Matrices. There has been extensive research on the core algorithmic problem of generating sparse random matrices that are efficiently solvable. For LSObvCompress, we utilize the random band matrices of Dietzfelbinger and Walzer [31] that is the most efficient to our knowledge.

Random band matrices are constructed such that each row consists of a random band with width w, and all entries outside of the band are zero. Formally, let m be the length of each row of the matrix. For each row, a band start index s is chosen randomly from $[m - w + 1]$, and each entry within the band, i.e. in range $[s, s + w)$, is a uniformly random bit from $\{0,1\}$. All other entries outside the range $[s, s + w)$ remain 0.

Intuitively, random band matrices are solvable in $O(nw)$ time because the generated random matrix is "almost diagonal" after the rows are sorted by the band start positions. Furthermore,

each row reduction operation maintains an invariant where the number of non-zero entries per rows is $O(w)$ making Gaussian elimination very efficient.

Adaptation for LSObvCompress. Unfortunately, we are unable to directly apply random band matrices for LSObvCompress. Going back to the linear system framework presented in Section 3.1, the client will solve the linear system associated with the matrix $\mathbf{M}_{\mathsf{c}(I)}$. Recall that I is the subset of non-zero plaintexts, $\mathbf{M}$ is the chosen random matrix and $\mathbf{M}_{\mathsf{c}(I)}$ is the sub-matrix of $\mathbf{M}$ consisting of all the column vectors whose indices appear in I. Suppose we chose $\mathbf{M}$ to be a random band matrix. Unfortunately, $\mathbf{M}_{\mathsf{c}(I)}$ is not guaranteed to be a random band matrix. In particular, it is possible that I (and, thus, the columns) are chosen such that each matrix row will have a band much smaller than length w or be all zero. In this case, it is unclear if the matrix $\mathbf{M}_{\mathsf{c}(I)}$ still has a unique solution.

Instead, we will choose our matrix $\mathbf{M}$ using an adaptation of random band matrices to ensure that $\mathbf{M}_{\mathsf{c}(I)}$ is still efficiently solvable for any choice of non-zero plaintext indices I. To do this, we will instead choose $\mathbf{M}$ to be the tranpose of random band matrices. In other words, we will generate each column vector of $\mathbf{M}$ to consist of a random band of width w. To do this, we imagine both parties share two hash functions $\mathsf{H}_1 : [n] \to [m - w + 1]$ and $\mathsf{H}_2 : [n] \to \{0,1\}^w$. For the i-th column of the shared matrix $\mathbf{M}$, $\mathsf{H}_1(i)$ denotes the start of the band and $\mathsf{H}_2(i)$ chooses the random w-bit band.

Next, we can consider any subset of non-zero plaintexts I and the associated sub-matrix $\mathbf{M}_{\mathsf{c}(I)}$. As each column vector consists of a random w-length band, $\mathbf{M}_{\mathsf{c}(I)}$ remains a transpose of a random band matrix. As the column and row rank of any matrix is identical, we can rely on the analysis of Dietzfelbinger and Walzer [31] to see that $\mathbf{M}_{\mathsf{c}(I)}$ will have a unique solution with high probability for any choice of I.

The only caveat is that we cannot apply the running time analysis of the random band row matrix construction, as the bands are constructed column-wise instead of row-wise. Nonetheless, we show that solving the system remains practically efficient with this modification in our experiments (see Section 7.1). Intuitively, this is because a transpose of a random band row matrix remains similar to a random band row matrix after the columns are sorted by the band start positions. The maximum band width across the entire rows is not much larger than the column band width w, which allows the linear system to be solved efficiently just as in the random band row matrix construction. See Figure 3 for an illustration. We prove the following theorem (see Appendix D for the proof):

Theorem 1. *Consider a $m \times t$ matrix with $m = (1+\epsilon)t$ where each column consists of a single random w-bit band. For constant $\epsilon > 0$ and band length $w = O(\lambda + \log t)$, the random band matrix has column rank n and executing Gaussian elimination after sorting the columns by the starting location of the band runs in time $O(tw)$ except with probability $2^{-\lambda}$.*

We now formally present LSObvCompress using random band matrices. See Algorithms 1 and 3 for the description of the oblivious compression and decompression algorithms.

Next, we analyze the properties of LSObvCompress showing that it combines the good compression rates and efficient encoding/decoding times of our prior two attempts.

Failure Probability and Compression Rate. For the failure probability, we note that LSObvCompress fails only when $\mathbf{M}_{\mathsf{c}(I)}$ does not have a unique solution or that unique solution cannot be found. By Theorem 1, we know this occurs with probability at most $2^{-\lambda}$ assuming that $w = O(\lambda/\epsilon + \log n)$.

Algorithm 1 LSObvCompress.ObvCompress algorithm

Input: $\mathbf{pk}_{\mathcal{E}}, \tilde{\mathbf{c}}, t, R$: public additively homomorphic encryption key, vector of n ciphertexts, number of non-zero plaintext entries, and randomness.

Output: $\hat{\mathbf{c}}$: compressed encoding of $\tilde{\mathbf{c}}$.

$m \leftarrow (1+\epsilon)t$
$\mathbf{M} \leftarrow 0^{m \times n}$
for $i = 1, \ldots, n$ **do**
  $\mathbf{v_i} \leftarrow \mathsf{GenRandVec}(i, m; R)$
  $\mathbf{M}[:][i] \leftarrow \mathbf{v_i}$ ▷ Set the ith column to $\mathbf{v_i}$
$\hat{\mathbf{c}} \leftarrow \mathbf{M} \cdot \tilde{\mathbf{c}}$ ▷ HE add using $\mathcal{E}.\mathsf{Eval}$ and $\mathbf{pk}_{\mathcal{E}}$
return $\hat{\mathbf{c}}$

Algorithm 2 GenRandVec algorithm

Input: i, m, R: column index, column vector length, and randomness

Output: $\mathbf{v_i}$: generated random column vector

$w \leftarrow$ band width
$s \leftarrow \mathsf{H_1}(R \,||\, i)$ ▷ Random value from $[m - w + 1]$
$\mathbf{u} \leftarrow \mathsf{H_1}(R \,||\, i)$ ▷ Random w-bit band.
$\mathbf{v_i} \leftarrow 0^m$
for $j = 0, \ldots, w-1$ **do**
  $\mathbf{v_i}[s+j] \leftarrow \mathbf{u}[j]$
return $\mathbf{v_i}$

In our experiments, we will use concrete parameters for w and ϵ for various values of t to obtain 2^{-40} error probability. We point readers to Section 7.1 for more details. For the compression rate, our experiments show that ϵ may be as small as 0.05. As a result, LSObvCompress obtains compression rates that are only 5% larger than optimal.

Running Time. We start by analyzing the compression algorithm that computes the matrix multiplication of $\mathbf{M}$ and the input ciphertext vector $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]^T$. As $\mathbf{M}$ is a binary matrix with at most nw non-zero entries, this can be performed using at most nw ciphertext-ciphertext additions. For decompression, we note that the main cost is solving the linear system $\mathbf{M}_{\mathsf{c}(I)}$ that requires $O(tw)$ time by Theorem 1 that is corroborated by our experiments (see Section 7.1).

Noise Growth for SHE. Recall that for the applications to batch PIR and labeled PSI, we will initialize LSObvCompress, using lattice-based SHE schemes. Therefore, noise growth is an important factor to consider. Suppose that the input ciphertexts $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]^T$ each have error at most $\mathsf{Err}(\tilde{c}_i) \leq e$. We note that the compression algorithm requires computing the sum of at most w ciphertexts. Therefore, each ciphertext in the compressed output has error at most $O(w \cdot e)$ as ciphertext-ciphertext additions only incur linear noise growth (see Appendix A for more details). As decompression is done after decryption, we do not need to worry about noise growth for decompression.

Algorithm 3 LSObvCompress.Decompress algorithm

Input: $\mathbf{sk}_{\mathcal{E}}, \hat{\mathbf{c}}, I, R$: secret additively homomorphic encryption key, compressed encoding of ciphertexts, set of non-zero plaintext indices, and randomness.

Output: $\{i, p_i\}_{i \in I}$: original non-zero plaintext values.

$m \leftarrow (1+\epsilon)t$

$\mathbf{M}_{\mathsf{c}(I)} \leftarrow 0^{m \times t}$ ▷ Initialize all zero matrix.

for $i_j \in I = \{i_1, \ldots, i_t\}$ **do**

$\mathbf{v}_{i_j} \leftarrow \mathsf{GenRandVec}(i_j, m; R)$

$\mathbf{M}_{\mathsf{c}(I)}[:][j] \leftarrow \mathbf{v}_{i_j}$ ▷ Set the jth column to $\mathbf{v}_{i_j}$

$\hat{\mathbf{p}} \leftarrow$ decryption of $\hat{\mathbf{c}}$ using $\mathcal{E}.\mathsf{Dec}$ and $\mathbf{sk}_{\mathcal{E}}$

$p_I \leftarrow \mathsf{SolveLinearSystem}(\mathbf{M}_{\mathsf{c}(I)}, \hat{\mathbf{p}})$

if $p_I = \perp$ **then**

return $\perp$

$\mathbf{p} \leftarrow \emptyset$

for $i_j \in I = \{i_1, \ldots, i_t\}$ **do**

$\mathbf{p} \leftarrow \mathbf{p} \cup \{(i_j, (p_I)_j)\}$

return p

Algorithm 4 SolveLinearSystem algorithm

Input: $\mathbf{M}, \hat{\mathbf{p}}$: LHS matrix, RHS values to solve for

Output: $\mathbf{p}$: solution to the linear system $\mathbf{M} \cdot \mathbf{p} = \hat{\mathbf{p}}$

$(\mathbf{M}^{\pi}, \pi) \leftarrow$ column sorting of the matrix $\mathbf{M}$ in ascending band start positions, along with the corresponding permutation that produces the column sorted matrix (e.g. $\mathbf{M}^{\pi}[:][i] = \mathbf{M}[:][\pi(i)]$)

$\mathbf{p}^{\pi} \leftarrow$ execute Gaussian elimination on $\mathbf{M}^{\pi}$ and $\hat{\mathbf{p}}$, $\perp$ if no unique solution

if $\mathbf{p}^{\pi} = \perp$ **then**

return $\perp$

$\mathbf{p} \leftarrow 0^t$

for $i = 1 \ldots t$ **do**

$\mathbf{p}[\pi(i)] \leftarrow p^{\pi}[i]$

return p

3.4 Comparison with Sparse Random Linear Codes [52, 55]

Liu and Tromer [55] implicitly study oblivious ciphertext compression. They observe that Sparse Random Linear Codes (SRLCs) [52], which use matrices $\mathbf{M} \in \mathbb{F}^{m \times n}$ where each column has a small number of non-zero entries drawn randomly from $\mathbb{F}$ can be used. However, each entry is sampled independently, in an unstructured way, which results in larger encodings than with LSObvCompress. Indeed, they show that such matrices can be sampled with full rank with high probability only if $m = O(t \log^2 t \log \lambda)$, which is larger than $m = 1.05t$ of LSObvCompress. Moreover, because SRLCs are unstructured, Gaussian elimination takes $O(t^3)$ time, resulting in slower $O(t^3)$ decoding time, compared to the $O(t \cdot \lambda)$ decoding time of LSObvCompress. Finally, since SRLCs use elements drawn randomly from $\mathbb{F}$, when used with FHE, large parameters must be used to handle the noise when multiplying ciphertexts by these large elements. However, LSObvCompress only uses elements from $\{0, 1\}$, which means that ciphertexts are only added together, resulting in minimal noise growth.

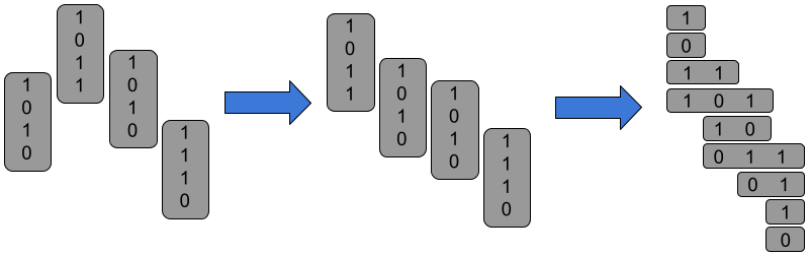

Figure 3: Example of a random band column matrix construction with band width $w = 4$. Second diagram shows the matrix after sorting the columns by the band start positions. Third diagram shows the random band row matrix view of the constructed matrix. In this example, the maximum band row width is 3.

4 Oblivious Ciphertext Decompression

We show that similar ideas that we used to solve the oblivious ciphertext compression problem may also be used to solve the oblivious ciphertext decompression problem. As a reminder, in this problem, the compressor is given a plaintext vector, $\mathbf{p} = [p_1, \ldots, p_n]^T$, and t relevant indices $I \subset [n]$. The goal is for the decompressor to decode ciphertexts $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]^T$ such that $\tilde{c}_i$ is an encryption of p_i for all relevant indices $i \in I$. There are no requirements for any $i \notin I$.

Description of LSObvDecompress**.** Essentially, we will apply the ideas of LSObvCompress, but in reverse. That is, we will start with a matrix $\mathbf{M}$ of dimension $n \times m$ (that both the compressor and decompressor can generate based on shared randomness), where $m = (1 + \epsilon)t$ is the encoding length for some constant $\epsilon > 0$. Then, based on relevant row indices $I = \{i_1, \ldots, i_t\} \subset [n]$, the compressor will solve the linear system formed by a $(t \times m)$-dimensional sub-matrix $\mathbf{M}_{\mathsf{r}(I)}$ of $\mathbf{M}$ and vector $\mathbf{p}_I = [p_{i_1}, p_{i_2}, \ldots, p_{i_t}]$, to obtain compressed plaintext vector $\hat{\mathbf{p}}$ of dimension m. Specifically, the compressor will solve the linear system for $\hat{\mathbf{p}}$ satisfying $\mathbf{M}_{\mathsf{r}(I)} \cdot \hat{\mathbf{p}} = \mathbf{p}_I$ using sub-matrix $\mathbf{M}_{\mathsf{r}(I)} = [\mathbf{M}_{i_1}^T, \ldots, \mathbf{M}_{i_t}^T]$.

Afterwards, the vector $\hat{\mathbf{p}}$ is encrypted entry-wise. The encrypted version of $\hat{\mathbf{p}}$ is the final encoding that we denote $\hat{\mathbf{c}}$. If the linear system according to $\mathbf{M}_{\mathsf{r}(I)}$ is not solvable, then the encoding fails and the compressor outputs any m encryptions. In applications, this is the point where we can utilize packing techniques where multiple plaintext values may be encrypted into a single ciphertext (as done in [14]).

For oblivious decompression, the decompressor computes $\mathbf{M} \cdot \hat{\mathbf{c}}$ homomorphically. Intuitively, this gives the decompressor ciphertext vector $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_n]^T$ such that for each $i_j \in I$, the underlying plaintext of $\tilde{c}_{i_j}$ is $\mathbf{v}_{i_j} \cdot \hat{\mathbf{p}}$, which is exactly p_{i_j}, as desired. For every $\tilde{c}_{i_j}$ where $i_j \notin I$, the underlying plaintext will be some arbitrary linear combination of the entries of $\hat{\mathbf{p}}$, but recall that these values need not be correct.

For the choice of matrix $\mathbf{M}$, we can in fact generate it similarly as in LSObvCompress as a random band matrix of dimension $n \times m$, where each row consists of a single random band of length w. Note, this is the original random band matrix construction [31] without modification.

We present the pseudocode for LSObvDecompress in Algorithms 5 and 6.

Failure Probability. From above, we saw that the encoding is correct as long as the compressor can solve the linear system associated with $\mathbf{M}_{\mathsf{r}(I)}$. For the failure probability, we can simply calculate the probability that $\mathbf{M}_{\mathsf{r}(I)}$ does have a unique solution (or it cannot be found). As $\mathbf{M}$ is a random band matrix, we know that $\mathbf{M}_{\mathsf{r}(I)}$ is also a random band matrix. Therefore, if we set the band length $w = O(\lambda/\epsilon + \log t)$ and $\mathbf{M}_{\mathsf{r}(I)}$ to be a $t \times (1 + \epsilon)t$, then $\mathbf{M}_{\mathsf{r}(I)}$ has a unique solution.

Algorithm 5 LSObvDecompress.Compress algorithm

Input: $\mathbf{sk}_{\mathcal{E}}, \mathbf{p} = [p_1, \ldots, p_n]^T, R$: secret additively homomorphic encryption key, plaintext values and randomness.

Output: $\hat{\mathbf{c}}$: compressed ciphertexts.

Compute $I = \{i \mid p_i \neq 0\} \subseteq [n]$.
If $|I| > t$, abort.
If $|I| < t$, arbitrarily add indices to I until $|I| = t$.
$m \leftarrow (1+\epsilon)t$
$\mathbf{M}_{\mathsf{r}(I)} \leftarrow 0^{t \times m}$ ▷ Initialize all zero matrix.
for $i \in I$ **do**
    $\mathbf{M}_i \leftarrow \mathsf{GenRandVec}(i, m, R)^T$
$\hat{\mathbf{p}} \leftarrow \mathsf{SolveLinearSystem}(\mathbf{M}_{\mathsf{r}(I)}, \mathbf{p}_I)$
$\hat{\mathbf{c}} \leftarrow$ encryption of $\hat{\mathbf{p}}$ using $\mathcal{E}.\mathsf{Dec}$ and $\mathbf{sk}_{\mathcal{E}}$
return $\hat{\mathbf{c}}$

Algorithm 6 LSObvDecompress.ObvDecompress algorithm

Input: $\hat{\mathbf{c}}, R$: compressed ciphertexts and randomness.

Output: $\tilde{\mathbf{c}}$: decompressed ciphertexts.

$m \leftarrow (1+\epsilon)t$
for $i \in [n]$ **do**
    $\tilde{c}_i \leftarrow \mathsf{GenRandVec}(i, m, R) \cdot \hat{\mathbf{c}}$
$\tilde{\mathbf{c}} \leftarrow [\tilde{c}_1, \ldots, \tilde{c}_n]$
return $\tilde{\mathbf{c}}$

Compression Rate. We show that ϵ may be as small as 0.05 using experimental evaluation (see Section 7.1). Note, this is only 5% larger than the minimum of t ciphertexts since t plaintext values must be correctly encoded.

Running Time. The compression algorithm requires solving the linear system associated to $\mathbf{M}_{\mathsf{r}(I)}$ that is a $t \times (1+\epsilon)t$ random band matrix. This can be done in $O(tw)$ time using only plaintext operations. Additionally, the resulting vector must be encrypted using $O(m) = O(t)$ time.

Decompression simply requires computing the matrix-vector multiplication $\mathbf{M} \cdot \hat{\mathbf{c}}$. As each row has at most w one entries, this requires $O(nw)$ homomorphic additions.

Obliviousness. Note that $\hat{\mathbf{c}}$ is always a length-m ciphertext vector. Reducing to the security of the underlying encryption scheme $\mathcal{E}$, we can replace each of these ciphertexts with encryptions of 0 meaning $\hat{\mathbf{c}}$ is independent of input plaintexts.

Noise Growth for SHE. Recall that for the applications to batch PIR and labeled PSI, we will initialize LSObvCompress, using lattice-based SHE schemes. Therefore, noise growth is an important factor to consider. We note that the compression algorithm is performed in plaintext without any homomorphic operations. Therefore, we only consider noise growth for decompression. The compressed input consists of m fresh SHE ciphertexts. Decompression adds at most w ciphertexts. If the input ciphertexts $\tilde{\mathbf{c}} = [\tilde{c}_1, \ldots, \tilde{c}_m]^T$ have error $\mathsf{Err}(\tilde{c}_i) \leq e$ for all $i \in [m]$, then each output ciphertext has error at most $O(w \cdot e)$. See Appendix A for further details.

5 Batch PIR

We will present three single-server PIR schemes using our compression techniques. We refer readers to Section 7.2 for our experimental evaluation to choose the best option for various settings of database size, entry size and batch size. We also present an improved two-server scheme.

5.1 Single-Server: Compressed Responses

In this section, we present our improved single-server batch PIR with compressed responses that apply for large entries that cannot leverage vectorization techniques [59].

Cuckoo Hashing Batch PIR Framework. In this section, we review the Cuckoo Hashing Batch PIR Framework by Angel *et al.* [14]. In a naive batch PIR scheme, the server would process each of the ℓ queries on the entire n database entries, resulting in a total of $O(n\ell)$ server operations. To reduce server computation, Angel *et al.* [14] presented a batch PIR framework that cleverly utilizes cuckoo hashing to encode both the batch query and the database entries. To date, this is the most practically efficient approach to constructing a batch PIR scheme. Our batch PIR will be built directly from this framework.

In this framework, the server setup works by creating $B \geq \ell$ independent single-query PIR servers and replicating each of the n database entries appropriately to a subset of $\alpha \geq 1$ servers. Consider a sparse database $D = \{(k_1, v_1), \ldots, (k_n, v_n)\} \in (\mathcal{K} \times \mathcal{V})^n$. Concretely, the choice of the α-subset is determined by the individual database entry (k_i, v_i) and α independent hash functions $H_1, \ldots, H_\alpha : \mathcal{K} \to [B]$ mapping keys to one of the B servers. In particular, (k_i, v_i) will be replicated to the servers indexed by $H_1(k_i), \ldots, H_\alpha(k_i)$. The total number of entries across all B servers will be $n\alpha$.

The hash functions that will be shared between the client and the server so that the client may also perform batch queries. Given a batch query $\{q_1, \ldots, q_\ell\}$, the client performs cuckoo hashing to map the ℓ query keys into the B buckets. In particular, each bucket will contain at most one query key after cuckoo hashing. Then, the client constructs a single-query PIR request for each of the B PIR servers. For empty buckets, the client will construct dummy "zero" requests such that the response to a dummy request will be a ciphertext that encrypts zero. Concretely, $B - \ell$ dummy requests. Finally, the server will process the B independent single-query PIR requests and send B responses back to the client.

Concrete Instantiation. Angel *et al.* [14] empirically determined that setting $B = 1.5\ell$ and $\alpha = 3$ results in an appropriate balance between the failure probability of the client allocation procedure, and efficiency. We notice that the request and response size in the cuckoo hashing framework is larger than the naive approach. In the naive approach, the request and response size are merely ℓ ciphertexts whereas the framework requires 1.5ℓ ciphertexts. This is 50% larger than the number of responses in the naive approach.

Client Mapping or Keyword PIR. One subtlety of this framework is that each of the B independent single-query PIR servers consists of a sparse database. We note that this is true regardless of whether the original batch PIR problem consists of a dense database where $\mathcal{K} = [n]$ or a sparse database where $\mathcal{K}$ could be much larger. In earlier works (such as [14, 59]), it was suggested to use $O(n)$ client mappings to convert from database indices to bucket indices. Recent work [64] instead directly uses state-of-the-art keyword PIR schemes to avoid linear client storage.

Throughout the rest of our work, we will follow this approach and use single-query keyword PIR protocols for each of the B buckets.

Keyword PIR Framework [64]. We provide a brief overview of the keyword PIR framework of [64] (used for each of the B buckets), paying special attention to those details relevant to our final batch PIR scheme.

- As in recent standard PIR schemes [14, 13, 60, 58], the framework represents the n-entry database as a $d_1 \times d_2 \times \cdots \times d_z$ hypercube, where $d_1 \cdots d_z \geq n$.
- The query algorithm for some key k creates z vectors of length $d_1, \ldots, d_z$, homomorphically encrypts them, then uploads them to the server.
- The server takes in an encoding $\mathbf{E}$ of the database, a two-dimensional matrix of size $d_1 \times \lceil n/d_1 \rceil$, and homomorphic encryptions of vectors $\mathbf{v}_1, \ldots, \mathbf{v}_z$. It first applies $\mathbf{v}_1$ to $\mathbf{E}$ to obtain a $\lceil n/d_1 \rceil$ vector, arranges this vector into a $d_2 \times \lceil n/(d_1 d_2) \rceil$ matrix and applies $\mathbf{v}_2$ to obtain a vector of size $\lceil n/(d_1 d_2) \rceil$, and repeats this for all z dimensions (where for the last dimension, the vector from the previous step will not be arranged into a matrix and instead, an inner product with $\mathbf{v}_z$ will be performed). At the end of this process, the server obtains a ciphertext encrypting the queried entry of the client.
- The server then sends this ciphertext to the client, who can decrypt it to obtain their queried entry.

Our Construction. To reduce communication in batch PIR, we will apply $\mathsf{LSObvCompress}$ to reduce the server response communication in the cuckoo hashing framework.

Namely, recall that for the $B - \ell$ buckets which do not have an associated key, the client will construct dummy "zero" requests such that the corresponding response ciphertext will encrypt zero. This can be done by setting, e.g., $\mathbf{v}_z$ to the zero-vector, since in the last step of the response algorithm, the server computes the inner product of $\mathbf{v}_z$ with some vector to obtain the final response ciphertext. Therefore, after the server processes the $B = 1.5\ell$ requests, it obtains $B = 1.5\ell$ responses of which ℓ consist of encrypted entries, and the rest are encrypted zeros. Thus, the server can apply the compression of $\mathsf{LSObvCompress}$ with $n = B$ and $t = \ell$ to obtain compressed ciphertexts. Of course, the client knows the indices of the ℓ real requests. As a result, the client can execute the decompression of $\mathsf{LSObvCompress}$ to obtain the requested entries.

Our construction therefore results in response size with overhead as small as $1.05\times$ the optimal, with minimal added computation, instead of the $1.5\times$ overhead in response size of [14].

Noise Growth. For noise growth, we will assume that the keyword PIR framework [64] is applied using recent PIR schemes from SHE composition [60, 58]. We perform the same noise analysis for prior PIR schemes in Appendix B and see that our new PIR scheme increases the noise growth by a $O(w)$ multiplicative factor. See Appendix C for analysis.

5.2 Single-Server: Compressed Requests

Next, we apply $\mathsf{LSObvDecompress}$ to compress requests for single-server batch PIR schemes.

Our Construction. In our framework using the keyword PIR from [64], the client generates $B = 1.5\ell$ requests, each containing z vectors. However, we only need correct answers from ℓ

requests. Thus, the client can combine all $B \cdot z$ request vectors into one long vector, and for relevant indices I consisting only of entries corresponding to the z vectors for each of the ℓ important requests, apply LSObvDecompress. This will result in a compressed request with size overhead only $1.05\times$ compared to the naive batch PIR, which the client can then encrypt and send to the server. This is in contrast to the $1.5\times$ overhead in request size of [14]. We also utilize in our construction the request packing techniques from [14] to fit multiple requests into a single ciphertext.

The server will first apply the request ciphertext packing decoding and, then run decompression from LSObvDecompress to obtain the B encrypted requests. Note, only the ℓ important requests will be correct. This is sufficient as the remaining 0.5ℓ dummy requests are ignored by the client anyways. The remainder of the server processing and client decrypting remains identical.

Noise Growth. We show that applying LSObvDecompress increases noise by an $O(w)$ multiplicative factor (see Appendix C for further details).

Why not both request and response compression? Theoretically, one can apply compression for both requests and responses simultaneously. In practice, the noise growth is $O(w^2)$ multiplicative factor (see Appendix C for the analysis) that is large. We were unable to find parameters where request and response compression beat either of the PIR schemes above. We leave it as an open problem to find better SHE/PIR schemes enabling both request and response compression. The full details of using both may be found in Appendix E.

5.3 Single-Server: Vectorized Responses

We present a method to compress responses in conjunction with the recent vectorization techniques of Mughees and Ren [59]. The vectorization techniques [59] utilize *Single-Instruction-Multiple-Data* (SIMD) techniques. SIMD encodes multiple database entries into a single ciphertexts (leveraging additional structure of the SHE scheme) and operates on all of them simultaneously.

Our Construction. The core idea of utilizing LSObvCompress to compress responses remains the same, but we wish to leverage that multiple entries fit into a single ciphertext. To do this, we present a vectorized version of LSObvCompress that optimally packs multiple entries into a single ciphertext. If d entries fit into a single ciphertext, our vectorized LSObvCompress sends only $\lceil 1.05\ell/d \rceil$ to the client. In contrast, $\lceil 1.5\ell/d \rceil$ ciphertexts are encoded in [59].

At a high level, vectorized LSObvCompress works nearly identically as the variant of LSObvCompress described in Section 3.3. The only difference is that we apply techniques from [59] to rotate ciphertexts and pack multiple entries into a ciphertext before performing compression. Due to lack of space, we defer the full description to Appendix F.

5.4 Two-Server: Compressed Responses

Next, we use LSObvCompress to compress responses for two-server batch PIR. In this setting, the client sends requests to both servers. Each server holds a copy of the database and cannot communicate with each other. They then send individual responses back to the client, who uses both to reconstruct the requested entry. The same correctness and privacy conditions are required. Privacy is considered with respect to each individual server assuming non-collusion.

Two-Server PIR. To date, the most concretely efficient two-server PIR schemes are built using distributed point functions (DPF) [43, 20, 46]. A point function $f_i : \mathcal{K} \to \{0, 1\}$ satisfies $f_i(x) = 1$

if and only if $x = i$. DPFs enable secret sharing f_i amongst the two servers using two functions, f_i^0 and f_i^1, satisfying $f_i(x) = f_i^0(x) + f_i^1(x)$ for all $x \in \mathcal{K}$. Both f_i^0 and f_i^1 must not individually reveal anything about f_i.

To perform a two-server keyword PIR query for key k, the client uses DPFs to create secret shares of f_k, f_k^0 and f_k^1, that are sent to each of the two servers. Suppose the database consists of n key-value pairs, $D = \{(k_1, v_1), \ldots, (k_n, v_n)\}$. For each server $j \in \{0, 1\}$, the j-th server computes

$$z_j = f_k^j(k_1) \cdot v_1 + \ldots + f_k^j(k_n) \cdot v_n.$$

Finally, the client receives z_0 and z_1 and computes the final answer $z_0 + z_1$. If $k = k_i$, then $z_0 + z_1 = v_i$. Otherwise, $z_0 + z_1 = 0$ when $k \notin \{k_1, \ldots, k_n\}$.

There is a small issue that the client cannot distinguish between $v_i = 0$ and $k \notin \{k_1, \ldots, k_n\}$. To fix this, we can ensure that zero is not a valid entry. For example, one can simply append a 1-bit to the end of each entry, $v_1, \ldots, v_n$.

Two-Server Batch PIR. To our knowledge, the most efficient two-server batch PIR remains the cuckoo hashing framework of Angel *et al* [14]. In the concrete instantiation, we use a two-server, single-query, keyword PIR for each of the $B = 1.5\ell$ buckets when performing a batch query for ℓ entries.

Our Construction. Our goal is to utilize LSObvCompress to reduce the number of responses from $B = 1.5\ell$ that are sent by both servers. We will use the two-server keyword PIR based on the DPF of [20]. Note that this PIR does not use encryption and, instead, relies on the non-collusion of the two servers for security. The encryption scheme $\mathcal{E}$ for LSObvCompress in this setting is additive secret sharing. Homomorphic additions will simply be addition operations by each server.

First, recall that in order to use LSObvCompress, the 0.5ℓ dummy requests must results in additive sharings (encryptions) of 0. We will add $(k_0, 0)$ is added to each of the B buckets for special key k_0. The client will issue a keyword PIR query for k_0 for each of the 0.5ℓ dummy buckets. The servers upon receipt of these requests for all B buckets will then compute the corresponding responses. Consider the i-th response z_i^0 and z_i^1 for both servers. Let $I' \subset [B]$ be the indices of the real, non-dummy requests. For all $i \in [B]$, $z_i = z_i^0 + z_i^1$ is the i-th requested entry. If $i \notin I'$, then $z_i = z_i^0 + z_i^1 = 0$.

We can thus apply LSObvCompress as follows. Both servers will use LSObvCompress to compress their responses $\mathbf{z}^0 = [z_1^0, \ldots, z_B^0]$ and $\mathbf{z}^1 = [z_1^1, \ldots, z_B^1]$. Recall that this is done by computing the matrix-vector multiplications $\hat{\mathbf{z}}^0 = \mathbf{M} \cdot \mathbf{z}^0$ and $\hat{\mathbf{z}}^1 = \mathbf{M} \cdot \mathbf{z}^1$ where $\mathbf{M}$ is the transpose of a random band matrix. The client will compute $\hat{\mathbf{z}} = \hat{\mathbf{z}}^0 + \hat{\mathbf{z}}^1 = \mathbf{M} \cdot (\mathbf{z}^0 + \mathbf{z}^1)$ that is a compression of the requested entries, $\mathbf{z}^0 + \mathbf{z}^1$. Finally, the client runs the decompression portion of LSObvCompress on $\hat{\mathbf{z}}$ to obtain the non-dummy queried entries, z_i for all $i \in I$.

6 Labeled PSI

In this section, we show our techniques may be used to build protocols for labeled PSI in the unbalanced setting. Recall that the receiver has a set X and the sender has a labeled set $\{(y, L_y) \mid y \in Y\}$. Note, that one may interpret this as batch keyword PIR with the receiver as the client and the sender as the server. The only difference is that PSI requires privacy for both parties' inputs. Therefore, we need to also enable privacy for the sender (server) input. We present two improved constructions using our compression techniques.

Batch Keyword PIR and Oblivious PRF. We can use the generic transformation from [36] combining our single-server batch keyword PIR with any oblivious pseudorandom function (OPRF) to build our labeled PSI protocol. An OPRF allows the receiver to input set X and learn the set of pseudo-random outputs $\{F_k(x) \mid x \in X\}$, where F is a PRF, and k is known to the sender. For security, both the sender and receiver should learn nothing else (beyond the size of $|X|$). Our scheme provides full security against a malicious receiver and privacy against a malicious sender, as in prior works [23, 28].

At a high level, the protocol goes as follows. First, the sender generates a private key k and evaluates the OPRF on its input set Y. The labels $\{L_y \mid y \in Y\}$ are encrypted using keys derived from the OPRF evaluation. The sender and receiver execute the OPRF protocol on the receiver's input and the sender's private key. Finally, the sender and receiver execute a batch PIR using the receiver's output of the OPRF protocol to retrieve the encrypted labels. Afterwards, the receiver decrypts labels in the intersection. We point readers to Appendix G.1 for more details and analysis.

We note that there are subtleties in the security argument due to the usage of the keyword PIR construction [64]. In particular, the database encoding algorithm from that keyword PIR construction has non-negligible failure probability. One option is to ensure that the database encoding algorithm fails negligibly by increasing the band length parameter (using the analysis from [19]). However, we show that this is unnecessary – since the input to the database encoding algorithm is just key-value pairs where each key is a random hash of the corresponding item in Y output by the OPRF and each value is a pseudo-random encryption of the label under a random key output by the OPRF, this failure probability is not a function of the sender's set Y and thus reveals nothing about it. Therefore, it suffices to use the keyword PIR construction unmodified with non-negligible encoding failures. See Appendix G.1 for the formal security proof.

Improving Oblivious Polynomial Evaluation. Prior works [24, 23, 28] built labeled PSI using oblivious polynomial evaluation (OPE). We can apply LSObvDecompress to OPE protocols for reducing request sizes by 30%. Although, we note prior work [59] showed that batch PIR approaches result in more efficient labeled PSI protocols compared to OPE. Nevertheless, OPE is an interesting application of our compression algorithms. See Appendix G.2 for more details.

7 Experimental Evaluation

We perform experimental evaluation for our new compression algorithms, LSObvDecompress and LSObvCompress, as well as their improvements to batch PIR and labeled PSI. Finally, we also benchmark our protocols for the real world application of anonymous messaging.

Experimental Setup. We implemented our experimental evaluations with around 3000 lines of C++ code. All our experiments are performed using Ubuntu PCs with 96 cores, 3.7 GHz Intel Xeon W-2135 and 128 GB of RAM with only single-threaded execution. The AVX2 and AVX-512 instruction sets with SIMD instructions are enabled. The results are the average of at least 10 experimental trials with standard deviation less than 10% of the averages. Our implementations will target error probability 2^{-40} and 128 bits of computational security. Server monetary costs are computed using Amazon EC2 savings plan pricing of t2.2xlarge instances [4] of \$0.09 per GB of traffic and \$0.021 per CPU hour at the time. We will utilize SHA256 as the hash function and AES-GCM-256 as the encryption scheme with 32 byte keys. Unless otherwise specified, we will use the compression parameter $\epsilon = 0.05$ for our experiments.

Interpreting the Experimental Results. As our compression schemes are general schemes that can be instantiated on various protocols, they incur additional computational overhead compared to the ones that don't use our compression schemes. To assess concrete tradeoffs between the computational overhead and the communication reduction, we will use the Amazon EC2 server monetary cost model which measures the communication and computational efficiency as a dollar cost. We note that this model has been used in prior works [64, 19] for this exact purpose.

7.1 Oblivious Ciphertext Compression

We first evaluate the performance of LSObvCompress and LSObvDecompress in isolation and report results in Figure 4.

Setup. In our experiments, we will use Regev encryption [65] as the underlying scheme using the implementation from Spiral [9]. We fix the plaintext size to 8 KB and the ciphertext size to 20 KB. In the figure, t corresponds to the number of non-zero entries and n corresponds to the total number of entries including the zero entries. We fix the fraction of zero entries to $0.5t$ (thus $n = 1.5t$). Note that this corresponds to the fraction of dummy requests/responses in the cuckoo hashing framework from [14]. In our evaluations, we will target two compression sizes of $1.05t$ and $1.07t$. Note that this results in 30% and 29% request/response size reduction respectively.

Results. We see that computation time increases with better compression rate as well as larger t and n. However, we claim that our LSObvCompress and LSObvDecompress remain practically efficient for many applications; as we will show in the next sections, the additional computational cost is a relatively small fraction of the entire protocol's computation time, and the significant reduction in the communication cost will justify these small additional computational overhead.

7.2 Single-Server Batch PIR

We evaluate the single-server batch PIR schemes from Section 5 using our compression techniques to reduce communication. We report our results in Figure 5.

Setup. We implement our compression algorithms on top of the open-source Spiral implementation [9]. We use $n = 1$ million database entries for all of our results. Baseline corresponds to Angel *et al* [14]'s batch PIR framework implemented on top of Spiral [58] without our compression techniques. The parameters for the baseline were chosen using the script provided by their open-source implementation. In our evaluations, we consider three batch sizes $\ell \in \{512, 1024, 2048\}$ (in the context of ciphertext compression/decompression, the batch size ℓ corresponds to the number of non-zero entries). We follow the batch PIR setup from [14] and fix the fraction of dummy requests/responses (i.e. zero entries) to 0.5ℓ. We target compression size of 1.05ℓ.

Results. Using our LSObvCompress, we see 30% response size reduction in exchange for a reasonable additional computational cost compared to state-of-the-art PIR for large entries without compression. We see that this small additional computation cost is justified by the reduction in the server monetary cost. In particular, LSObvCompress reduces the server monetary cost by up to 10% compared to the baseline.

Using our LSObvDecompress algorithm, we see 20-24% reduction in the request size with slight increase in computation and server monetary cost.

We were unable to integrate our vectorized version of LSObvCompress into the vectorized batch PIR protocol [59] as we are unaware of an open-source implementation. Nevertheless, we still

Compression Size	Sizes & Schemes	Compression Time	Decompression Time	Total Time
$1.05t$	$\mathbf{t = 512, n = 768}$			
	LSObvCompress	2.38 s	0.66 s	3.04 s
	LSObvDecompress	0.60 s	1.65 s	2.25 s
	$\mathbf{t = 1024, n = 1536}$			
	LSObvCompress	5.15 s	1.34 s	6.49 s
	LSObvDecompress	1.25 s	3.94 s	5.19 s
	$\mathbf{t = 2048, n = 3072}$			
	LSObvCompress	10.54 s	2.67 s	13.21 s
	LSObvDecompress	2.48 s	6.88 s	9.38 s
	$\mathbf{t = 4096, n = 6144}$			
	LSObvCompress	21.45 s	5.27 s	26.72 s
	LSObvDecompress	5.05 s	14.50 s	19.55 s
$1.07t$	$\mathbf{t = 512, n = 768}$			
	LSObvCompress	1.82 s	0.53 s	2.35 s
	LSObvDecompress	0.49 s	1.34 s	1.83 s
	$\mathbf{t = 1024, n = 1536}$			
	LSObvCompress	4.12 s	1.07 s	5.19 s
	LSObvDecompress	0.96 s	3.19 s	4.15 s
	$\mathbf{t = 2048, n = 3072}$			
	LSObvCompress	8.41 s	2.12 s	10.53 s
	LSObvDecompress	1.72 s	5.80 s	7.52 s
	$\mathbf{t = 4096, n = 6144}$			
	LSObvCompress	17.06 s	4.43 s	21.49 s
	LSObvDecompress	2.76 s	11.19 s	13.95 s

Figure 4: Evaluations of LSObvCompress and LSObvDecompress for different values of t (non-zero/relevant entries) and n (uncompressed input size). We fix the plaintext size to 8 KB and ciphertext size to 20 KB for all our results.

DB Entry Size	Batch Size & Schemes	Public Param Size	Request Size	Response Size	Total Server Time	Amortized Server Time	Total Client Time
8 KB	$\ell = \mathbf{512}$						
	Baseline	20.87 MB	1.81 MB	15.59 MB	**840 s**	**1.64 s**	**6.1 s**
	LSObvCompress	22.62 MB	1.81 MB	**10.92 MB**	890 s	1.74 s	6.7 s
	LSObvDecompress	20.87 MB	**1.40 MB**	15.59 MB	863 s	1.69 s	6.4 s
	$\ell = \mathbf{1024}$						
	Baseline	20.87 MB	3.35 MB	31.18 MB	**1,256 s**	**1.23 s**	**6.8 s**
	LSObvCompress	23.11 MB	3.35 MB	**21.84 MB**	1,369 s	1.34 s	8.2 s
	LSObvDecompress	20.87 MB	**2.55 MB**	31.18 MB	1,323 s	1.29 s	8.0 s
	$\ell = \mathbf{2048}$						
	Baseline	20.87 MB	3.96 MB	62.37 MB	**1,750 s**	**0.85 s**	**7.0s**
	LSObvCompress	23.43 MB	3.96 MB	**43.67 MB**	1,871 s	0.91 s	9.7 s
	LSObvDecompress	20.87 MB	**3.15 MB**	62.37 MB	1,812 s	0.89 s	9.4 s
16 KB	$\ell = \mathbf{512}$						
	Baseline	20.87 MB	1.81 MB	31.18 MB	**1,286 s**	**2.51 s**	**6.3 s**
	LSObvCompress	22.62 MB	1.81 MB	**21.84 MB**	1,348 s	2.63 s	8.4 s
	LSObvDecompress	20.87 MB	**1.40 MB**	31.18 MB	1,308 s	2.55 s	8.3 s
	$\ell = \mathbf{1024}$						
	Baseline	20.87 MB	3.35 MB	62.37 MB	**1,775 s**	**1.73 s**	**7.9 s**
	LSObvCompress	23.11 MB	3.35 MB	**43.69 MB**	1,929 s	1.88 s	9.6 s
	LSObvDecompress	20.87 MB	**2.55 MB**	62.37 MB	1,881 s	1.83 s	9.4 s
	$\ell = \mathbf{2048}$						
	Baseline	20.87 MB	3.96 MB	124.74 MB	**2,634 s**	**1.29 s**	**8.5 s**
	LSObvCompress	23.43 MB	3.96 MB	**87.34 MB**	2,773 s	1.35 s	12.4 s
	LSObvDecompress	20.87 MB	**3.15 MB**	124.74 MB	2,746 s	1.34 s	12.4 s

Figure 5: Evaluations of Spiral Batch PIR [14, 58] with and without our compression techniques, LSObvCompress and LSObvDecompress with $\epsilon = 0.05$. We fix the number of entries to $n = 1$ million for all our results.

implemented our vectorized version LSObvCompress from Section 5.3 using the SEAL library [66]. The results are presented in Appendix F.1.

Choosing the Right Protocol. In general, LSObvCompress provides the best communication and server monetary cost reduction. Thus, LSObvCompress will typically be the best option for most settings.

In certain settings, we note that LSObvDecompress may be useful where we wish to minimize upload communication from the client to the server. There are many natural settings where the upload costs/speed are more expensive/slower than the download costs/speed. For applications in these scenarios, it may be critical to save as much upload communication as possible that is achieved by LSObvDecompress.

Application: Anonymous Messaging. Angel and Setty [15] introduced Pung that built an anonymous messaging protocol using any single-server batch PIR (see Appendix H). In Figure 6, we report our results for retrieving 288-byte messages. We fix the number of database entries to $n = 1$ million and batch size to $b = 512$.

By using our improved batch PIR constructions, we obtain more communication-efficient versions of Pung. In particular, we see that LSObvCompress reduces the server monetary cost by 7%

Schemes	Public Param Size	Request Size	Response Size	Total Server Time	Total Client Time	Server Monetary Cost
Baseline	20.5 MB	0.86 MB	8.90 MB	**328 s**	**1.4 s**	$0.00279
LSObvCompress	23.1 MB	0.86 MB	**6.23 MB**	338 s	1.6 s	**$0.00260**
LSObvDecompress	20.5 MB	**0.66 MB**	8.90 MB	347 s	1.7 s	$0.00288

Figure 6: Instantiation of the Pung messaging system [15] using batch Spiral PIR with and without our compression techniques ($\epsilon = 0.05$). We fix the number of database entries to $n = 1$ million and batch size to $\ell = 512$. Each entry is of size 288 B.

Batch Size & Schemes	Response Size	Server Time	Client Time	Server Monetary Cost
$\ell =$ **512**				
Baseline	221 KB	**9.63** s	**0.01 s**	$0.000076
LSObvCompress	**155 KB**	9.69 s	0.08 s	**$0.000070**
$\ell =$ **1024**				
Baseline	442 KB	**9.76 s**	**0.01 s**	$0.000096
LSObvCompress	**310 KB**	9.92 s	0.16 s	**$0.000085**
$\ell =$ **2048**				
Baseline	885 KB	**9.79 s**	**0.01 s**	$0.000136
LSObvCompress	**619 KB**	10.09 s	0.33 s	**$0.000114**
$\ell =$ **4096**				
Baseline	1,769 KB	**9.81 s**	**0.01 s**	$0.000216
LSObvCompress	**1,238 KB**	10.53 s	0.78 s	**$0.000172**
$\ell =$ **8192**				
Baseline	3,539 KB	**9.82 s**	**0.01 s**	$0.000375
LSObvCompress	**2,477 KB**	11.13 s	1.80 s	**$0.000287**

Figure 7: Comparison of DPF based two server batch PIR protocol [2] with and without LSObvCompress ($\epsilon = 0.05$). We fix the number of database entries to $n = 1$ million and each entry size to 288 B for all our results.

compared to the baseline.

7.3 Two-Server Batch PIR

We implement our response-compressed two-server batch PIR from Section 5.4 on top of the two-server single-query PIR implementation in [2]. We report our results in Figure 7.

Setup. We fix the number of database entries to $n = 1$ million where each database entry is 288 bytes large. As in the single-server batch PIR experiment (Section 7.2), we fix the fraction of dummy responses to 0.5ℓ and target compression size of 1.05ℓ. We omit evaluating request sizes as they are the same for both schemes.

Results. We observe that using LSObvCompress can reduce response size by 30% in exchange for a small additional computational cost. However, the small additional computational cost is justified by the savings in the server monetary cost. Compared to the baseline, LSObvCompress can reduce the server monetary cost by up to 24%.

Label Size & Schemes	Total Online Comm.	Total Online Time	Server Monetary Cost
512 B			
Cong *et al.* [28]	33.2 MB	**169 s**	$0.00397
LSObvCompress	**11.4 MB**	304 s	**$0.00279**
1024 B			
Cong *et al.* [28]	66.1 MB	**331 s**	$0.00787
LSObvCompress	**11.6 MB**	355 s	**$0.00311**
1536 B			
Cong *et al.* [28]	103.6 MB	535 s	$0.01244
LSObvCompress	**11.9 MB**	**446 s**	**$0.00367**

Figure 8: Comparisons of Cong *et al.* [28]'s labeled PSI and our LSObvCompress-based PSI with $\epsilon = 0.05$. We fix the size of the sender's set to 1 million and the receiver's set to 512.

7.4 Labeled PSI

Next, we evaluate our labeled PSI built from our batch PIR in Section 5.1 and an OPRF protocol (see Appendix G.1.1 for details). We report our results in Figure 8.

Setup. Our implementation uses the single-server batch PIR implementation from Section 5.1 with the OPRF implementation from [1]. We fix the size of the sender's set to 1 million and receiver's set to 512 (note, in the context of batch PIR these corresponds to the number of database elements and the batch size respectively). We have used one of the default parameter sets available in their open-source implementations for Cong *et al* [28]'s scheme.

Results. Our scheme has 65-88% reduced communication over prior state-of-the-art works [28]. For smaller label size, our construction with Spiral [58] is slower, but we start to catch up and eventually outperform Cong *et al* [28] for larger label sizes. Note that even with these additional computation cost, we reduce the server monetary cost by 30-70%.

Due to the limitation of their open source implementation [1], we could not compare our construction on larger label sizes, but we expect our scheme to outperform significantly as the label sizes increase. In any case, our communication cost and server monetary cost is significantly smaller.

8 Related Works

Ciphertext Compression. Variants of ciphertext compression have been studied in the past. Liu and Tromer [55] implicitly studied oblivious ciphertext compression without explicitly defining the primitive. In their scheme, they use sparse linear random codes that result in larger encodings and slower decoding time (see Figure 1), and, if instantiated with a FHE scheme, larger parameters for that scheme. Angel *et al.* [14] used packing and vectorization techniques to reduce request communication in PIR. Mughees and Ren [59] also showed vectorization techniques may be used to reduce response communication in batch PIR. Fleischhacker *et al.* [35] studied a more challenging variant of our setting where neither the decompressor (client) nor the compressor (server) know the identity of the the non-zero ciphertexts. As a result, their schemes have worse compression rate and more expensive compression and decompression algorithms. The same problem was implicitly studied in [26].

PIR. Single-server PIR was first studied by Kushilevitz and Ostrovsky [53]. Follow-up works constructed PIR from various other assumptions [22, 62, 29, 54, 40]. More recent works have studied concretely efficient protocols from lattice-based homomorphic encryption [11, 15, 14, 39, 63, 13, 60, 12, 58, 57, 64].

PIR has also been studied in the setting of multiple, non-colluding servers. A line of work has studied the communication efficiency with information-theoretic security (see [27, 33, 32] and references therein). Recent works have studied concretely efficient two-server PIR with computational security using distributed point functions [43, 20, 46].

Batch PIR. Batch PIR has been studied heavily in the past. Beimel *et al.* [18] presented a method to reduce server computation using matrix multiplication. Groth *et al.* [45] presented a communication-optimal scheme adapting the scheme in [40]. Another line of work (see [49, 56, 47, 69] and references therein) presented batch codes that transforms any single-query PIR into a batch PIR. More recent work [15, 14] introduced probabilistic batch codes that result in the most concretely-efficient batch PIR schemes to date. Mughees and Ren [59] introduced vectorization techniques to reduce server responses for small database entries. Patel *et al.* [64] presented keyword PIR schemes that can remove the client mapping.

Labeled PSI. Labeled PSI is a variant where each identifier has an associated data label that should be retrieved. Labeled PSI is most often studied in the unbalanced setting where the receiver's set is much smaller than the sender's set. Many recent works [24, 23, 30, 51, 28] studied labeled PSI with sub-linear communication in the larger set. The same setting where the receiver only queries for a single item has been studied as symmetric PIR [42, 13].

9 Conclusions

In this work, we present state-of-the-art constructions for both batch PIR and labeled PSI with reduced communication costs compared to prior solutions. To do this, we identify a common task in both primitives that we denote as oblivious ciphertext compression where a compressor (server) is given n ciphertexts with only $t < n$ non-zero ciphertexts. The decompressor (client) knows the location of the t non-zero ciphertexts, but the compressor is unaware of this knowledge. We present LSObvCompress that enables compressions consisting of $1.05t$ ciphertexts that is only 5% larger than the minimum while requiring only additive homomorphism. Using LSObvCompress, we present batch PIR schemes with 30% smaller responses and labeled PSI protocols with 65-88% reduced communication and comparable computation.

References

[1] APSI: C++ library for Asymmetric PSI. https://github.com/microsoft/APSI.

[2] C++ DPF-PIR library. https://github.com/dkales/dpf-cpp.

[3] Certificate transparency. https://certificate.transparency.dev/.

[4] EC2 On-Demand Pricing. https://aws.amazon.com/ec2/pricing/on-demand/.

[5] Password monitor: Safeguarding passwords in microsoft edge. https://www.microsoft.com/en-us/research/blog/password-monitor-safeguarding-passwords-in-microsoft-edge/.

[6] Protect your accounts from data breaches with password checkup. https://security.googleblog.com/2019/02/protect-your-accounts-from-data.html.

[7] Protecting your device information with private set membership. https://security.googleblog.com/2021/10/protecting-your-device-information-with.html.

[8] Safe Browsing APIs (v4). https://developers.google.com/safe-browsing/v4.

[9] Spiral. https://github.com/menonsamir/spiral.

[10] Technology preview: Private contact discovery for Signal. https://signal.org/blog/private-contact-discovery/.

[11] Carlos Aguilar Melchor, Joris Barrier, Laurent Fousse, and Marc-Olivier Killijian. XPIR: Private information retrieval for everyone. *PoPETs*, 2016(2):155–174, April 2016.

[12] Ishtiyaque Ahmad, Yuntian Yang, Divyakant Agrawal, Amr El Abbadi, and Trinabh Gupta. Addra: Metadata-private voice communication over fully untrusted infrastructure. In *OSDI 21*, 2021.

[13] Asra Ali, Tancrède Lepoint, Sarvar Patel, Mariana Raykova, Phillipp Schoppmann, Karn Seth, and Kevin Yeo. Communication-computation trade-offs in PIR. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 1811–1828. USENIX Association, August 2021.

[14] Sebastian Angel, Hao Chen, Kim Laine, and Srinath T. V. Setty. PIR with compressed queries and amortized query processing. In *2018 IEEE Symposium on Security and Privacy*, pages 962–979. IEEE Computer Society Press, May 2018.

[15] Sebastian Angel and Srinath Setty. Unobservable communication over fully untrusted infrastructure. In *OSDI 16*, pages 551–569, 2016.

[16] Gilad Asharov, Moni Naor, Gil Segev, and Ido Shahaf. Searchable symmetric encryption: optimal locality in linear space via two-dimensional balanced allocations. In Daniel Wichs and Yishay Mansour, editors, *48th ACM STOC*, pages 1101–1114. ACM Press, June 2016.

[17] Michael Backes, Aniket Kate, Praveen Manoharan, Sebastian Meiser, and Esfandiar Mohammadi. Anoa: A framework for analyzing anonymous communication protocols. In *2013 IEEE 26th Computer Security Foundations Symposium*, pages 163–178, 2013.

[18] Amos Beimel, Yuval Ishai, and Tal Malkin. Reducing the servers computation in private information retrieval: PIR with preprocessing. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 55–73. Springer, Heidelberg, August 2000.

[19] Alexander Bienstock, Sarvar Patel, Joon Young Seo, and Kevin Yeo. Near-optimal oblivious key-value stores for efficient PSI, PSU and volume-hiding multi-maps. In *USENIX Security 2023*, 2023.

[20] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1292–1303. ACM Press, October 2016.

[21] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 868–886. Springer, Heidelberg, August 2012.

[22] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In Jacques Stern, editor, *EUROCRYPT'99*, volume 1592 of *LNCS*, pages 402–414. Springer, Heidelberg, May 1999.

[23] Hao Chen, Zhicong Huang, Kim Laine, and Peter Rindal. Labeled PSI from fully homomorphic encryption with malicious security. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1223–1237. ACM Press, October 2018.

[24] Hao Chen, Kim Laine, and Peter Rindal. Fast private set intersection from homomorphic encryption. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1243–1255. ACM Press, October / November 2017.

[25] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Tfhe: fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, 2020.

[26] Seung Geol Choi, Dana Dachman-Soled, S. Dov Gordon, Linsheng Liu, and Arkady Yerukhimovich. Compressed oblivious encoding for homomorphically encrypted search. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2277–2291. ACM Press, November 2021.

[27] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *Journal of the ACM (JACM)*, 45(6):965–981, 1998.

[28] Kelong Cong, Radames Cruz Moreno, Mariana Botelho da Gama, Wei Dai, Ilia Iliashenko, Kim Laine, and Michael Rosenberg. Labeled PSI from homomorphic encryption with reduced computation and communication. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1135–1150. ACM Press, November 2021.

[29] Ivan Damgård and Mats Jurik. A generalisation, a simplification and some applications of Paillier's probabilistic public-key system. In Kwangjo Kim, editor, *PKC 2001*, volume 1992 of *LNCS*, pages 119–136. Springer, Heidelberg, February 2001.

[30] Daniel Demmler, Peter Rindal, Mike Rosulek, and Ni Trieu. Pir-psi: Scaling private contact discovery. *Proceedings on Privacy Enhancing Technologies*, 2018(4):159–178, 2018.

[31] Martin Dietzfelbinger and Stefan Walzer. Efficient gauss elimination for near-quadratic matrices with one short random block per row, with applications. In *27th Annual European Symposium on Algorithms (ESA 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.

[32] Zeev Dvir and Sivakanth Gopi. 2-server PIR with sub-polynomial communication. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 577–584. ACM Press, June 2015.

[33] Klim Efremenko. 3-query locally decodable codes of subexponential length. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 39–44. ACM Press, May / June 2009.

[34] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Paper 2012/144, 2012. `https://eprint.iacr.org/2012/144`.

[35] Nils Fleischhacker, Kasper Green Larsen, and Mark Simkin. How to compress encrypted data. In *EUROCRYPT 2023, Part I*, pages 551–577. Springer, 2023.

[36] Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 303–324. Springer, Heidelberg, February 2005.

[37] Gayathri Garimella, Benny Pinkas, Mike Rosulek, Ni Trieu, and Avishay Yanai. Oblivious key-value stores and amplification for private set intersection. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 395–425, Virtual Event, August 2021. Springer, Heidelberg.

[38] Nethanel Gelernter and Amir Herzberg. On the limits of provable anonymity. WPES '13, page 225–236, New York, NY, USA, 2013. Association for Computing Machinery.

[39] Craig Gentry and Shai Halevi. Compressible FHE with applications to PIR. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 438–464. Springer, Heidelberg, December 2019.

[40] Craig Gentry and Zulfikar Ramzan. Single-database private information retrieval with constant communication rate. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *ICALP 2005*, volume 3580 of *LNCS*, pages 803–815. Springer, Heidelberg, July 2005.

[41] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, August 2013.

[42] Yael Gertner, Yuval Ishai, Eyal Kushilevitz, and Tal Malkin. Protecting data privacy in private information retrieval schemes. In *30th ACM STOC*, pages 151–160. ACM Press, May 1998.

[43] Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 640–658. Springer, Heidelberg, May 2014.

[44] Matthew Green, Watson Ladd, and Ian Miers. A protocol for privately reporting ad impressions at scale. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1591–1601. ACM Press, October 2016.

[45] Jens Groth, Aggelos Kiayias, and Helger Lipmaa. Multi-query computationally-private information retrieval with constant communication rate. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010*, volume 6056 of *LNCS*, pages 107–123. Springer, Heidelberg, May 2010.

[46] Syed Mahbub Hafiz and Ryan Henry. A bit more than a bit is more than a bit better: Faster (essentially) optimal-rate many-server PIR. *PoPETs*, 2019(4):112–131, October 2019.

[47] Ryan Henry. Polynomial batch codes for efficient IT-PIR. *PoPETs*, 2016(4):202–218, October 2016.

[48] Alexandra Henzinger, Matthew M. Hong, Henry Corrigan-Gibbs, Sarah Meiklejohn, and Vinod Vaikuntanathan. One server for the price of two: Simple and fast single-server private information retrieval. In *USENIX Security 2023*, 2023.

[49] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Batch codes and their applications. In László Babai, editor, *36th ACM STOC*, pages 262–271. ACM Press, June 2004.

[50] Stanislaw Jarecki and Xiaomin Liu. Fast secure computation of set intersection. In Juan A. Garay and Roberto De Prisco, editors, *SCN 10*, volume 6280 of *LNCS*, pages 418–435. Springer, Heidelberg, September 2010.

[51] Daniel Kales, Christian Rechberger, Thomas Schneider, Matthias Senker, and Christian Weinert. Mobile private contact discovery at scale. In Nadia Heninger and Patrick Traynor, editors, *USENIX Security 2019*, pages 1447–1464. USENIX Association, August 2019.

[52] Tali Kaufman and Madhu Sudan. Sparse random linear codes are locally decodable and testable. In *48th FOCS*, pages 590–600. IEEE Computer Society Press, October 2007.

[53] Eyal Kushilevitz and Rafail Ostrovsky. Replication is NOT needed: SINGLE database, computationally-private information retrieval. In *38th FOCS*, pages 364–373. IEEE Computer Society Press, October 1997.

[54] Helger Lipmaa. An oblivious transfer protocol with log-squared communication. In *International Conference on Information Security*, pages 314–328. Springer, 2005.

[55] Zeyu Liu and Eran Tromer. Oblivious message retrieval. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 753–783. Springer, Heidelberg, August 2022.

[56] Wouter Lueks and Ian Goldberg. Sublinear scaling for multi-client private information retrieval. In Rainer Böhme and Tatsuaki Okamoto, editors, *FC 2015*, volume 8975 of *LNCS*, pages 168–186. Springer, Heidelberg, January 2015.

[57] Rasoul Akhavan Mahdavi and Florian Kerschbaum. Constant-weight PIR: Single-round keyword PIR via constant-weight equality operators. In *USENIX Security 22*, pages 1723–1740, Boston, MA, 2022.

[58] Samir Jordan Menon and David J Wu. Spiral: Fast, high-rate single-server PIR via FHE composition. In *2022 IEEE Symposium on Security and Privacy*, 2022.

[59] M. Mughees and L. Ren. Vectorized batch private information retrieval. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 437–452, 2023.

[60] Muhammad Haris Mughees, Hao Chen, and Ling Ren. OnionPIR: Response efficient single-server PIR. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2292–2306. ACM Press, November 2021.

[61] Muhammad Haris Mughees, Gonçalo Pestana, Alex Davidson, and Benjamin Livshits. PrivateFetch: Scalable catalog delivery in privacy-preserving advertising, 2021.

[62] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In Jacques Stern, editor, *EUROCRYPT'99*, volume 1592 of *LNCS*, pages 223–238. Springer, Heidelberg, May 1999.

[63] Jeongeun Park and Mehdi Tibouchi. SHECS-PIR: Somewhat homomorphic encryption-based compact and scalable private information retrieval. In Liqun Chen, Ninghui Li, Kaitai Liang, and Steve A. Schneider, editors, *ESORICS 2020, Part II*, volume 12309 of *LNCS*, pages 86–106. Springer, Heidelberg, September 2020.

[64] Sarvar Patel, Joon Young Seo, and Kevin Yeo. Don't be dense: Efficient keyword PIR for sparse databases. In *USENIX Security 2023*, 2023.

[65] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[66] Microsoft SEAL (release 4.1). `https://github.com/Microsoft/SEAL`, January 2023. Microsoft Research, Redmond, WA.

[67] Vitaly Shmatikov and Ming-Hsiu Wang. Measuring relationship anonymity in mix networks. WPES '06, page 59–62, New York, NY, USA, 2006. Association for Computing Machinery.

[68] Nigel P Smart and Frederik Vercauteren. Fully homomorphic simd operations. *Designs, codes and cryptography*, 71:57–81, 2014.

[69] Kevin Yeo. Cuckoo hashing in cryptography: Optimal parameters, robustness and applications. In *CRYPTO 2023*, 2023.

A Somewhat Homomorphic Encryption (SHE)

In this section, we outline two classes of SHE schemes: one with small ciphertext expansion but large noise growth and the other with large ciphertext expansion and small noise growth. The first class (including Regev [65] and BFV [21, 34]) are SHE schemes with small ciphertext expansion (the ratio of ciphertext size to plaintext size), but large noise growth especially for ciphertext-ciphertext multiplication. In contrast, the second class of schemes (including GSW [41]) are SHE schemes with large ciphertext expansion, but very small noise growth for ciphertext-ciphertext expansion. In particular, GSW [41] ensures only additive noise growth whereas the first class of SHE schemes require multiplicative noise growth when performing homomorphic multiplications. Finally, it is shown that protocols can perform operations using ciphertexts from different classes. Recent PIR schemes [60, 58] rely on multiplying ciphertexts from each class resulting in small noise growth.

Regev and BFV Encryption. Many PIR schemes relied upon Regev encryption [65] equipped with homomorphic addition and its extension by Brakerski [21] as well as Fan and Vercauteren [34]

enabling homomorphic multiplication. These schemes are defined over a ring $\mathbb{R} = \mathbb{Z}/(x^n + 1)$ where n is the dimension of the polynomial along with a plaintext and ciphertext modulus q and t respectively. A plaintext value is a polynomial in $R \mod t$ and a ciphertext consists of $\tilde{c} = (c_0, c_1)$ where both polynomials are elements of $R \mod q$. For more information on the details of these schemes, we defer readers to prior works [65, 21, 34]. We will only use them in a blackbox manner with certain properties that we describe next.

First, we describe the noise growth of each homomorphic operation. For ciphertext-ciphertext addition, we note that noise growth is additive. In particular, if we consider two ciphertexts $\tilde{c}_1$ and $\tilde{c}_2$ with error $\mathsf{Err}(\tilde{c}_1)$ and $\mathsf{Err}(\tilde{c}_2)$, then the resulting error is $O(\mathsf{Err}(\tilde{c}_1) + \mathsf{Err}(\tilde{c}_2))$ after homomorphic addition. For ciphertext-plaintext multiplication (also known as absorption) with a ciphertext $\tilde{c}$ with error $\mathsf{Err}(\tilde{c})$ and any plaintext message m, the resulting error is $O(|m| \cdot \mathsf{Err}(\tilde{c}))$. Finally, for ciphertext-ciphertext multiplication, the noise growth becomes $O(t \cdot (\mathsf{Err}(\tilde{c}_1) + \mathsf{Err}(\tilde{c}_2)))$. Note, for a sequence of ciphertext-ciphertext multiplications, the noise growth would grow exponentially in the length of the sequence. As a result, recent PIR schemes avoid these operations.

Finally, these schemes have been shown to emit properties that enable packing multiple plaintext values into a single ciphertext that has been used to reduce PIR request sizes [14].

GSW Encryption. The second class of SHE scheme is the Gentry, Sahai and Waters scheme [41] that can be defined over the same polynomial ring $\mathbb{R} = \mathbb{Z}/(x^n + 1)$, plaintext space $\mathbb{R} \mod t$ and ciphertext space $\mathbb{R} \mod q$. The scheme is parameterized by a base B and length ℓ that provides a trade-off between noise growth and efficiency. Once again, we defer details of the GSW scheme to prior works [41]. Instead, we will only provide details about certain properties that will be leveraged in our work.

In particular, we will rely on the external product operation introduced by Chillotti *et al.* [25]. The input to the external product is a Regev/BFV ciphertext and a GSW ciphertext and the output is a Regev/BFV ciphertext containing the multiplication of the two input ciphertexts. The main advantage of the external product is that noise growth is linear and asymmetric. For a Regev/BFV ciphertext $\tilde{c}_1$ with error $\mathsf{Err}(\tilde{c}_1)$ and a GSW ciphertext $\tilde{c}_2$ with error $\mathsf{Err}(\tilde{c}_2)$, the output of the external product is a Regev/BFV ciphertext with noise $O(B \cdot \mathsf{Err}(\tilde{c}_2) + \mathsf{Err}(\tilde{c}_1))$ requiring $O(\ell)$ polynomial multiplications. As earlier stated, the choice of B and ℓ provide trade-offs between the noise growth and efficiency of the external product operation. Secondly, we note that the noise growth is asymmetric in the sense that noise grows linearly in the Regev/BFV ciphertext and the B multiplicative factor only affects the GSW ciphertext.

B Noise Analysis of PIR Schemes using SHE Composition

W give a high-level overview of recent works such as OnionPIR [60] and Spiral [58] that compose the two classes to obtain efficient single-server PIR constructions.

PIR from SHE Composition. Recent PIR schemes compose the two classes of SHE schemes to obtain fast computation and small communication. These include the more theoretical work of Gentry and Halevi [39] as well as recent practical PIR schemes of OnionPIR [60] and Spiral [58]. These families of PIR schemes enable larger levels of recursion than prior works (such as [14, 13]) due to their superior ablility to minimize noise growth by switching between SHE schemes.

At a high level, these PIR schemes operate as follows. The database is represented as a hypercube with dimension $d_1 \times d_2 \times \ldots \times d_z$. The first dimension is typically large such as $d_1 \in$

$\{128, 512, 1024\}$ while the other dimensions are the same and much smaller $d_2 = \ldots = d_z \in \{2, 4\}$. For convenience, we will denote $d = d_2 = \ldots = d_z$. For a PIR request, the client will upload an encrypted vector of length $d_1 + d_2 + \ldots + d_z = d_1 + (z-1)d$ specifying a single entry in each of the z dimensions. The client uploads these as Regev/BFV encryptions using the packing techniques from Angel *et al.* [14] that can be unpacked by the server. Afterwards, the server applies the first dimension using Regev/BFV encryptions. The remaining Regev/BGV encryptions in the client's request are converted to GSW ciphertexts. Using the result of the first dimension's processing, the remaining $z-1$ levels are handled using external products. Note, the final result is a Regev/BGV encryption that can be reduced using modulus switching before being returned the client.

Keyword PIR Framework [64]. Our work will rely upon the keyword PIR framework [64] that may be built from the recent PIR schemes using SHE composition. A high-level description of this framework is provided in Section 5.1. At a high level, the only difference is that the client request may contain different inputs. However, the server-side processing remains identical, which is the only critical part needed for analyzing noise growth.

Noise Growth. To compute the noise growth of this family of (keyword) PIR schemes, we will first make some assumptions without loss of generality. Suppose that the server has unpacked the request and obtained $d_1 + \ldots + d_z = d_1 + (z-1)d$ ciphertexts. We will assume that the d_1 Regev/BFV ciphertexts for the first level have error e_0. The remaining $(z-1)d$ GSW ciphertexts will have error e_1 that may be different due to translation to between Regev/BFV to GSW schemes. We will also assume that each database entry has norm at most ℓ. For large database entries, each entry is split into smaller parts each of norm at most ℓ.

Next, we can compute the noise growth of the above family of PIR schemes. We start by analyzing the first dimension processing where the result is n/d_1 BFV/Regev ciphertexts with error $O(d_1 \cdot \ell \cdot e_0)$ as each of the n/d_1 output ciphertexts are the result of summing $d_1 - 1$ ciphertext-plaintext multiplications with the original z BFV/Regev encryptions with error e_0.

We move onto the remaining $z-1$ dimensions. Consider the processing of the second dimension. The output will be $n/(d_1 \cdot d)$ ciphertexts where each ciphertext is the sum of d outputs from the external product operations. After the external product, the noise is $O(Be_1 + d_1 \ell e_0)$. As we do d additions, the noise of each ciphertext after the second dimension processing is $O(dBe_1 + d_1 d\ell e_0)$. Repeating the analysis for all z dimensions, we obtain the noise of the final ciphertext is

$$O\left(\sum_{i=1}^{z-1} d^i Be_1 + \sum_{i=1}^{z-1} d^i d_1 \ell e_0\right).$$

Assuming that $d \geq 2$, we get that the final noise is

$$O((n/d_1)Be_1 + n\ell e_0)$$

since $d \geq 2$ and $d_1 \cdot d^{z-1} = \Theta(n)$.

C Noise Analysis of Our PIR Schemes with Compression

In this section, we analyze the noise growth of our PIR scheme that utilizes LSObvDecompress for request compression and LSObvCompress for response compression. We will build on top of the analysis from Appendix B. Recall that if we assume that the d_1 ciphertexts for the first dimension

have error e_0 and the $(z-1)d$ ciphertexts for the other dimensions have error e_1, then the final ciphertext has noise $O((n/d_1)Be_1 + n\ell e_0)$.

Let us consider the result of executing $\mathsf{LSObvDecompress}$ for request compression. We note that this incurs a multiplicate $O(w)$ noise growth. As a result, we can imagine that the ciphertexts now have error $O(we_0)$ and $O(we_1)$ respectively. After the server processes the z dimensions, the resulting ciphertext has noise $O((n/d_1)Be_1w + n\ell e_0 w)$. Therefore, if we only apply $\mathsf{LSObvDecompress}$ this increases the noise growth by a multiplicative $O(w)$ factor.

Next, let us consider applying $\mathsf{LSObvCompress}$ additionally for response compression. Note, this would result in another $O(w)$ multiplicative factor in noise growth. The final ciphertexts would have noise $O((n/d_1)Be_1w^2 + n\ell e_0 w^2)$, which is $O(w^2)$ larger than before. Finally, we note that if one had applied only $\mathsf{LSObvCompress}$ for response compression without request compression, then the final noise would instead be $O((n/d_1)Be_1w + n\ell e_0 w)$, which is only a $O(w)$ multiplicative factor larger.

D Random Band Matrix Analysis

In this section, we analyze the variant of random band matrices from Section 3.3 where each column is generated with a random w-bit band. As these are transposes of random band matrices, we know these variants will have unique solutions. It remains to show that the running time of these random band matrices runs in time $O(nw)$ similar to the original random band matrices from [31].

To do this, we show that each row will consist of exactly one continguous section of non-zero entries of length $O(w)$. We couple the process of generating random band matrices with random column vectors as two-dimensional balls-into-bins allocation (see [16] for more details). In particular, we model each of the t columns as t lists of w items. There exists $m = (1+\epsilon)t$ entries corresponding to each of the m rows. Each of the t lists are assigned to a random entry from $[m-w+1]$. If the i-th list is assigned to entry $j \in [m-w+1]$, then one of the w items in the list are placed into each of the entries $\{j, j+1, \ldots, j+w-1\}$. Note, the maximum load of any of m entries is equivalent to the largest consecutive section of non-zero entries in any of the m rows of the generated random band matrix after sorting by column starting location.

Prior work [16] studied the setting where each of the t lists picked one of the m entries uniformly at random. We adapt the analysis for the slightly skewed distribution used for random band matrices in our work where only one of the first $m-w$ entries are chosen uniformly at random.

Proof of Theorem 1. We use the coupling described above. Therefore, it suffices for us to analyze only two-dimensional balls-into-bins allocations. We denote binary random variables $X_{i,j}$ to be whether the random band of the i-th column will overlap with the j-th row. Therefore, $X_{i,j} = 1$ if this event is true and $X_{i,j} = 0$ otherwise. Note, $X_{i,j} = 1$ if and only if the i-th column's random band starts in the set of row indices $\{j-w+1, j-w+2, \ldots, j\}$. In other words, $\mathsf{E}[X_{i,j} = 1] \leq w/(m-w+1)$. Let B_j be the total number of columns whose random bands overlap with the j-th row. By linearity of expectation, we get that

$$B_j = \sum_{i \in [t]} \mathsf{E}[X_{i,j}] \leq \frac{tw}{m-w+1}.$$

Note that each $X_{i,j}$ is an independent random variable. Therefore, we can apply Chernoff bounds

to get that

$$\Pr\left[B_j > 3 \cdot \frac{tw}{m-w+1}\right] \leq 2^{-tw/(m-w+1)}.$$

Next, we apply a Union bound over all m rows to get that B_j for all $j \in [m]$ is upper bounded by the same value with probability at most $m \cdot 2^{-tw/(m-w+1)}$. Finally, by noting that $m = (1+\epsilon)t$ for some constant $\epsilon > 0$ and picking $w = O(\lambda + \log t)$, we get that each row has a band length of at most $O(w)$ except with probability $2^{-\lambda}$. □

E Single-Server Batch PIR with Request and Response Compression

In this section, we present a single-server batch PIR scheme that uses both LSObvDecompress and LSObvCompress to compress requests and responses respectively.

We leverage one important aspect of the keyword PIR framework from [64], presented in Section 5. In this framework, the smallest dimension of the request vectors, say z w.l.o.g., is typically very small, of size $d_z = 2$ or $d_z = 4$. As described in Section 5, if the client wishes to construct a dummy "zero" request, it can set $\mathbf{v}_z$ to the all-0 vector of length 2 or 4, and all other vectors $\mathbf{v}_1, \ldots, \mathbf{v}_{z-1}$ arbitrarily. This is because when $\mathbf{v}_z$ is applied at the last level to obtain the final ciphertext, it will always produce a ciphertext that encrypts zero.

Our Construction. We use a similar approach as Section 5.1 and Section 5.2. However, we must modify request compression to be compatible with response compression. Recall that LSObvCompress for response compression requires that the dummy responses are zero encryptions.

To achieve this, we will set the last dimension of all 0.5ℓ dummy requests to be the all-zero vector. In the keyword PIR framework [64], the resulting response will be a zero encryption. The relevant indices become all ℓ real request vectors and the vector corresponding to the last dimension for each of the 0.5ℓ dummy requests. Afterwards, the client can execute LSObvDecompress to compress the request. The server decompresses the request using LSObvDecompress, processes the requests to compute $B = 1.5\ell$ responses and compresses using LSObvCompress as there are at most ℓ zero entries. Finally, the client decompresses to obtain the ℓ entries.

Efficiency. The client compresses a vector of length $1.5\ell \cdot (d_1 + \ldots + d_z)$ with $|I| = \ell \cdot (d_1 + \ldots + d_z) + 0.5\ell d_z$. Therefore, the compressed request consist of $\lceil 1.05|I|/r \rceil$ if r requests can fit into a single ciphertext using the request packing techniques in [14]. The compressed response has the identical size as the PIR scheme from Section 5.1.

F Vectorized Batch PIR with Compression

In this section, we present the full description of our construction of vectorized batch PIR that leverage LSObvCompress for response compression (we briefly sketched the ideas in Section 5.3). We start by presenting an overview of the prior vectorized batch PIR [59] before presenting our construction.

Vectorized Batch PIR. Mughees and Ren [59] proposed a batch PIR scheme that cleverly utilizes ciphertext vectorization to improve computation and reduce response sizes. In particular, they utilize *Single-Instruction-Multiple-Data* (SIMD) techniques for efficiently performing homomorphic

Algorithm 7 LSObvCompress.ObvCompress algorithm

Input: $\mathbf{pk}_{\mathcal{E}}, \tilde{c}, t, R$: public homomorphic encryption key, vector of n SIMD ciphertexts with all slots populated with the same plaintext, number of non-zero plaintext entries, and randomness.
Output: $\hat{\mathbf{c}}$: compressed ciphertexts of $\tilde{c}$.

$m \leftarrow (1+\epsilon)t$
$\mathbf{M} \leftarrow 0^{m \times n}$
for $i = 1, \ldots, n$ **do**
  $\mathbf{v_i} \leftarrow \mathsf{GenRandVec}(i, m; R)$ ▷ Random column band
  $\mathbf{M}[:][i] \leftarrow \mathbf{v_i}$ ▷ Set i-th column to $\mathbf{v_i}$
$d \leftarrow$ number of SIMD slots per ciphertext
$\hat{\mathbf{c}} \leftarrow$ length $\lceil (1+\epsilon)t/d \rceil$ ciphertexts encrypting 0 slots
for $i = 1 \ldots n$ **do**
  for $j \in$ non-zero indices of $\mathbf{M}[:][i]$ **do**
    $a \leftarrow \lfloor (j-1)/d \rfloor + 1$
    $b \leftarrow ((j-1) \mod d) + 1$
    $msk \leftarrow$ one-hot binary mask with b-th slot set to 1
    $\hat{\mathbf{c}}[a] \leftarrow \hat{\mathbf{c}}[a] + msk \cdot \tilde{c}[i]$
return $\hat{\mathbf{c}}$

operations [68]. SIMD encodes multiple database entries into a single ciphertext and operates on all of them simultaneously. In more detail, each plaintext polynomial consists of multiple SIMD slots in which multiple database entries can be encoded. Homomorphic operations can be applied to the ciphertexts and the plaintexts in SIMD fashion, i.e. single ciphertext-plaintext absorption corresponds to SIMD slot-wise ciphertext-plaintext absorption. Additionally, there are ciphertext rotation operations to manipulate and rotate the SIMD slots.

The vectorized batch PIR scheme [59] also builds on top of the cuckoo hashing based framework by Angel *et al.* [14]. In the cuckoo hashing framework, consider the point where the server completes processing the B PIR requests and holds the B PIR responses. Each valid response ciphertext will contain the desired entry in an arbitrary SIMD slot. The vectorized batch PIR merges multiple response ciphertexts using SIMD operations to reduce the total response size.

We now describe the merging of response ciphertexts in more detail. Let ℓ be the number of queries in the batch, $B = 1.5\ell$ be the number of single-query PIR buckets, and d be the number of SIMD slots. For simplicity, we will assume d to be a power of two and each database entry can be encoded in a single SIMD slot. Let $\tilde{c}_1, \ldots \tilde{c}_B$ be the response ciphertexts after the PIR servers process the queries, where each $\tilde{c}_i$ encrypts a length d vector of the form $[0, \ldots, 0, p_i, 0, \ldots, 0]$. Note that $\tilde{c}_i$ will contain at most one non-zero SIMD slot. If $\tilde{c}_i$ is the PIR response of a dummy request, then every slot of $\tilde{c}_i$ will be zero. Note that the server knows that at most ℓ PIR responses contain a non-zero slot and the remaining $B - \ell$ PIR responses will encrypt an entirely zero vector.

The vectorized scheme [59] converts $[0, \ldots, 0, p_i, 0, \ldots, 0]$ into $[p_i, \ldots, p_i]$ using ciphertext rotations. Without loss of generality, suppose that p_i is in the first slot. The server first rotates the ciphertext by 1 position to obtain a new ciphertext that contains p_i in the second slot. It then homomorphically adds the two ciphertexts to obtain a ciphertext that has p_i in the first two slots. Afterwards, it rotates the resulting ciphertext by 2 positions to obtain a new ciphertext that contains p_i in the third and the fourth slots. Again, it homomorphically adds the two ciphertexts to

obtain a ciphertext that has p_i in the first four slots. Repeating this $\lceil \log_2 d \rceil$ times, we will obtain a ciphertext that has p_i in each of the d SIMD slots. This process works regardless of the original position of p_i. Afterwards, each $\tilde{c}_i = [p_i, \ldots, p_i]$ with the same value in every slot.

Next, the scheme masks each ciphertext to ensure that only a single, but predictable, slot may contain a non-zero value. The goal is to transform $\tilde{c}_i = [p_i, \ldots, p_i]$ to contain p_i only in the slot with i-th index. For example, we will multiply $\tilde{c}_1 \cdot [1, 0, \ldots, 0]$ to obtain $[p_1, 0, \ldots, 0]$. When $d < n$, we note that the i-th index would rotate from the last slot back to the first slot. For example, the $\tilde{c}_{d+1}$ will also be multiplied by $[1, 0, \ldots, 0]$ as the $(d+1)$-th slot is equivalent to the first slot. Afterwards, we get the following:

$$\begin{aligned}
\tilde{c}_1 &= [p_1, 0, 0, \ldots, 0] \\
\tilde{c}_2 &= [0, p_2, 0, \ldots, 0] \\
&\ldots \\
\tilde{c}_d &= [0, 0, 0, \ldots, p_d] \\
\tilde{c}_{d+1} &= [p_{d+1}, 0, 0, \ldots, 0] \\
\tilde{c}_{d+2} &= [0, p_{d+2}, 0, \ldots, 0] \\
&\ldots
\end{aligned}$$

Finally, merge each consecutive group of d ciphertexts into one by homomorphic additions to obtain the following:

$$[p_1, p_2, \ldots, p_d], [p_{d+1}, p_{d+2}, \ldots, p_{2d}], \ldots$$

Thus, this results in $\lceil B/d \rceil$ response ciphertexts.

The vectorized scheme by Mughees and Ren [59] can significantly reduce response size for small entries and large number of SIMD slots d. However, it still shares the same inefficiency as in Angel *et al.*'s framework. Out of the $B = 1.5\ell$ PIR requests, $B - \ell = 0.5\ell$ requests will correspond to dummy requests. Thus, 0.5ℓ PIR responses will be zero and at most ℓ will contain a non-zero slot. The vectorized scheme pessimistically encodes all 1.5ℓ responses requiring $\lceil B/d \rceil = \lceil 1.5\ell/d \rceil$ response ciphertexts. This means that these "dummy", zero responses must still occupy SIMD slots in the final responses. We show that we can apply LSObvCompress to effectively remove these dummy, zero PIR responses.

Our Construction. In the context of LSObvCompress, the input ciphertext vector consists of $n = B = 1.5\ell$ ciphertexts to compress, where there are at most $t = \ell$ ciphertexts that encrypt non-zero plaintext entries. Our goal is to compress this down to $(1+\epsilon)\ell$ SIMD slots, which implies we will need a total of $h = \lceil (1+\epsilon)\ell/d \rceil$ ciphertexts.

The main idea behind our usage of LSObvCompress stays unchanged from the large entry setting in Section 5.1. We present a vectorized version of LSObvCompress that efficiently packs multiple plaintext values into a single ciphertext. For each of the B response ciphertexts $\tilde{c}_i$, suppose we completed the first step of the vectorized PIR scheme [59]. Thus, all SIMD slots are populated with the plaintext entry, $\tilde{c}_i = [p_i, \ldots, p_i]$.

Next, we will apply our technique to compress using the matrix $\mathbf{M}$ that is described in Section 3.3 where $\mathbf{M}$ is the transpose of a random band matrix with dimensions $B \times m$ where $m = (1 + \epsilon)\ell$. For convenience, we will denote the plaintext vector as $p = [p_1, \ldots, p_B]^T$. Using ideas from LSObvCompress, our goal is to compute the $m = (1+\epsilon)\ell$ SIMD slot values from the matrix-vector

multiplication, $M \cdot p$. However, we would like to do this such that these are encoded in the slots of $h = \lceil m/d \rceil = \lceil (1+\epsilon)\ell/d \rceil$ ciphertexts.

At a high level, our vectorized version of LSObvCompress will encode the $m = (1+\epsilon)\ell$ SIMD slots such that the first d values of $M \cdot p$ will be in the first ciphertext, the next d values of $M \cdot p$ will be in the second ciphertext and so forth. We compute this vectorized encoding as follows. For each ciphertext $i \in [B]$, let $j_{i_1}, \ldots, j_{i_{g(i)}}$ be the indices of non-zero entries in the i-th column of $\mathbf{M}$. At a high level, we may imagine that the SIMD slots of the h ciphertexts are flattened, i.e. the left-hand-side column vector is of length $d \cdot h$. With this formulation, computing and adding $\mathbf{M}[:][i] \cdot p_i$ corresponds to adding the encryption of p_i to flattened slots with indices $j_{i_1}, \ldots, j_{i_{g(i)}}$. More precisely, for each $j_{i_k} \in \{j_{i_1}, \ldots, j_{i_{g(i)}}\}$, we compute a pair of indices $(a, b) = (\lfloor (j_{i_k} - 1)/d \rfloor + 1, (j_{i_k} - 1 \mod d) + 1)$, and homomorphically add to $\hat{\mathbf{c}}_a$ a multiplication of $\tilde{c}_i$ by a one-hot binary mask that is 1 only at the bth slot.

The decompression works identically, except the algorithm now flattens the decrypted SIMD slots to $(1+\epsilon)\ell$ individual plaintext entries before solving the linear system.

We formally present our vectorized version of LSObvCompress in Algorithm 7. All differences with the original algorithm are highlighted in blue. We note the changes only enable vectorization and the core compression ideas remain identical. One can obtain the prior LSObvCompress algorithm by setting the number of SIMD slots to be $d = 1$.

We point out that the compression algorithm will incur $O(B \cdot w)$ homomorphic operations (where w is the column band width), which is asymptotically the same running time as the original LSObvCompress with $B = n$ input ciphertexts. The decompression algorithm requires $O(\ell \cdot w)$ time.

F.1 Experimental Evaluation

We implement vectorized LSObvCompress using Microsoft SEAL [66] library and present our results in Figure 9. Following prior works [59], we use the polynomial degree of 8192, ciphertext modulus of 200 bits, and plaintext modulus of 20 bits for our experimental setup. We fix the entry size to 256 bytes. Note that the results in Figure 9 correspond to ideal cases where the noise level budget is sufficient enough to accommodate additional homomorphic operations incurred by our scheme. In particular, this means that we may not be able to directly plug in our scheme to the vectorized PIR scheme [59], as different set of parameters may have to be chosen to handle the extra noise growth.

G Supplementary Material for Labeled PSI

We start by presenting the formal functionality of unbalanced labeled private set intersection (PSI) in Figure 10. Afterwards, we present our two constructions for labeled PSI. The first uses a generic transformation relying on our improved batch PIR from Section 5. The second improves upon oblivious polynomial evaluation protocols using LSObvDecompress for request comrpession.

G.1 Labeled PSI from Batch Keyword PIR and Oblivious PRF

Here, we recall the generic composition from [36] to build labeled PSI from batch keyword PIR and an oblivious pseudorandom function (OPRF). First, we describe the OPRF instantiation that we

Compression Size	Size	Compression Time	Decompression Time	Total Time
$1.05t$	$\mathbf{t = 512, n = 768}$	12.8 s	0.1 s	12.9 s
	$\mathbf{t = 1024, n = 1536}$	25.1 s	0.2 s	25.3 s
	$\mathbf{t = 2048, n = 3072}$	50.5 s	0.5 s	60.0 s
	$\mathbf{t = 4096, n = 6144}$	100.1 s	0.9 s	101.0 s
$1.07t$	$\mathbf{t = 512, n = 768}$	10.3 s	0.1 s	10.4 s
	$\mathbf{t = 1024, n = 1536}$	20.9 s	0.2 s	21.1 s
	$\mathbf{t = 2048, n = 3072}$	41.6 s	0.4 s	42.0 s
	$\mathbf{t = 4096, n = 6144}$	83.8 s	0.9 s	84.7 s

Figure 9: Performance of vectorized LSObvCompress evaluated on various values of t and n. We use fixed entry size of 256 B.

Parameters: There are two parties, a receiver and a sender. The honest receiver and sender have respective set sizes $n_{\mathrm{X}}, n_{\mathrm{Y}}$. If the receiver or sender is maliciously corrupt, then their set size is n'_{X} or n'_{Y}, respectively.

Functionality:

- On input $(\textsc{Receive}, \mathsf{sid}, X)$ from the receiver where $X \subseteq \{0,1\}^*$, ensure that $|X| \leq n_{\mathrm{X}}$ if the receiver is honest and $|X| \leq n'_{\mathrm{X}}$ otherwise. Give $(\textsc{Receiver-Input}, \mathsf{sid})$ to the sender.
- Thereafter, on input $(\textsc{Send}, \mathsf{sid}, (Y, \{L_y \in \{0,1\}^\ell \mid y \in Y\}))$ from the sender where $Y \subseteq \{0,1\}^*$, ensure that $|Y| \leq n_{\mathrm{Y}}$ if the sender is honest and $|Y| \leq n'_{\mathsf{Y}'}$ otherwise. Give output $(\textsc{Output}, \mathsf{sid}, \{(x, L_x) \mid x \in X \cap Y\}$ to the receiver.

Figure 10: Ideal Functionality $\mathcal{F}_{\mathsf{ul\text{-}psi}}$ of unbalanced labeled private set intersection.

use for the transformation. The formal oblivious PRF functionality used in our unbalanced labeled PSI protocol can be found in Figure 11.

G.1.1 OPRF with Malicious Security

An OPRF allows the receiver to input set X and learn the set of pseudo-random outputs $\{F_k(x) \mid x \in X\}$, where F is a PRF, and k is known to the sender. For security, both the sender and receiver should learn nothing else (except the sender learns the size of X).

In this paper, we will use the Diffie-Hellman based OPRF protocol of [50] that computes the function $F_\alpha(x) = H'(H(x)^\alpha)$, where H, H' are hash functions modeled as random oracles. We take a description of this OPRF almost verbatim from [23]: Let G be a cyclic group with order q, where the One-More-Gap-Diffie-Hellman (OMGDH) problem is hard. H is a random oracle hash function with range $\mathbb{Z}_q^*$. The sender has a key $\alpha \in \mathbb{Z}_q^*$ and the receiver has a set of inputs X. In the OPRFRequest algorithm formally described in Algorithm 8, the receiver first samples $\beta_i \leftarrow \mathbb{Z}_q^*$ for each $i = 1 \ldots |X|$ and sends $\{H(x_i)^{\beta_{x_i}} \mid i = 1 \ldots |X|\}$ to the sender. Next, in the OPRFAnswer algorithm formally described in Algorithm 9, the sender on input its PRF key α responds with $\{(H(x_i)^{\beta_i})^\alpha \mid i = 1 \ldots |X|\}$. Finally, in the OPRFProc algorithm formally described

Parameters: There are two parties, a sender and a receiver. The receiver has set size n if honest and n' otherwise. Let $\mathsf{out} \in \mathbb{Z}$ be the bit length.

Functionality:

- The functionality first samples random function $F : \{0,1\}^* \to \{0,1\}^{\mathsf{out}}$.
- Subsequently, on input $(\textsc{Sender}, \mathsf{sid}, Y)$ from the sender where $Y \subseteq \{0,1\}^*$, the functionality returns $\{F(y) \mid y \in Y\}$ to the sender.
- Next, on input $(\textsc{Receive}, \mathsf{sid}, X)$ from the receiver where $X \subseteq \{0,1\}^*$, ensure that $|X| \leq n$ if the receiver is honest and $|X| \leq n'$ otherwise. The functionality returns $(\textsc{Receiver-Input}, \mathsf{sid})$ to the sender.
- Thereafter, on input $(\textsc{Send}, \mathsf{sid})$ from the sender, the functionality returns $\{F(x) \mid x \in X\}$ to the receiver.

Figure 11: Ideal Functionality $\mathcal{F}_{\mathsf{oprf}}$ for batch Oblivious PRF.

Algorithm 8 OPRFRequest algorithm

Input: X: receiver's (ordered) set of items.
Output: $((\beta_1, \ldots, \beta_{|X|}), \mathsf{req})$: list of temporary exponents and list of requests.

```
req ← ⊥
for i = 1 ... |X| do
    β_i ← Z_q^*
    req.insert(H(X[i])^{β_i})
return ((β_1, ..., β_{|X|}), req)
```

in Algorithm 10, the receiver outputs $H'(H(x_i)^\alpha) = H'((H(x_i)^{\beta_i})^\alpha)^{1/\beta_i})$ for each $x \in X$ (each output consists of a hash and an encryption key, to be used in the unbalanced labeled PSI protocol).

The outer hash function H' is used to map the group element to a sufficiently long bit string, and is modeled as a random oracle to help facilitate extraction in the malicious setting. In particular, by observing the queries made to $H(x_i)$, the simulator can collect a list of pairs $\{(x_i, H(x_i)\}$ which are known to the receiver. From this set the simulator can compute the set $A = \{(x_i, H(x_i)^\alpha)\}$. For some subset of the $H(x_i)$, the receiver sends $\{H(x_i)^{\beta_i}\}$ to the simulator, who sends back $\{H(x_i)^{\beta_i \alpha}\}$. For the receiver to learn the OPRF value for x_i, it must send $H(x_i)^\alpha$ to the random oracle H'. At this time, the simulator extracts x_i if $(x_i, H(x_i)^\alpha) \in A$. Although this OPRF does not facilitate extracting all x_i at the time the first message is sent, extraction is performed before the receiver learns the OPRF value, which will be sufficient for our purposes.

In the context of our unbalanced labeled PSI protocol, this OPRF has the property that the sender can use the same key with multiple receivers. This allows the sender, who has a large and often relatively static set, to pre-process its set only once.

Algorithm 9 $\mathsf{OPRFAnswer}$ algorithm

Input: (α, req): PRF key and list of requests.
Output: resp: list of responses.

$\mathsf{resp} \leftarrow \perp$
for $i = 1 \dots |\mathsf{req}|$ **do**
  $\mathsf{resp.insert}(\mathsf{req}[i]^{\alpha})$
return resp

Algorithm 10 $\mathsf{OPRFProc}$ algorithm

Input: $(X, (\beta_1, \dots, \beta_{|X|}), \mathsf{resp})$: receiver's (ordered) input set, list of temporary exponents and list of responses.
Output: X': mapping from items to OPRF outputs.

$X' \leftarrow \perp$
for $i = 1 \dots |\mathsf{resp}|$ **do**
  $(\hat{x}_i, k_i) \leftarrow H'(\mathsf{resp}[i]^{1/\beta_i})$
  $X'.\mathsf{mapInsert}(X[i], (\hat{x}_i, k_i))$
return X'

G.1.2 Unbalanced Labeled PSI Transformation

Using a maliciously secure OPRF and a batch keyword PIR scheme, we now describe the generic construction of [36] to get unbalanced labeled PSI. At a high-level, the sender and receiver use an OPRF to compute a (pseudo-random) hashed item and encryption key associated with each real item in their sets. The sender then builds a *pseudo-database* of key-value pairs with its hashed items as the keys and the encryption of the corresponding real labels under the corresponding encryption keys as the values. Then, the receiver uses the batch keyword PIR protocol to query for the hashed items in its set, and decrypts the labels using the corresponding encryption keys from the OPRF.

Algorithm 11 formally describes the algorithm the sender uses to build their pseudo-database of hashed items and encrypted labels. First, the sender invokes $\mathcal{F}_{\mathsf{oprf}}$ on each $y \in Y$ and receive the output $(\hat{y}, k_y)$. Then from these outputs and the set of labels $\{L_y \mid y \in Y\}$ it constructs a pseudo-database $\mathsf{DB} = \{(\hat{y}, \mathsf{Enc}(k_y, L_y)) \mid y \in Y\}$.

Next, the receiver's query algorithm is formally described in Algorithm 12. The receiver invoke $\mathcal{F}_{\mathsf{oprf}}$ on each input item x to receive $(\hat{x}, k_x)$ Then, it sends a batch keyword PIR request on input $Q' = \{\hat{x} \mid x \in X\}$.

Then, the server's answer algorithm, formally described in Algorithm 13 answers the receiver's batch keyword PIR request using its pseudo-database DB.

Finally, the algorithm the receiver uses to decrypt the server's response is formally described in Algorithm 14. The receiver uses the batch keyword PIR to obtain the payload for each $\hat{x}$ that was in DB, $\mathsf{Enc}(k_x, L_x)$, and decrypts it with the corresponding encryption key k_x from the OPRF invocation to output $\{(x, L_x) \mid \hat{x} \in X' \cap Y'\} = \{(x, L_x) \mid x \in X \cap Y\}$.

Now, we prove the following theorem showing that our labeled PSI is private against a malicious sender and secure against a malicious receiver.

Theorem 2. ULPSI *securely realizes* $\mathcal{F}_{\mathsf{ul\text{-}psi}}$ *with privacy against a malicious sender and security against a malicious receiver in the* $\mathcal{F}_{\mathsf{oprf}}$*-hybrid model.*

Algorithm 11 BuildPseudoDB algorithm

Input: $\{(y, L_y) \mid y \in Y\}$: Sender's input set of items and associated labels.
Output: DB: pseudo-database of hashed items and encrypted labels.

```
DB ← ⊥
Invoke F_oprf on input Y to receive {(ŷ, k_y) | y ∈ Y}
for (y, L_y) : y ∈ Y do
    ct_y ← Enc(k_y, L_y)
    DB.mapInsert(ŷ, ct_y)
return DB
```

Algorithm 12 ULPSIQuery algorithm

Input: X: Receiver's input set of items.
Output: (K, req): mapping from items to encryption keys and batch keyword PIR request.

```
K ← ⊥
Q ← ⊥
Invoke F_oprf on input X to receive {(x̂, k_x) | x ∈ X}
for x ∈ X do
    K.mapInsert(x, k_x)
    Q.insert(x̂)
req ← BatchKWPIR.query(Q)
return (K, req)
```

Proof. We recall from [36] that privacy against a malicious sender follows immediately since the adversary learns nothing from the invocation of $\mathcal{F}_{\mathsf{oprf}}$, and what the adversary receives from the receiver from the keyword batch PIR invocation reveals nothing about its query (and thus its set X). Correctness in the real world if the sender is semi-honest easily follows from the correctness of the underlying keyword batch PIR as well as the correctness and security of $\mathcal{F}_{\mathsf{oprf}}$: If $x \notin Y$, then $\hat{x}$ will not be a key of DB (except with negligible probability), due to the pseudo-randomness of the OPRF. Otherwise, if $x \in Y$ then $\hat{x}$ will be a key of DB from the correctness of the OPRF, and so the keyword batch PIR will return the corresponding encrytped label $\mathsf{Enc}(k_x, L_x)$, which the receiver can decrypt with k_x (also correctly computed from the OPRF) to get L_x.

We now prove security against a malicious receiver, by describing a simulator $\mathcal{S}$.

- $\mathcal{S}$ first emulates $\mathcal{F}_{\mathsf{oprf}}$ and when the malicious receiver invokes it on input X, it returns random $\{(\hat{x}, k_x) \mid x \in X\}$.
- Then, it sends X to the unbalanced labeled PSI functionality, $\mathcal{F}_{\mathsf{ul\text{-}psi}}$, and receives $\{(x, L_x) \mid x \in X \cap Y\}$.
- Finally, it builds a size n_{Y} pseudo-database DB using first the key-value pairs $\{(\hat{x}, \mathsf{Enc}(k_x, L_x)) \mid x \in X \cap Y\}$ and then $(r, \mathsf{Enc}(k, 0))$, for random r, k, for the remaining $n_{\mathrm{Y}} - |X \cap Y|$. If it fails, it outputs $\perp$ to the receiver.
- Then, on input the keyword batch PIR request from the malicious receiver, the simulator returns the honest response using pseudo-database DB to the receiver.

Algorithm 13 ULPSIAnswer algorithm

Input: $(\mathsf{req}, \mathsf{DB})$: batch keyword PIR request and pseudo-database.
Output: resp: batch keyword PIR response.

$\mathsf{resp} \leftarrow \mathsf{BatchKWPIR.answer}(\mathsf{DB}, \mathsf{req})$
return resp

Algorithm 14 ULPSIDecrypt algorithm

Input: (K, resp): mapping from items to encryption keys and batch keyword PIR response.
Output: int: output intersection items and associated labels.

$\mathsf{int} \leftarrow \perp$
$\{(x, \mathsf{ct}_x) \mid x \in X \cap Y\} \leftarrow \mathsf{BatchKWPIR.decrypt}(\mathsf{resp})$
for $(x, \mathsf{ct}_x) : x \in X \cap Y$ **do**
    $L_x \leftarrow \mathsf{Dec}(K[x], \mathsf{ct}_x)$
    $\mathsf{int.insert}((x, L_x))$
return int

To see why the simulation works, we proceed with a hybrid argument. Hybrid $\mathcal{H}_0$ is the real world. Hybrid $\mathcal{H}_1$ is the real world except for all $y \notin X$ (where this X is the set input by the receiver to $\mathcal{F}_{\mathsf{oprf}}$), the receiver replaces the key-value pairs of y in DB with $(r, \mathsf{Enc}(k, 0))$ for random r, k. It is easy to see that $\mathcal{H}_0$ is indistinguishable from $\mathcal{H}_1$ because (i) r, k are outputs of $\mathcal{F}_{\mathsf{oprf}}$ unknown to the receiver, and are thus uniformly random; (ii) by reducing to the security of the encryption scheme, replacing $\mathsf{Enc}(k, L_y)$ with $\mathsf{Enc}(k, 0)$ is indistinguishable. Note at this point that even though the keyword batch PIR database encoding algorithm may fail with non-negligible probability, we have shown that this failure is simply a function of the pseudorandom ciphertexts of the database, and not the underlying items $\{(y, L_y) \mid y \in Y \setminus X\}$. Observe that $\mathcal{H}_1$ is in fact the ideal (simulated) world, and thus the proof is complete.

□

G.2 Improving Oblivious Polynomial Evaluation with LSObvDecompress

Now we show that we can use LSObvDecompress to reduce the receiver-to-sender communication of the unbalanced labeled PSI scheme of [28]. We first provide an overview of their scheme.

G.2.1 Overview of [28]

As in the construction from batch keyword PIR and OPRF of Section G.1, the receiver and sender first both run an OPRF on their items to obtain a hash and an encryption key, the latter of which the sender uses to encrypt the corresponding item label. The receiver thus obtains $X' = \{(\hat{x}, k_x) \mid x \in X\}$ and the sender obtains $\mathsf{DB} = \{(\hat{y}, \mathsf{Enc}(k_y, L_y)) \mid y \in Y\}$. Then, they use the same cuckoo hashing technique of Angel *et al.* [14] in the batch PIR setting, in which using three hash functions h_1, h_2, h_3, the sender places each item $\hat{y}$ of DB in three different bins out of $1.5 \cdot |X|$ total bins, and the receiver places each item $\hat{x}$ of X' in a single one of the $1.5 \cdot |X|$ total bins so that no bin has more than one item. Then, for every bin B in which there is at most one $\hat{x}$, the receiver and sender essentially compute the intersection $\hat{x} \cap \{\hat{y} \mid y \in Y \cap B\}$.

More specifically, the sender first interpolates the polynomial satisfying the following:

$$G(x) = \begin{cases} \mathsf{Enc}(k_y, L_y) & \text{if } x = \hat{y} : y \in Y \cap B \\ \text{random field element} & \text{otherwise} \end{cases}$$

Then, using a FHE scheme $\mathcal{E}$, the receiver encrypts $\hat{x}$ and sends the ciphertext to the sender (where the receiver encrypts 0 for empty bins), who returns the homomorphic evaluation of the polynomial G on $\hat{x}$. Next, the receiver first decrypts this returned FHE ciphertext, then using k_x from the OPRF output on x, attempts to decrypt the inner ciphertext. If AEAD is used for this ciphertext, then if $\hat{x} \in \{\hat{y} \mid y \in Y \cap B\}$, the receiver will obtain L_x; otherwise, the decryption will output $\perp$.

It is clear that correctness holds. For security, the sender only sees OPRF queries for $|X|$ inputs, which reveal nothing about X, and then $1.5 \cdot |X|$ FHE ciphertexts. The receiver only sees AEAD ciphertexts encrypted with unknown random keys (from the OPRF) for items that are not in its set.

G.2.2 Applying LSObvDecompress and LSObvCompress

As we observed with the cuckoo hashing framework for batch PIR, for the $0.5 \cdot |X|$ bins in which there is no $\hat{x}$, it does not matter what encrypted value the receiver gives to the sender. Thus, we can use LSObvDecompress to compress the $1.5 \cdot |X|$ total ciphertexts with respect to only the $|X|$ indices in which the corresponding bin has some $\hat{x}$, resulting in only $(1 + \varepsilon)$ encrypted values sent by the receiver. Since the sender ends up just getting $(1 + \epsilon) \cdot |X|$ ciphertexts, security holds as before. As we show in our experiments (see Section 7.1), ϵ can be as small as 0.05. Therefore, we may reduce the receiver-sender communication of the [28] scheme as much as 30% as long as the noise budget is sufficient.

Similarly, we can apply LSObvCompress to reduce the sender's responses. Because LSObvCompress crucially relies on the fact that the plaintexts of irrelevant indices are 0, we append a dummy point that evaluates to zero to each bin. For the $0.5 \cdot |X|$ dummy bins, the receiver can simply send these zero evaluation points as the query.

We point to Section 7.2 for our experimental evaluation.

H Supplementary Material for Anonymous Messaging (Pung)

Pung, introduced by Angel and Setty [15], is a private messaging protocol that allows parties to exchange messages with each other using a central server, without the server learning anything, including metadata, about the messages. The metadata that is kept private includes which parties are actually engaging in conversation, the number of messages in such conversations, when conversations start, etc. Formally, Pung achieves *relationship unobservability under explicit retrieval*, for which we provide the definition later. Pung also achieves the standard definitions of message integrity and privacy; integrity of ciphertexts under chosen plaintext attacks (INT-CTXT) and indistinguishability under adaptive chosen ciphertext attacks (IND-CCA2), respectively.

At a high-level, Pung is a round-based protocol that proceeds as follows: First, we assume that every pair of users that wish to communicate with each other share a symmetric key ahead of time (or at least know each others' public keys ahead of time, from which such a shared key can be computed; see [15, §6]). This symmetric key is used to derive (i) an encryption key k_E and (ii)

a PRF key k_L used to encrypt and generate pseudo-random labels, respectively, for each message sent between the parties.

In each round r, when a user Alice wants to send a message m to Bob, she simply creates the corresponding label $\ell = \mathsf{PRF}(k_L, (r, \mathsf{Bob}))$ and ciphertext $c = \mathsf{Enc}(k_E, m)$ using her shared keys with Bob and uploads the tuple (ℓ, c) to the server. Every user uploads to the server a fixed number s of label-ciphertext tuples in each round (to prevent traffic analysis attacks). If a user wants to send less than s real messages in a given round, they upload extra random tuples (which are indistinguishable from real label-ciphertext tuples) to make up for the difference. The server collects all such label-ciphertext tuples in round r and creates a PIR database based on them.

When Bob wishes to retrieve Alice's message (and others) in some round $r' \geq r$, he derives the same label ℓ, and combines it with the labels of the other messages he wishes to retrieve in this round to create a batch keyword PIR request query, which he then sends to the server. We assume that in each round r, every user also retrieves a fixed number ρ of messages (again, to prevent traffic analysis attacks). If a user wants to retrieve less than ρ messages in a given round, they simply insert extra dummy queries into the batch to make up for the difference.

Then, the server answers Bob's (and every other users') batch keyword PIR request for this round and sends the corresponding response back to him. From this he computes Alice's ciphertext c and decrypts it using k_E (along with all other real messages he wishes to decrypt).

For security, in every round, the server only receives a fixed number of (pseudo-)random tuples from each user in the send step, and then fixed-length batch keyword PIR requests from each user in the retrieval step. Thus, the information that the server receives from and computation it does for each user in every round hides the users' real intentions. As such, the server cannot tell which users are even sending and retrieving real messages in each round, nor how many real messages, or who is communicating with whom, so security holds (see [15, §A] for a formal proof).

It is easy to see that executing batch keyword PIR results in the primary computational and communication overhead of Pung – all other steps are simple symmetric cryptographic operations. We therefore benchmark our batch keyword PIR scheme using parameters relevant to the Pung setting in Section 7.2 where we show the expected improvements.

H.1 Formal Security

In this section, we present the formal security definition which Pung satisfies, *relationship unobservability under explicit retreival* (UO-ER). We more or less take this definition verbatim from [15, §A] (which itself builds on definitions from [38, 17, 67]).

We start by introducing an abstract protocol π that models communication through a mailbox service (e.g., key-value store). π exposes two functions: $\textsc{Send}(i, r, m)$ and $\textsc{Retrieve}(i, r)$. $\textsc{Send}(i, r, m)$ takes the recipient's id i, a round r and a message m ($\perp$ if the user is idle), and deposits m in a mailbox accessible to i during roudn r. $\textsc{Retrieve}$ retrieves any message sent to i during round r or returns $\perp$ if there is no message. We define UO-ER, based on the following security game.

Security game for UO-ER The game consists of a setup-simulation-guess protocol played by a challenger $\mathcal{C}$ and an adversary. An instance of the game, $G^b_{\mathcal{A},\pi,n,t}(1^\lambda) = b'$, is parameterized by the actions of the adversary, $\mathcal{A}$; the abstract protocol, π; the number of correct users n; the security parameter, λ; the number of rounds for which π runs, t; and the correct output of the game, b. The actual output of the game is the adversary's guess b'. The adversary wins the game if his guess is

correct: $b' = b$.

High-Level Outline. The game is a standard indistinguishability game. In the *setup* phase, $\mathcal{A}$ specifies two scenarios, M^0 and M^1, that describe the behavior of correct users (what messages they send and retrieve). In the *simulation* step $\mathcal{C}$ chooses one scenario at random (M^b) and simulates the actions of correct users under this scenario. $\mathcal{C}$ does so by translating all send and retrieve instructions in M^b to the corresponding SEND and RETRIEVE functions in π, and providing all resulting application-layer packets to $\mathcal{A}$. At the end of the simulation, $\mathcal{A}$ chooses to either move forward onto the *guess* step and issue an answer or it can ask for the setup and simulation steps to be rerun from scratch (i.e., $\mathcal{A}$ specifies two new scenarios M^0 and M^1, and $\mathcal{C}$ simulates the new scenario M^b). If after a polynomial number of iterations $\mathcal{A}$ has not issued a guess, the default value of 0 is assigned to the game's output.

Setup. $\mathcal{A}$ specifies two scenarios M^0 and M^1, each containing a total of $n \cdot t$ entries. Each entry corresponds to a pair of send (S) and retrieve (R) actions performed by a correct user during a single round. The values for these actions are: $S(i, r) - \{j, r_S, m_{i \to j}\}$; and $R(i, r) = \{i, r_R\}$. i, j are the ids of correct users (not necessarily distinct); r_S and r_R are arbitrary rounds; r is the current round (used by $\mathcal{C}$ to know which action to simulate next); $m_{i \to j}$ is the plaintext message sent from i to j, and it could be $\perp$ to indicate that the user does not send a message in round r.

There are two peculiarities with the way $\mathcal{A}$ constructs scenarios. First, M^0 and M^1 describe only the actions of *correct* users. This is because malicious users do not follow π, and cannot be simulated by $\mathcal{C}$. Malicious users' actions play a role during simulation, where $\mathcal{A}$ is invoked as an oracle. Second, correct users send only to other correct users. This is consistent with our goal of relationship unobservability, which provides meaningful privacy only if both the sender and the recipient of a given communication are honest.

Simulation. $\mathcal{A}$ provides the two scenarios that it generates to $\mathcal{C}$. If this is the first iteration, $\mathcal{C}$ flips a random coin and obtains bit b. Otherwise, $\mathcal{C}$ continues to use the previously derived bit b. $\mathcal{C}$ then chooses scenario M^b and follows the protocol in Figure 12. After each call to π's functions, packets going in and out of application layer queues are collected; they are given to $\mathcal{A}$ at the end of the simulation. Note that while the simulation does not explicitly allow $\mathcal{A}$ to drop, reorder, replay, insert, or modify messages from correct users, [at least some of] this power is implicit in the ability of $\mathcal{A}$ to fully specify both scenarios. $\mathcal{A}$ can, however, adpatively perform these actions on behalf of malicious users.

Once the simulation is over, $\mathcal{A}$ can either issue an answer or ask for a rerun with new scenarios (but b is kept unchanged). This allows $\mathcal{A}$ to adapt its strategy ac cross iterations. This process can repeat a number of times that is polynomial in the security parameter λ, after which the game automatically outputs 0 as $\mathcal{A}$'s guess.

Guess. $\mathcal{A}$ outputs a guess b' indicating that scenario $M^{b'}$ was simulated. $\mathcal{A}$ wins the game if it guesses $b' = b$.

Definition 3. *Protocol π provides* UO-ER *if given security parameter λ, for all probabilistic polynomial time algorithms $\mathcal{A}$, for any polynomial number of rounds t and honest users n, there exists a negligible function* negl *such that:*

$$|\Pr[G^0_{\mathcal{A},\pi,n,t}(1^\lambda) = 1] - \Pr[G^1_{\mathcal{A},\pi,n,t}(1^\lambda) = 1]| \leq \mathsf{negl}(1^\lambda).$$

This definition states that if π provides UO-ER, then an adversary gains no meaningful advantage from observing network packets. In other words, the probability of $\mathcal{A}$ distinguishing between

Function SIMULATE($\mathcal{A}, \pi, n, t, M^b$):

packets $\leftarrow \{\}$ // network traces from all users

For $i = 0 \ldots n-1$:

$\pi_i \leftarrow new(\pi)$ // setup an instance of π for user i

For $r = 0 \ldots t-1$

For $i = 0 \ldots n-1$

packets $\overset{\text{insert}}{\longleftarrow} \pi_i$.SEND($M^b.S(i,r)$)

packets $\overset{\text{insert}}{\longleftarrow} \mathcal{A}$ simulates malicious senders

For $i = 0 \ldots n-1$

packets $\overset{\text{insert}}{\longleftarrow} \pi_i$.RETRIEVE($M^b.R(i,r)$)

packets $\overset{\text{insert}}{\longleftarrow} \mathcal{A}$ simulates malicious retrievers

Return packets

Figure 12: Simulation performed by challenger $\mathcal{C}$. $\mathcal{A}$ is the adversary's algorithm; π is an explicit communication protocol; M^b is the scenario to simulate; n is the number of honest users in the scenario; and t is the total number of rounds for which to run π.

Alice communicating with Bob, and Alice communicating with Charlie (or not communicating at all) is negligibly better than a random guess (or any prior $\mathcal{A}$ may have obtained through other channels).

Note that the above game and definition only capture one message per round. This can be addressed by treating multiple entries in a scenario (M) as different send-retrieve paris from the same user during a round. This requires increasing the number of entries per round in the scenarios from n to $k \cdot n$, where k responds to clients' retrieval-rate (the size of their batch keyword PIR queries). $\mathcal{A}$ can thus distinguish between M^0 and M^1 if it can distinguish any one message, which is the expected behavior of extending UO-ER to multiple messages.

Theorem 3 ([15])**.** *Pung's protocol provies UO-ER.*