

Optimal Single-Server Private Information Retrieval

Mingxun Zhou Wei-Kai Lin Yiannis Tselekounis Elaine Shi*

Carnegie Mellon University

Abstract

We construct a single-server pre-processing Private Information Retrieval (PIR) scheme with optimal bandwidth and server computation (up to poly-logarithmic factors), assuming hardness of the Learning With Errors (LWE) problem. Our scheme achieves amortized $\tilde{O}_\lambda(\sqrt{n})$ server and client computation and $\tilde{O}_\lambda(1)$ bandwidth per query, completes in a single roundtrip, and requires $\tilde{O}_\lambda(\sqrt{n})$ client storage. In particular, we achieve a significant reduction in bandwidth over the state-of-the-art scheme by Corrigan-Gibbs, Henzinger, and Kogan (Eurocrypt'22): their scheme requires as much as $\tilde{O}_\lambda(\sqrt{n})$ bandwidth per query, with comparable computational and storage overhead as ours.

1 Introduction

Imagine that a server holds large public database DB indexed by $0, 1, \dots, n-1$, e.g., the repository of DNS entries or a collection of webpages. A client wants to fetch the i -th entry of the database. Although the database is public, the client wants to hide which entry it is interested in. Chor, Goldreich, Kushilevitz, and Sudan [CGKS95, CKGS98] first formulated this problem as Private Information Retrieval (PIR), and since then, a long line of works have focused on constructing efficient PIR schemes [CG97, Cha04, GR05, CMS99, KO97, Lip09, OS07, Gas04, DG16, PR93, DCIO98, BLW17, BGI16, PPY18, IKOS04, Hen16, HH17, ACLS18, IKOS06, LG15, DHS14, CK20, CHK22, KCG21, dCP22].

The good news is that PIR schemes with *poly-logarithmic* bandwidth are well-known [CG97, Cha04, GR05, CMS99, KO97, Lip09, OS07, PR93, BLW17, BGI16, PPY18, DG16, IKOS04, Hen16], either in the single-server or multi-server settings. The bad news is that in the classical PIR setting *without pre-processing*, all known schemes suffer from prohibitive server computation overhead: the server(s) must (in aggregate) perform computation that is linear in the database size n to answer each query. Intuitively, if there is an entry that the server does not look at, it leaks information that the client is not interested in that entry. Beimel, Ishai, and Malkin [BIM00] formalized this intuition into an elegant lower bound, showing that any PIR scheme without pre-processing must incur $\Omega(n)$ server computation per query.

Recognizing this inherent limitation, Beimel et al. [BIM00] introduce a new model for PIR that allows *pre-processing*, and they were the first to show that the linear-computation lower bound can be circumvented with the help of pre-processing. Subsequently, a line of works further explored PIR in the preprocessing model [CK20, CHK22, PY22, SACM21], culminating in the recent works by Corrigan-Gibbs, Henzinger, and Kogan [CHK22] and by Shi et al. [SACM21]. Corrigan-Gibbs, Henzinger, and Kogan [CHK22] proved that in the single-server and pre-processing setting, we can construct a PIR scheme with amortized $\tilde{O}_\lambda(\sqrt{n})$ server and client computation per query,

*Author ordering is randomized.

Table 1: **Comparison of single-server PIR schemes.** Q is the batch size for batch PIR, m is the number of clients, n is the database size, and $\epsilon \in (0, 1)$ is some suitable constant. “BW” means bandwidth per query. “CRA” means the composite residuosity assumption, ϕ -hiding is a number-theoretic assumption described in [CMS99], “OLDC” means oblivious locally decodable codes, and “VBB” means virtual-blackbox obfuscation.

Scheme	Assumpt.	Adaptive	BW	Per-query time		Extra space	
				Client	Server	Client	Server
Standard [Cha04, CMS99, GR05]	CRA or ϕ -hiding or LWE	✓	$\tilde{O}(1)$	$\tilde{O}(1)$	$O(n)$	0	0
Batch PIR [ACLS18, IKOS04]	same as above	✗	$\tilde{O}(1)$	$\tilde{O}(1)$	$O(\frac{n}{Q})$	0	0
[CHR17, BIPW17]	OLDC	✓	n^ϵ	n^ϵ	n^ϵ	$O(1)$	mn
[BIPW17]	OLDC, VBB	✓	n^ϵ	n^ϵ	n^ϵ	0	n
[CK20]	LWE	✓	$\tilde{O}_\lambda(\sqrt{n})$	$\tilde{O}_\lambda(\sqrt{n})$	$\tilde{O}_\lambda(n)$	$\tilde{O}_\lambda(\sqrt{n})$	0
[CHK22]	LWE	✓	$\tilde{O}_\lambda(\sqrt{n})$	$\tilde{O}_\lambda(\sqrt{n})$	$\tilde{O}_\lambda(\sqrt{n})$	$\tilde{O}_\lambda(\sqrt{n})$	0
Ours	LWE	✓	$\tilde{O}_\lambda(1)$	$\tilde{O}_\lambda(\sqrt{n})$	$\tilde{O}_\lambda(\sqrt{n})$	$\tilde{O}_\lambda(\sqrt{n})$	0

while requiring $\tilde{O}_\lambda(\sqrt{n})$ client storage. Here, we use $\tilde{O}_\lambda(\cdot)$ to hide $\text{poly}(\lambda, \log n)$ factors where λ is the security parameter. Corrigan-Gibbs et al. [CHK22] also showed that their scheme achieves optimality up to polylog factors in terms of server computation, assuming $\tilde{O}(\sqrt{n})$ client storage. Unfortunately, their scheme suffers from $\tilde{O}_\lambda(\sqrt{n})$ bandwidth overhead which is significantly worse than classical PIR schemes without pre-processing. On the other hand, Shi et al. [SACM21] showed that in a setting with two non-colluding servers, we can construct a PIR scheme that incurs only $\tilde{O}_\lambda(1)$ online bandwidth and $\tilde{O}_\lambda(\sqrt{n})$ server and client computation per query, while requiring $\tilde{O}_\lambda(\sqrt{n})$ client storage. Both of these schemes support *unbounded* number of queries after a one-time pre-processing, and in the above, the cost of the pre-processing is amortized to each query.

While the two schemes [SACM21, CHK22] achieve similar server and client computation overhead, Shi et al. [SACM21] has the advantage that it achieves $\tilde{O}_\lambda(1)$ online bandwidth — although unfortunately, this is achieved at the price of requiring two non-colluding servers. Notably, Shi et al.’s scheme is known to be optimal up to polylog factors even in the two-server setting, in terms of bandwidth and server computation, assuming that the client can only download roughly \sqrt{n} amount of data during the offline pre-processing phase [CK20].

Given the state of the art, we ask whether we can achieve the best of both worlds. Specifically, we ask the following natural question — the same open question was also raised by Corrigan-Gibbs et al. in their recent work [CHK22]:

Can we construct a *single-server* pre-processing PIR scheme that achieves (near) *optimality* in both *server computation* and *bandwidth*?

1.1 Our Contributions

We provide an affirmative answer to the aforementioned question by proving the following theorem:

Theorem 1.1. *Assume that the Learning With Errors (LWE) assumption holds. Then, there exists*

a single-server pre-processing PIR scheme that achieves amortized $\tilde{O}_\lambda(1)$ bandwidth, and $\tilde{O}_\lambda(\sqrt{n})$ server and client computation per query, and requires $\tilde{O}_\lambda(\sqrt{n})$ client storage.

More specifically, in our scheme, there is a one-time pre-processing phase with the same overheads in all dimensions as Corrigan-Gibbs [CHK22] (up to polylog factors). During the offline pre-processing, the client and the server engage in $\tilde{O}_\lambda(\sqrt{n})$ communication, the server performs $\tilde{O}_\lambda(n)$ computation, and the client performs $\tilde{O}_\lambda(\sqrt{n})$ computation. In Theorem 1.1 above, the cost of the pre-processing is *amortized* to the subsequent queries. After the one-time pre-processing, we can support an unbounded number of queries, and for each query, we incur the same costs as stated in Theorem 1.1 in the *worst case*. Our actual construction makes use of two cryptographic primitives: fully homomorphic encryption (FHE) [Gen09, GSW13] and privately programmable pseudorandom functions [BLW17, PS18, KW21], both of which have known instantiations assuming LWE.

Near optimality. Our scheme is optimal up to polylog factors in terms of server computation and bandwidth, in light of the lower bounds proven in recent works [CK20, CHK22]. Specifically, Corrigan-Gibbs and Kogan [CK20] showed that for any pre-processing PIR scheme where the server stores only the original database, it must be that $C \cdot T \geq \Omega(n)$ where C is the bandwidth incurred during the offline pre-processing and T is the online server time per query. The recent work of Corrigan-Gibbs, Henzinger, and Kogan [CHK22] proved that for any pre-processing PIR scheme that supports unbounded number of dynamic queries and assuming the server stores only the original database, it must be that $S \cdot T \geq \Omega(n)$ where S is client’s storage and T is the online server time per query.

Although in the main body, we focus on the most interesting special case where the parameters S and T are balanced, in Appendix A, we discuss how to achieve a smooth tradeoff between S and T . In particular, for any function $f(n) \in [\log^c n, n/\log^c n]$ for some suitable positive constant c , we give a scheme that requires only $\tilde{O}_\lambda(f(n))$ client space, and achieves $\tilde{O}_\lambda(n/f(n))$ online server and client time per query, and $\tilde{O}_\lambda(1)$ bandwidth per query. Therefore, we achieve near optimality for every choice of client space.

Comparison with prior schemes. Table 1 compares our scheme against various prior works. In this table, we focus on schemes in the *single-server* setting, and for pre-processing PIR schemes, we amortize the pre-processing overhead over an unbounded number of subsequent queries. Among these schemes, batch PIR schemes [ACLS18, IKOS04, Hen16] must have a large batch size of Q to achieve the stated amortized performance, and fail in the scenario when the queries are generated *adaptively* and arrive one by one. We discuss additional related work in Section 1.2.

1.2 Additional Related Work

We now review some additional related work. Besides being first to define PIR with pre-processing, Beimel et al. [BIM00] additionally showed how to construct a preprocessing PIR with polylogarithmic online bandwidth assuming polylogarithmically many non-colluding servers, and $\text{poly}(n)$ server storage. Unlike our work as well as the recent works by Corrigan-Gibbs et al. [CHK22, CK20], the scheme by Beimel et al. [BIM00] employs a *public* pre-processing, where the pre-processing results in no client-side secret state. In fact, in their scheme [BIM00], the server pre-processes the database, resulting in a $\text{poly}(n)$ -sized encoding of the database which is then stored by the server. The very recent work of Persiano and Yeo [PY22] proved that for any PIR scheme with *public* pre-processing, it must be that $T \cdot R \geq \Omega(n \log n)$ where T is the server computation per query

and R is size of the additional state computed by the public pre-processing. In comparison, our work considers a *private* pre-processing model, i.e., the end of the pre-processing, the client stores some secret state not seen by the server. This model matches well with a “subscription model” in practice. For example, every client that needs private DNS service can subscribe with the provider, and during subscription, they perform the one-time pre-processing.

Besides the single-server PIR scheme from FHE mentioned in Table 1, the work of Corrigan-Gibbs and Kogan [CK20] also propose another scheme assuming only linearly homomorphic encryption, which requires $O(n^{2/3})$ bandwidth and client computation and $O(n)$ server computation per query, as well as $O(n^{2/3})$ client storage. Further, the work of Corrigan-Gibbs, Henzinger, and Kogan [CHK22] additionally suggests a single-server PIR scheme assuming only linearly homomorphic encryption, incurring $O(\sqrt{n})$ bandwidth and client computation, and $O(n^{3/4})$ server computation per query, requiring $O(n^{3/4})$ client storage.

Hamlin et al. [HOWW19] suggested a related notion called *private anonymous data access* (PANDA). PANDA is a form of preprocessing PIR which requires an additional *third-party trusted setup* besides the client and the servers; and moreover, the server storage and time grow w.r.t. the number of corrupt clients. In applications (e.g., private DNS) that involve a potentially unbounded number of mutually distrustful clients, PANDA schemes would be unsuitable.

A line of works have explored the concrete efficiency of PIR schemes [ACLS18, MCR21, KCG21, PPY18, GI14, ?]. In particular, the work of Angel et al. [ACLS18] relies on batching to amortize the linear server computation over a batch of queries. Kogan and Corrigan-Gibbs [KCG21] gives a practical instantiation of the two-server pre-processing PIR scheme described in their earlier work [CK20], with a new trick that removes the k -fold parallel repetition. For their private blacklist application, it turns out that the database is somewhat small, and therefore, they are willing to incur $\Theta(n)$ client-side computation per online query, in exchange for logarithmic bandwidth. The work of Patel et al. [PPY18] explores how to rely on a stateful client to improve the concrete performance of PIR schemes. Our work focuses on the asymptotical overhead, and we leave it to future work to consider concretely efficient instantiations that preserve our asymptotical performance.

Some works have considered achieving sublinear server time by relaxing the security definition. Toledo et al. [TDG16] suggested to relax the security definition to differential privacy, to improve the server time to sublinear, assuming that a large number of servers are available. Al-bab et al. [AIVG20] also considered the differential privacy notion, and they can achieve sublinear amortized server computation in a batched setting.

2 Technical Roadmap

2.1 Starting Point: Optimal 2-Server Scheme By Shi et al.

2.1.1 An Inefficient Toy Scheme

Our starting point is the nearly optimal 2-server scheme by Shi et al. [SACM21], and we will explore how to coalesce the two servers into one. To understand their scheme, it helps to start out with the following toy scheme which is a slight variant of the strawman schemes described in recent works [CK20, SACM21]. Henceforth, we use the notations **Right** and **Left** to denote two non-colluding servers. Let \mathcal{D}_n be some distribution from which we can sample random sets of expected size \sqrt{n} — at this moment, the reader need not care what exactly the distribution \mathcal{D}_n is.

Inefficient Toy 2-Server Scheme: Single-Copy Version

Offline preprocessing.

(DB[k] denotes the k-th bit of the database)

- Client samples \sqrt{n} sets $S_1, S_2, \dots, S_{\sqrt{n}} \subseteq \{0, 1, \dots, n-1\}$ from the distribution \mathcal{D}_n .
- Client sends the resulting sets $S_1, \dots, S_{\sqrt{n}}$ to Left. For each set $j \in [\sqrt{n}]$, Left responds with the parity bit $p_j := \bigoplus_{k \in S_j} \text{DB}[k]$ of indices in the set.
- Client stores the hint table $T := \{T_j := (S_j, p_j)\}_{j \in [\sqrt{n}]}$.

Online query for index $x \in \{0, 1, \dots, n-1\}$.• **Query:** (Client \Leftrightarrow Right)

1. Find an entry $T_j := (S_j, p_j)$ in its hint table T such that $x \in S_j$. Let $S^* := S_j$ if found, else let S^* be a fresh random set containing x .
2. Send the set $S := \text{ReSamp}(S^*, x)$ to Right, where $\text{ReSamp}(S^*, x)$ outputs a set almost identical to S^* , except that the coins used to determine x 's membership are re-tossed.
3. Upon obtaining a response $p := \bigoplus_{k \in S} \text{DB}[k]$ from Right, output the candidate answer $\beta' := p_j \oplus p$ or $\beta' := 0$ if no such T_j was found earlier.
4. Client obtains the true answer $\beta := \text{DB}[x]$ — the full scheme will repeat this single-copy scheme $k = \omega(\log \lambda)$ times, and β is computed as a majority vote among the k candidate answers, which is guaranteed to be correct except with negligible probability.

• **Refresh** (Client \Leftrightarrow Left)

1. Client samples a random set S' , and sends S' to Left.
2. Left responds with $p' := \bigoplus_{k \in S'} \text{DB}[k]$. Let $\tilde{p} = p' \oplus \beta$ if $x \notin S'$, else let $\tilde{p} = p'$. If a table entry T_j containing x was found and consumed earlier, Client replaces T_j with $(S' \cup \{x\}, \tilde{p})$.

In this 2-server toy scheme, during the offline phase, the client samples \sqrt{n} sets each of expected size \sqrt{n} from some distribution \mathcal{D}_n . It downloads the parities of all these sets from the Left server. It stores all these sets as well as the parity of each set in a local hint table. During the online phase, to query an index $x \in \{0, 1, \dots, n-1\}$, the client looks up its hint table and finds a set S^* that contains x , whose parity is p_j . It then resamples the coins that determine whether x is in the set or not. It sends the resampled set to the Right server, which returns the client the parity p' . The client computes $\beta' = p' \oplus p_j$ as the candidate answer. If we choose the distribution \mathcal{D}_n carefully, then, with significant probability, the $\text{ReSamp}(x)$ will *remove the element x from the set, without adding or removing any other element*. In this case, the candidate answer β' would be correct. If we can ensure that each single copy has $2/3$ correctness probability, then we can amplify the correctness probability to $1 - \text{negl}(\lambda)$ through parallel repetition using $\omega(\log \lambda)$ copies and majority voting. Finally, once we consume a hint from the table, we need to replenish it. To achieve this, the client samples a random set S' , and obtains its parity p' from the Left server. The client replaces the consumed entry with the set $S' \cup \{x\}$ and its parity which can be computed knowing p' and $\beta = \text{DB}[x]$.

Privacy. Privacy w.r.t. the Left server is easy to see. Basically, the Left server sees \sqrt{n} random sets sampled from \mathcal{D}_n during the offline phase, and during each online query, it sees an additional random set also sampled from \mathcal{D}_n . Privacy w.r.t. the Right server can be proven using an inductive argument. Initially, the client's hint table consists of \sqrt{n} random sets sampled independently from

\mathcal{D}_n . Suppose that at the end of the i -th, the client’s hint table satisfies the above distribution. Then, during the i -th query that requests some index $x \in \{0, 1, \dots, n-1\}$, if some hint (S_j, p_j) is matched, i.e., $S_j \ni x$, then, the distribution of S_j is the same as sampling from \mathcal{D}_n subject to containing x . Therefore, the set sent to the **Right** server, i.e., **ReSamp** (S_j) has the same distribution as sampling at random from \mathcal{D}_n . Further, notice that the client replaces the consumed entry with another set sampled at random subject to containing x . Thus, at the end of the i -th query, the client’s hint table still has \sqrt{n} i.i.d. sets sampled from \mathcal{D}_n .

Inefficiency of the toy scheme. In the toy scheme, both the server and the client perform roughly \sqrt{n} computation per query. However, the online bandwidth to each of the two servers is roughly \sqrt{n} , and the client storage is $O(n)$.

2.1.2 Compressing the Bandwidth and Client Storage

Pseudorandom sets with private ReSamp. Shi et al. [SACM21] suggested an idea to improve the efficiency of the toy scheme in the two-server setting. To achieve this, they introduce a cryptographic object called a pseudorandom set (PRSet), allowing us to succinctly represent a pseudorandom set of size roughly \sqrt{n} with a short key of $\text{poly}(\lambda)$ bits. In this way, the client can store a key in place of each set, and send a key to the server in place of the full description of a set. Their PRSet scheme must support the following operations:

- $\text{sk} \leftarrow \mathbf{Gen}(1^\lambda, n)$: samples a key sk that generates a pseudorandom set emulating the distribution \mathcal{D}_n ;
- $S \leftarrow \mathbf{Set}(\text{sk})$: given a key sk , enumerate the set S ;
- $\text{sk} \leftarrow \mathbf{Member}(\text{sk}, x)$: test if an element $x \in \{0, 1, \dots, n-1\}$ is contained in $\mathbf{Set}(\text{sk})$;
- $\text{sk}' \leftarrow \mathbf{ReSamp}(\text{sk}, x)$: given a key sk , generates a related key sk' that effectively resamples the coins that are used to determine whether x is in the set or not, while preserving all other coins¹;

Designing such a PRSet scheme turns out to be non-trivial, since we need to satisfy the following properties simultaneously.

- *Privacy of ReSamp.* The resampled key output by $\mathbf{ReSamp}(\text{sk}, x)$ must hide the point x that is being resampled.
- *Efficient membership test and set enumeration.* The membership test algorithm $\mathbf{Member}(\text{sk}, x)$ must complete in $\tilde{O}_\lambda(1)$ running time and the set enumeration algorithm $\mathbf{Set}(\text{sk})$ must complete in $\tilde{O}_\lambda(\sqrt{n})$ time.

Shi et al. [SACM21] show how to rely on a privately puncturable pseudorandom function [BKM17, CC17, BTWV17] to construct a PRSet scheme that supports a private **ReSamp** operation. Further, to satisfy efficient membership test and efficient set enumeration simultaneously, they carefully crafted a distribution \mathcal{D}_n that the PRSet scheme emulates. Notably, whether two elements are in the set may not be independent in the distribution \mathcal{D}_n . Such weak dependence between elements brings additional possibilities of errors. In particular, $\mathbf{ReSamp}(\text{sk}, x)$ may accidentally remove other elements besides x . If $\mathbf{ReSamp}(\text{sk}, x)$ either fails to remove x or ends up removing additional elements besides x , the resulting PIR scheme would be incorrect. Shi et al. [SACM21] made sure that the probability of such error is small, such that each single copy of the PIR scheme still has 2/3 correctness.

¹Shi et al. [SACM21] referred to **ReSamp** as **Punct** since the operation is implemented by calling the puncturing operation of the underlying privately puncturable PRF.

Optimal 2-server PIR scheme. With such a PRSet scheme, we can easily modify the aforementioned toy scheme to compress the client storage and bandwidth [SACM21]. Specifically, during the offline phase, the client sends \sqrt{n} PRSet keys to the Left server. The Left server uses the set enumeration algorithm **Set** to enumerate the sets and sends the client their parity bits. The client now stores a hint table where each entry is of the form (sk_i, p_i) where sk_i is a PRSet key that can be used to generate a set of size roughly \sqrt{n} , and p_i is the parity bit as before. During an online query for $x \in \{0, 1, \dots, n-1\}$, the client finds a sk^* in its hint table such that $\text{Member}(\text{sk}^*, x) = 1$, and it sends the outcome of $\text{ReSamp}(\text{sk}^*, x)$ to the Right server. If such a key is not found, the client simply samples a random $\text{sk}' \leftarrow \text{Gen}(1^\lambda, n)$ and sends it to the server. The client computes the candidate answer the same way as before. What is most interesting is how to perform the refresh operation to replenish the consumed key. This is achieved in the following manner:

- Sample $\text{sk}' \leftarrow \text{Gen}(1^\lambda, n)$ subject to $\text{Member}(\text{sk}', x) = 1$, and send the outcome of $\text{ReSamp}(\text{sk}', x)$ to the Left server.
- The Left server enumerates the set using the **Set** algorithm and sends the client the parity bit p' . The client replaces the consumed entry with $(\text{sk}', p' \oplus \beta)$ where $\beta = \text{DB}[x]$ is the true answer to the current query.

2.2 Highlights of Our Construction and Proof Techniques

Corrigan-Gibbs and Kogan [CK20] proposed a technique to transform a two-server pre-processing PIR scheme into a single-server scheme, and the technique was further extended by Corrigan-Gibbs, Henzinger, and Kogan [CHK22]. Unfortunately, we cannot directly apply their technique to transform the scheme by Shi et al. as explained below.

The idea in earlier work [CK20, CHK22] is to get rid of the Left server and redirect the queries originally destined for the Left server instead to the Right server, but now encrypted under a fully homomorphic encryption (FHE) scheme. The Right server now evaluates the answers to the query through homomorphic evaluation. We are, however, faced with one important technicality. An FHE scheme supports homomorphic evaluation only of circuits and not of RAM programs. Given an FHE-encrypted PRSet key sk , the server wants to compute the encrypted parity of the set generated by sk — to implement this computation as a circuit, we would need a circuit of size at least n . This would make the server computation linear again. Fortunately, the following critical observation, first made by Corrigan-Gibbs et al. [CHK22], saves the day.

Observation. Although homomorphically evaluating the parity of a single set takes a linear-sized circuit, we can batch-evaluate the parity bits of $\Theta(\sqrt{n})$ sets in a circuit of size $\tilde{O}(n)$, leveraging oblivious sort. With batch evaluation, the amortized cost per set is only $\tilde{O}(\sqrt{n})$.

Idea 1: Batched refresh operations. Our first idea is inspired by the work of Corrigan-Gibbs et al. [CHK22]. To leverage this batch homomorphic evaluation idea, we must get rid of the online refresh phase which refreshes the consumed hints one by one. Instead, we first consider a Q -bounded scheme that supports only $Q = \sqrt{n}$ queries — in this way, we can hope to front-load all Q refresh operations upfront during the pre-processing phase. Given a Q -bounded scheme, it is easy to obtain a scheme supporting *unbounded* number of queries. We can simply rerun the offline setup every Q queries, and amortize the cost of the periodic setup over each query — in fact, it is also not hard to deamortize the periodic setup and spread the work across time.

In summary, through batching the refresh operations, we can hope to achieve $\tilde{O}_\lambda(\sqrt{n})$ amortized server computation per refresh operation.

Idea 2: a pseudorandom set scheme supporting Add and ReSamp. If we front-load all Q refresh operations upfront during the offline pre-processing, one technicality arises. Recall that during a query for $x \in \{0, 1, \dots, n-1\}$, we must replenish the consumed entry with a set sampled subject to containing the queried element x . During the offline pre-processing, however, we do not have foreknowledge of x . Therefore, we can only hope to sample (pseudo-)random sets (represented by keys) during the offline pre-processing, and add the element x to the set during the online phase.

This means that we need a new PRSet that supports not only **ReSamp**, but also an **Add** operation. Specifically, given a PRSet key sk , the client should be able to call $\text{sk}' \leftarrow \mathbf{Add}(\text{sk}, x)$ and then call $\text{rsk} \leftarrow \mathbf{ReSamp}(\text{sk}', y)$, and send the resulting rsk to the server. Further, for privacy, the resulting rsk must hide both x and y . To construct such a PRSet scheme, we need a cryptographic primitive called privately programmable pseudorandom functions [BLW17, PS18, KW21], which is stronger than the privately puncturable pseudorandom functions employed by Shi et al.

New proof techniques. For the optimal two-server scheme of Shi et al. [SACM21], they have a relatively simple privacy proof, but the correctness proof is rather involved due to technical reasons related to analyzing the underlying randomized process. Interestingly, for our optimal single-server scheme, it is quite the opposite: the privacy proof turns out to be very tricky, and the correctness proof is simpler. At a high level, the challenges in the privacy proof arise due to the way the probability analysis is intertwined with the cryptography. We devise new proof techniques that allow us to decompose the computational reductions and the analysis of the underlying randomized process.

To help the reader understand the technicalities of our privacy proof and our new ideas, we give an informal proof roadmap in Section 6.1.

3 Preliminaries

3.1 Privately Programmable Pseudorandom Functions

Intuitively, a privately programmable pseudorandom function [BLW17, PS18, KW21] is a pseudorandom function (PRF) with one extra capability: it allows one to create a *programmed key* that forces the PRF's outcomes at most L distinct input points $\{x_i\}$ to be a set of pre-determined values $\{v_i\}$. For security, we want to guarantee the privacy of the programmed inputs. Specifically, if the set of output values $\{v_i\}$ are randomly chosen, then the programmed key should not leak more information about the set of input points programmed. Further, the programmed key should not leak the original PRF's evaluation outcomes at the programmed inputs prior to the programming.

3.1.1 Syntax

Let \mathcal{X} denote the input domain and let \mathcal{V} denote the output range, whose sizes may depend on the security parameter λ . A programmable pseudorandom function is a tuple (**Gen**, **Eval**, **Prog**, **PEval**) of efficient, possibly randomized algorithms with the following syntax:

- **Gen**($1^\lambda, L$): given the security parameter λ and an upper bound, L , on the number of programmable inputs, output a master secret key msk .
- **Eval**(msk, x): given the master secret key msk and an input $x \in \mathcal{X}$, output the evaluation result $v \in \mathcal{V}$ on the input x .
- **Prog**($\text{msk}, P = \{(x_i, v_i)\}$): given the master secret key msk and a set P containing up to L pairs $(x_i, v_i) \in \mathcal{X} \times \mathcal{V}$, where all x_i 's must be distinct, output a programmed key sk_P .

<u>RealPPRF$_{\mathcal{A}}(1^\lambda, L)$:</u>	<u>IdealPPRF$_{\mathcal{A}, \text{Sim}}(1^\lambda, L)$:</u>
$P := \{(x_i, v_i)\}_{i \in [L']} \leftarrow \mathcal{A}(1^\lambda, L)$ <i>// require: $L' \leq L$</i>	$P := \{(x_i, v_i)\}_{i \in [L']} \leftarrow \mathcal{A}(1^\lambda, L)$ <i>// require: $L' \leq L$</i>
$\text{msk} \leftarrow \mathbf{Gen}(1^\lambda, L)$	$\text{sk}_P \leftarrow \mathbf{Sim}(1^\lambda, P, L)$
$\text{sk}_P \leftarrow \mathbf{Prog}(\text{msk}, P)$	$\text{sk}_P \rightarrow \mathcal{A}$
$\text{sk}_P \rightarrow \mathcal{A}$	repeat
repeat	$x \leftarrow \mathcal{A}$
$x \leftarrow \mathcal{A}$	If $x \notin \{x_i\}_{i \in [L']}$ then $\mathbf{PEval}(\text{sk}_P, x) \rightarrow \mathcal{A}$
$\mathbf{Eval}(\text{msk}, x) \rightarrow \mathcal{A}$	Else $v \xleftarrow{\$} \mathcal{V}, v \rightarrow \mathcal{A}$
until \mathcal{A} halts	until \mathcal{A} halts

Figure 1: The real and ideal experiments for simulation security.

- $\mathbf{PEval}(\text{sk}_P, x)$: given a programmed key sk_P and an input $x \in \mathcal{X}$, output the evaluation outcome, $v \in \mathcal{V}$, over the input x .

Correctness of programming. A programmable function satisfies correctness if for all $\lambda, L = \text{poly}(\lambda) \in \mathbb{N}$, all sets of up to L pairs $P := \{(x_i, v_i)\} \subseteq \mathcal{X} \times \mathcal{V}$ (with distinct x_i s), we have the following:

1. For every $i \in [|P|]$,

$$\Pr \left[\mathbf{PEval}(\text{sk}_P, x_i) \neq v_i \mid \begin{array}{l} \text{msk} \leftarrow \mathbf{Gen}(1^\lambda, L) \\ \text{sk}_P \leftarrow \mathbf{Prog}(\text{msk}, P) \end{array} \right] \leq \text{negl}(\lambda), \text{ and}$$

2. For any x' not in P , we have

$$\Pr \left[\mathbf{PEval}(\text{sk}_P, x') \neq \mathbf{Eval}(\text{msk}, x') \mid \begin{array}{l} \text{msk} \leftarrow \mathbf{Gen}(1^\lambda, L) \\ \text{sk}_P \leftarrow \mathbf{Prog}(\text{msk}, P) \end{array} \right] \leq \text{negl}(\lambda).$$

We note that Peikert and Shiehian [PS18] did not define the second correctness condition above, but their proof shows that the second condition also holds.

3.1.2 Security Definitions

Definition 3.1 (Simulation security). A programmable function is *simulation secure*, if there is a probabilistic polynomial-time (PPT) simulator Sim such that for any PPT adversary \mathcal{A} and any polynomial $L(\lambda)$,

$$\left\{ \text{RealPPRF}_{\mathcal{A}}(1^\lambda, L) \right\}_{\lambda \in \mathbb{N}} \stackrel{c}{\approx} \left\{ \text{IdealPPRF}_{\mathcal{A}, \text{Sim}}(1^\lambda, L) \right\}_{\lambda \in \mathbb{N}},$$

where RealPPRF and IdealPPRF are the respective views of \mathcal{A} in the executions of Figure 1 and “ $\stackrel{c}{\approx}$ ” denotes computational indistinguishability.

<u>RealPPRFPriv$\mathcal{A}(1^\lambda, L)$:</u>	<u>IdealPPRFPriv$\mathcal{A}, \text{Sim}(1^\lambda, L)$:</u>
$\{x_i\}_{i \in [L']} \leftarrow \mathcal{A}(1^\lambda, L)$ // require: $L' \leq L$	$\{x_i\}_{i \in [L']} \leftarrow \mathcal{A}(1^\lambda, L)$ // require: $L' \leq L$
$\{v_i\}_{i \in [L']} \xleftarrow{\$} \mathcal{V}$	$\text{sk} \leftarrow \text{Sim}(1^\lambda, L)$
$P := \{(x_i, v_i)\}_{i \in [L']}$	$\text{sk} \rightarrow \mathcal{A}$
$\text{msk} \leftarrow \mathbf{Gen}(1^\lambda, L), \text{sk} \leftarrow \mathbf{Prog}(\text{msk}, P)$	
$\text{sk} \rightarrow \mathcal{A}$	

Figure 2: The real and ideal experiments for private programmability.

Definition 3.2 (Private programmability). A programmable function is *privately programmable*, if there is a PPT simulator Sim such that for any PPT adversary \mathcal{A} and any polynomial $L(\lambda)$,

$$\left\{ \text{RealPPRFPriv}_{\mathcal{A}}(1^\lambda, L) \right\}_{\lambda \in \mathbb{N}} \stackrel{c}{\approx} \left\{ \text{IdealPPRFPriv}_{\mathcal{A}, \text{Sim}}(1^\lambda, L) \right\}_{\lambda \in \mathbb{N}},$$

where RealPPRFPriv and IdealPPRFPriv are the respective views of \mathcal{A} in the executions of Figure 2.

Last but not the least, we define an additional security property, i.e., the ordinary pseudorandomness notion for the PRF. We prove that pseudorandomness is implied by private programmability — however, defining this notion explicitly will facilitate our proofs later.

Definition 3.3 (Pseudorandomness). We say that a programmable pseudorandom function satisfies pseudorandomness iff for every probabilistic polynomial-time adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that the following holds:

$$\left| \Pr[\text{msk} \leftarrow \mathbf{Gen}(1^\lambda, L) : \mathcal{A}^{\text{Eval}(\text{msk}, \cdot)} = 1] - \Pr[\text{rf} \xleftarrow{\$} \mathcal{RF} : \mathcal{A}^{\text{rf}(\cdot)} = 1] \right| \leq \text{negl}(\lambda)$$

where \mathcal{RF} denotes the family of random functions that map the input domain \mathcal{X} to the output range \mathcal{V} .

Fact 3.4. *Suppose that a programmable PRF scheme satisfies private programmability, then it also satisfies pseudorandomness.*

Proof. Let q be the maximum number of queries made by the pseudorandomness adversary \mathcal{A} . We consider a sequence of hybrids H_0, H_1, \dots, H_q . In H_j where $j \in \{0, 1, \dots, q\}$, for the first j *distinct* queries made by \mathcal{A} , return to \mathcal{A} truly random answers, and for the remaining queries, return the outcomes of the PRF evaluation. Note that if \mathcal{A} makes any repeat query, it always gets the same answer as before.

It suffices to show that no probabilistic polynomial-time \mathcal{A} can distinguish H_i and H_{i+1} for any $i \in \{0, 1, \dots, q-1\}$. To show this, consider an intermediate hybrid H'_i . In H'_i , the first i distinct queries are answered with true randomness, and the remaining queries are answered using a simulated key generated by $\text{sk} \leftarrow \text{Sim}(1^\lambda, L)$.

We first show that H_{i+1} is computationally indistinguishable from H'_i . Suppose that there is an efficient adversary \mathcal{A} that can distinguish H_i and H_{i+1} . We can construct an efficient reduction \mathcal{B} that breaks the private programmability of the underlying PRF. \mathcal{B} answers the first $i-1$ distinct queries from \mathcal{A} using true randomness. When \mathcal{A} submits the i -th distinct query x_i , and submits $\{x_i\}$ to its own challenger. It gets back from its challenger sk . For all remaining queries including

the i -th query, it uses $\mathbf{PEval}(\mathbf{sk}, \cdot)$ to answer to \mathcal{A} . If \mathcal{B} is playing $\mathbf{RealPPRFPriv}$, then \mathcal{A} 's view is identically distributed as H_{i+1} , else if \mathcal{B} is playing $\mathbf{IdealPPRFPriv}$, then \mathcal{A} 's view is identically distributed as H'_i .

Next, we show that H'_i is computationally indistinguishable from H_i . Consider H''_i in which all but the first i queries are answered using a key \mathbf{sk} generated as follows: $\mathbf{msk} \leftarrow \mathbf{Gen}(1^\lambda, L)$, $\mathbf{sk} \leftarrow \mathbf{Prog}(\mathbf{msk}, \emptyset)$. H_i is statistically indistinguishable from H''_i due to the correctness of the programmable PRF. H''_i is computationally indistinguishable from H'_i through a straightforward reduction to the private programmability of the PRF.

Summarizing the above, H_i is computationally indistinguishable from H_{i+1} and this suffices for proving the claim. \square

3.1.3 Construction

In our syntax and security definitions above, we want the programmable PRF to support programming *at most* L inputs. By contrast, Peikert and Shiehian [PS18] gave a construction of privately programmable PRFs where the \mathbf{Prog} function must program *exactly* L inputs. Similarly, in their security definitions, the admissible adversary \mathcal{A} is required to satisfy $L' = L$ (as opposed to $L' \leq L$ in our case).

Given a privately programmable PRF construction that programs exactly L inputs, we now show how to construct a new scheme that allows programming *up to* L inputs. In our PIR construction later, we want the PRF's input domain to contain all strings of length up to some parameter $\ell \in \mathbb{N}$. We use the notation $\{0, 1\}^{\leq \ell}$ to denote all strings of length up to ℓ .

Let $\mathbf{PRF}' := (\mathbf{Gen}', \mathbf{Eval}', \mathbf{Prog}', \mathbf{PEval}')$ denote a privately programmable PRF whose input domain is $\mathcal{X}' = \{0, 1\}^{\leq \ell+1}$, i.e., all strings of length up to $\ell + 1$, and whose output range is \mathcal{V} , supporting programming exactly L inputs. We now construct a privately programmable PRF scheme denoted \mathbf{PRF} whose input domain is $\mathcal{X} = \{0, 1\}^{\leq \ell}$, i.e., all strings of length up to ℓ , and whose output range is \mathcal{V} , i.e., the same as that of \mathbf{PRF}' .

- $\mathbf{Gen}(1^\lambda, L)$: let $\mathbf{msk} \leftarrow \mathbf{Gen}'(1^\lambda, L)$, and output \mathbf{msk} ;
- $\mathbf{Eval}(\mathbf{msk}, x)$: output $\mathbf{Eval}'(\mathbf{msk}, x||0)$;
- $\mathbf{Prog}(\mathbf{msk}, P = \{(x_i, v_i)\}_{i \in [L']})$:
 - choose $L - L'$ distinct strings of length at most $\ell + 1$ that ends with 1, denoted $x'_1, \dots, x'_{L-L'}$;
 - for $j \in [L - L']$, choose $v_j \xleftarrow{\$} \mathcal{V}$ at random;
 - call $\mathbf{sk} \leftarrow \mathbf{Prog}'(\mathbf{msk}, \{(x_i||0, v_i)\}_{i \in [L']} \cup \{(x'_j, v_j)\}_{j \in [L-L']})$, and output \mathbf{sk} .
- $\mathbf{PEval}(\mathbf{sk}, x)$: let $v \leftarrow \mathbf{PEval}'(\mathbf{sk}, x||0)$ and output v .

Claim 3.5. *Suppose that the underlying programmable PRF' that maps $\{0, 1\}^{\ell+1}$ to \mathcal{V} satisfies correctness, simulation security, and private programmability. Then, the above PRF which maps $\{0, 1\}^\ell$ to \mathcal{V} also satisfies correctness, simulation security, and private programmability.*

Proof. We prove the properties one by one.

Correctness. If x is one of the programmed inputs, then correctness of \mathbf{PEval} over x follows directly from the correctness of \mathbf{PEval} of the underlying \mathbf{PRF}' . Now consider the case when x is not one of the programmed inputs. Observe that \mathbf{Prog} algorithm may program the underlying \mathbf{PRF}' at $L' - L$ strings of length at most $\ell + 1$ ending with 1. However, calling $\mathbf{PRF.PEval}(\mathbf{sk}, x)$ results

in calling $\text{PRF}'.\mathbf{PEval}(\text{sk}, x||0)$, and $x||0$ cannot be a programmed input for the underlying PRF' . Therefore, correctness of \mathbf{PEval} over a non-programmed input x follows from the correctness of the underlying PRF' .

Simulation security. $\text{PRF.Sim}(1^\lambda, P, L)$ is constructed as follows:

- Parse $P = \{(x_i, v_i)_{i \in [L']}\}$. Choose $L - L'$ strings of length at most $\ell + 1$ ending with 1, denoted $\{x'_j\}_{j \in [L-L']}$. Choose $v'_j \xleftarrow{\$} \mathcal{V}$ for $j \in [L - L']$.
- Output $\text{PRF}'.\text{Sim}(1^\lambda, \{x_i||0, v_i\}_{i \in [L']} \cup \{x'_j, v'_j\}_{j \in [L-L]}, L)$.

Suppose there is an efficient adversary \mathcal{A} that can break the simulation security of PRF with non-negligible probability, we can construct an efficient reduction \mathcal{B} that breaks the simulation security of the underlying PRF' with non-negligible probability. When \mathcal{A} submits $\{x_i, v_i\}_{i \in [L']}$ where $L' \leq L$, \mathcal{B} chooses $L - L'$ strings of length at most $\ell + 1$ ending with 1, denoted $\{x'_j\}_{j \in [L-L']}$. \mathcal{B} also chooses $v'_j \xleftarrow{\$} \mathcal{V}$ for $j \in [L - L']$. \mathcal{B} submits to its own challenger $\{x_i||0, v_i\}_{i \in [L']} \cup \{x'_j, v'_j\}_{j \in [L-L]}$, and obtains sk from its challenger. It forwards sk to \mathcal{A} . Now, whenever \mathcal{A} queries the point $x \in \{0, 1\}^{\leq \ell}$, \mathcal{B} forwards $x||0$ to its own challenger, obtains v , and forwards it to \mathcal{A} . Finally, \mathcal{B} outputs whatever \mathcal{A} outputs.

If \mathcal{B} is playing the game RealPPRF for the underlying PRF' , then \mathcal{A} 's view is identically distributed as in RealPPRF for PRF . If \mathcal{B} is playing the game IdealPPRF for the underlying PRF' , then \mathcal{A} 's view is identically distributed as in IdealPPRF for PRF .

Private programmability. Let $\text{PRF.Sim}(1^\lambda, L) = \text{PRF}'.\text{Sim}(1^\lambda, L)$. Suppose there is an efficient adversary \mathcal{A} that can break the private programmability of PRF with non-negligible probability, we can construct an efficient reduction \mathcal{B} that breaks the private programmability of the underlying PRF' with non-negligible probability. When \mathcal{A} submits $\{x_i\}_{i \in [L']}$ where $L' \leq L$, \mathcal{B} chooses $L - L'$ strings of length at most $\ell + 1$ ending with 1, denoted $\{x'_j\}_{j \in [L-L']}$. \mathcal{B} submits to its own challenger $\{x_i||0\}_{i \in [L']} \cup \{x'_j\}_{j \in [L-L]}$, and obtains sk from its challenger. It then outputs whatever \mathcal{A} outputs.

If \mathcal{B} is playing the game RealPPRFPriv of the underlying PRF' , then \mathcal{A} 's view is identically distributed as RealPPRFPriv of PRF . If \mathcal{B} is playing IdealPPRFPriv of the underlying PRF' , then \mathcal{A} 's view is identically distributed as IdealPPRFPriv of PRF . \square

3.2 Single-Server Private Information Retrieval

We define a single-server private information retrieval (PIR) scheme in the pre-processing setting. In a single-server PIR scheme, we have two stateful machines called the client and the server. The scheme consists of two phases:

- **Offline setup.** The offline setup phase is run only once upfront. The client receives nothing as input, and the server receives a database $\text{DB} \in \{0, 1\}^n$ as input. The client sends a single message to the server, and the server responds with a single message.
- **Online queries.** This phase can be repeated multiple times. Upon receiving an index $x \in \{0, 1, \dots, n - 1\}$, the client sends a single message to the server, and the server responds with a single message. The client performs some computation and outputs an answer $\beta \in \{0, 1\}$.

Correctness. Given a database $\text{DB} \in \{0, 1\}^n$, where the bits are indexed $0, 1, \dots, n - 1$, the correct answer for a query $x \in \{0, 1, \dots, n - 1\}$ is the x -th bit of DB .

For correctness, we require that for any q, n , that are polynomially bounded in λ , there is a negligible function $\text{negl}(\cdot)$, such that for any database $\text{DB} \in \{0, 1\}^n$, for any sequence of queries $x_1, x_2, \dots, x_q \in \{0, 1, \dots, n - 1\}$, an honest execution of the PIR scheme with DB and queries x_1, x_2, \dots, x_q , returns all correct answers with probability $1 - \text{negl}(\lambda)$.

Privacy. We say that a single-server PIR scheme satisfies privacy, iff there exists a probabilistic polynomial-time simulator Sim , such that for probabilistic polynomial-time adversary \mathcal{A} acting as the server, \mathcal{A} 's views in the following two experiments are computationally indistinguishable:

- **Real:** an honest client interacts with \mathcal{A} who acts as the server and may arbitrarily deviate from the prescribed protocol. In every online step t , \mathcal{A} may adaptively choose the next query $x_t \in \{0, 1, \dots, n - 1\}$ for the client, and the client is invoked with x_t ;
- **Ideal:** the simulated client Sim interacts with \mathcal{A} who acts as the server. In every online \mathcal{A} may adaptively choose the next query $x_t \in \{0, 1, \dots, n - 1\}$, and Sim is invoked without receiving x_t .

3.3 Fully Homomorphic Encryption

A fully homomorphic encryption scheme (FHE) with respect to a class of circuits \mathcal{C} is a tuple $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Eval}, \mathbf{Dec})$ of efficient, possibly randomized algorithms, with the following syntax:

- $\mathbf{Gen}(1^\lambda)$: receives the security parameter λ and outputs a secret key fsk .
- $\mathbf{Enc}(\text{fsk}, m)$: receives a secret key fsk and message m , and outputs a ciphertext c .
- $\mathbf{Eval}(\text{Circ}, c_1, \dots, c_d)$: receives a circuit $\text{Circ} \in \mathcal{C}$ with d inputs, as well as d ciphertexts and outputs a ciphertext c .
- $\mathbf{Dec}(\text{fsk}, c)$: receives a secret key fsk and ciphertext c , and outputs a plaintext m .

Correctness. Let \mathcal{C} be a class of circuits, Circ be an arbitrary element in \mathcal{C} , and d be the input size of Circ . A FHE scheme $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Eval}, \mathbf{Dec})$ is correct with respect to \mathcal{C} , if $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ is a correct encryption scheme, and there is a negligible function $\text{negl}(\cdot)$ such that for every security parameter λ , for all messages m_1, \dots, m_d , for any $\text{Circ} \in \mathcal{C}$, the following holds with at least $1 - \text{negl}(\lambda)$ probability: $\text{fsk} \leftarrow \mathbf{Gen}(1^\lambda)$, for $i \in [d]$, $c_i \leftarrow \mathbf{Enc}(\text{fsk}, m_i)$, $c' \leftarrow \mathbf{Eval}(\text{Circ}, c_1, \dots, c_d)$, then, it must be that $\mathbf{Dec}(\text{fsk}, c') = \text{Circ}(m_1, \dots, m_d)$.

Semantic security. We say that an FHE scheme $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Eval}, \mathbf{Dec})$ is semantically secure iff $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ is semantically secure.

3.4 The Distribution \mathcal{D}_n

For convenience, we often write $x \in \{0, 1, \dots, n - 1\}$ as a binary string, i.e., $x \in \{0, 1\}^{\log n}$.

Our pseudorandom set emulates the same distribution \mathcal{D}_n that was defined earlier in Shi et al. [SACM21]. Specifically, to define the distribution \mathcal{D}_n , imagine that we have a random oracle $\text{RO}(\cdot)$ that is sampled at random upfront — our actual PRSet scheme later will replace the RO with a PRF so our construction does not need an RO. Henceforth, let $B := \lceil 2 \log \log n \rceil$. An element $x \in \{0, 1\}^{\log n}$ is in the set iff for every $i \in [\frac{\log n}{2} + B]$, $\text{RO}((0^B || x)[i : \cdot])$ returns 1 — in other words,

if hashing every sufficiently long suffix of the string $0^B||x$ using the random oracle RO gives back 1. Throughout the paper, we write $\log = \log_2$, and assume that $\log n$ is an even integer — this is without loss of generality since we can always round it up to an even number incurring only constant blowup.

Efficient membership test and set enumeration. One important observation about the distribution \mathcal{D}_n is that the decisions regarding whether two elements x and y are in the set or not can be weakly dependent — as Shi et al. [SACM21] pointed out, this property is important for simultaneously ensuring efficient membership test and efficient set enumeration. Clearly, to test if an element $x \in \{0, 1\}^{\log n}$ is in the set or not, we only need to make $\frac{\log n}{2} + B$ calls to the RO.

Enumerating all elements in the set can be accomplished by making roughly $\sqrt{n} \cdot \text{poly} \log n$ calls to RO with at least $1 - o(1)$ probability. Let $\ell \geq \frac{1}{2} \log n + 1$, and let Z_ℓ be the set of all strings z of length exactly ℓ , such that using RO to “hash” all suffixes of z of length at least $\frac{1}{2} \log n + 1$, outputs 1. To enumerate the set generated by RO, we can start with $Z_{\frac{1}{2} \log n + 1}$ which takes at most $2^{\frac{1}{2} \log n + 1}$ calls to generate. Then, for each $\ell := \frac{1}{2} \log n + 2$ to $\log n$, we will generate Z_ℓ from $Z_{\ell-1}$. This can be accomplished by enumerating all elements $z' \in Z_{\ell-1}$, and checking whether $\text{RO}(0||z') = 1$ and $\text{RO}(1||z') = 1$. Finally, for every element $z \in Z_{\log n}$, we check if it is the case that for every $j \in [B]$, $0^j||z$ hashes to 1. If so, the element z is in the set.

Useful properties of \mathcal{D}_n . We will need to use the following useful facts about the distribution \mathcal{D}_n all of which were proven by Shi et al. [SACM21].

Fact 3.6. *For any fixed $x \in \{0, 1, \dots, n-1\}$, $\Pr_{S \leftarrow \mathcal{D}_n} [x \in S] = \frac{1}{\sqrt{n} \cdot 2^B}$. Moreover, $\mathbb{E}_{S \leftarrow \mathcal{D}_n} [|S|] \leq \frac{\sqrt{n}}{\log^2 n}$.*

Henceforth, let \mathcal{D}_n^{+x} be the following distribution: sample $S \leftarrow \mathcal{D}_n$ subject to $x \in S$. Given $x, y \in \{0, 1\}^{\log n}$, we say that x and y are *related*, if they share a common suffix of length at least $\frac{1}{2} \log n + 1$. Given a set $S \subseteq \{0, 1, \dots, n-1\}$, let $N_{\text{related}}(S, x)$ be the number of elements in S that are related to x .

Fact 3.7 (Number of related elements in sampled set). *Fix an arbitrary element $x \in \{0, 1, \dots, n-1\}$. Then,*

$$\mathbb{E}_{S \leftarrow \mathcal{D}_n^{+x}} [N_{\text{related}}(S, x)] \leq \frac{1}{\log n}$$

Fact 3.8 (Coverage probability). *Let $m \geq 6\sqrt{n} \cdot \log^3 n$. For any fixed $x \in \{0, 1, \dots, n-1\}$, $\Pr_{S_1, \dots, S_m \leftarrow \mathcal{D}_n^m} [x \notin \cup_{i \in [m]} S_i] \leq 1/n$.*

Henceforth, let $\text{EnumTime}(\text{RO})$ denote the number of RO calls made by the aforementioned set enumeration algorithm to enumerate the set generated by RO.

Fact 3.9 (Efficient set enumeration). *Suppose that $n \geq 4$. For any fixed $x \in \{0, 1, \dots, n-1\}$,*

$$\Pr_{\text{RO} \leftarrow \mathcal{D}_n^{+x}} [\text{EnumTime}(\text{RO}) > 6\sqrt{n} \log^5 n] \leq 1/\log n$$

4 Privately Programmable Pseudorandom Set

4.1 Definition

In our Privately Programmable Pseudorandom Set (PRSet) scheme, we can sample a key \mathbf{sk} that defines a pseudorandom set. We can support two operations on the key: we can call $\mathbf{Add}(\mathbf{sk}, x)$ to force x to be added to the set, we can also call $\mathbf{ReSamp}(\mathbf{sk}, x)$ to cause the decision whether x is in the set or not to be resampled. The key output by a \mathbf{ReSamp} operation is said to be *final*, i.e., we cannot perform any more operations on it. By contrast, keys output by either \mathbf{Gen} or \mathbf{Add} are said to be *intermediate*, i.e., we can still perform more operations on them. Henceforth, we use the notation \mathbf{rsk} to denote a final key and \mathbf{sk} to denote an intermediate key. Jumping ahead, later in our PIR scheme, the client always sends to the server a final key during an online query; however, the client locally stores a set of intermediate keys.

- $\mathbf{sk} \leftarrow \mathbf{Gen}(1^\lambda, n)$: given the security parameter 1^λ and the universe size n , samples a secret key \mathbf{sk} ;
- $S \leftarrow \mathbf{Set}(\mathbf{rsk})$: a deterministic algorithm that outputs a set S given a final secret key \mathbf{rsk} ;
- $b \leftarrow \mathbf{Member}(\mathbf{sk}, x)$: given an intermediate secret key \mathbf{sk} and an element $x \in \{0, 1, \dots, n-1\}$, output a bit indicating whether $x \in \mathbf{Set}(\mathbf{sk})$;
- $\mathbf{sk}_{+x} \leftarrow \mathbf{Add}(\mathbf{sk}, x)$: given an intermediate secret key \mathbf{sk} and an element $x \in \{0, 1, \dots, n-1\}$, output a secret key \mathbf{sk}_{+x} such that $x \in \mathbf{Set}(\mathbf{sk}_{+x})$;
- $\mathbf{rsk}_{-x} \leftarrow \mathbf{ReSamp}(\mathbf{sk}, x)$: given an intermediate secret key \mathbf{sk} and an element $x \in \{0, 1, \dots, n-1\}$, output a final key \mathbf{rsk}_{-x} that “resamples” the decision whether x is in the set or not.

We note that a PRSet scheme is parametrized by a family of distributions \mathcal{D}_n . The pseudorandom set generated by the PRSet scheme should emulate the distribution \mathcal{D}_n — we will define this more formally shortly.

Jumping ahead, later in our application, for each PRSet key sampled using \mathbf{Gen} , we perform at most one \mathbf{Add} operation on the key before we perform \mathbf{ReSamp} and obtain a final key.

Efficiency requirements. Our PRSet scheme samples pseudorandom sets of size roughly \sqrt{n} . We want an efficient set enumeration algorithm $\mathbf{Set}(\mathbf{rsk})$ that takes time roughly \sqrt{n} (rather than linear in n). Additionally, we want that the membership test $\mathbf{Member}(\mathbf{sk}, x)$ to complete in polylogarithmic time.

Remark 4.1. We do not give security definitions to our PRSet. Jumping ahead, the privacy proof of our PIR scheme actually opens up the PRSet scheme and relies on the properties of the underlying PRF directly. Nonetheless, abstracting out the PRSet helps to make the description of our PIR scheme conceptually cleaner.

4.2 Construction

We now present our PRSet construction. As mentioned, we assume that for each key sampled through \mathbf{Gen} , at most one \mathbf{Add} operation can be performed on the key before we call \mathbf{ReSamp} which produces a final key.

Intuition for our PRSet. In our pseudorandom set, we simply replace the RO with a PRF function, such that its description can be compressed using a short key.

Our pseudorandom set supports two additional operations:

- The **Add**(sk, x) operation modifies the secret key sk such that the element $x \in \{0, 1\}^{\log n}$ is forced to be in the set. In our construction, this is done in the most naïve way: simply attach the element x to the secret key. This will be fine in our PIR construction since the intermediate key generated by **Add** is stored only on the client side and never sent to the server. Therefore, we do not need the resulting key to hide the point x that is added.
- The **ReSamp**(sk, x) operation takes in an intermediate key that is either the output of **Gen** or the output of a previous **Add** operation, and it resamples the decision whether the element $x \in \{0, 1\}^{\log n}$ is in the set or not. In our PIR scheme later, this resampled key will be sent to the server during online queries. Therefore, we want the resulting key to hide not only the element x that is being resampled, but also the element x' that was added earlier should the input key sk be the result of a previous **Add**(\cdot, x') operation.

In our construction, this is accomplished in the following way. First, we sample at random the answers $\{v_i\}_{i \in [\frac{\log n}{2} + B]}$ — we want to force the PRF's evaluation at points $\{(0^B || x)[i :]\}_{i \in [\frac{\log n}{2} + B]}$ to be the values $\{v_i\}_{i \in [\frac{\log n}{2} + B]}$. Next, if the input key sk is the result of a previous **Add**(\cdot, x') operation, for any point $(0^B || x')[i :]$ where $i \in [\frac{\log n}{2} + B]$, if $(0^B || x')[i :] \neq (0^B || x)[i :]$, then we want to force the PRF's evaluation on $(0^B || x')[i :]$ to be 1. Finally, we call the underlying PRF's **Prog** function, to force the aforementioned outcomes on all the relevant points. Clearly, the total number of constraints to be forced is at most $L = 2(\frac{\log n}{2} + B)$.

Detailed construction. We describe our PRSet construction below.

PRSet Scheme

Parameters: $B := \lceil 2 \log \log n \rceil$, $L = 2(\frac{\log n}{2} + B)$.

- $\text{sk} \leftarrow \mathbf{Gen}(1^\lambda, n)$: call $\text{msk} \leftarrow \text{PRF.Gen}(1^\lambda, L)$, and output $\text{sk} := (\text{msk}, \perp)$.
- $S \leftarrow \mathbf{Set}(\text{rsk})$: Same as the set enumeration algorithm in Section 3.4, except that the calls to $\text{RO}(\cdot)$ are now replaced with calls to $\text{PRF.PEval}(\text{rsk}, \cdot)$.
- $b \leftarrow \mathbf{Member}(\text{sk}, x)$:
 1. Parse $\text{sk} := (\text{msk}', x')$. Write $x \in \{0, 1\}^{\log n}$ as a binary string and let $z := 0^B || x$. If $x' \neq \perp$, write $x' \in \{0, 1\}^{\log n}$ as a binary string and let $z' := 0^B || x'$.
 2. Output 1 if for every $i \in [\frac{\log n}{2} + B]$, the following holds: either $\text{PRF.Eval}(\text{msk}', z[i :]) = 1$ or $(x' \neq \perp$ and $z[i :] = z'[i :])$. Else, output 0.
- $\text{sk}_{+x} \leftarrow \mathbf{Add}(\text{sk}, x)$: parse $\text{sk} := (\text{msk}', \perp)$, and output $\text{sk}_{+x} := (\text{msk}', x)$.
- $\text{rsk}_{-x} \leftarrow \mathbf{ReSamp}(\text{sk}, x)$:
 1. Parse $\text{sk} := (\text{msk}', x')$, and write $x \in \{0, 1\}^{\log n}$ as a binary string and let $z := 0^B || x$.
 2. Sample uniformly random $v \xleftarrow{\$} \{0, 1\}^{\frac{\log n}{2} + B}$, and let $P := \{(z[i :], v[i])\}_{i \in [\frac{\log n}{2} + B]}$.
 3. If $x' \neq \perp$, do the following. Write $x' \in \{0, 1\}^{\log n}$ as a binary string, and let $z' := 0^B || x'$. For $i \in [\frac{\log n}{2} + B]$, if $z'[i :] \neq z[i :]$, add the constraint $(z'[i :], 1)$ to the set P .

4. Compute $\text{rsk}_{-x} \leftarrow \text{PRF.Prog}(\text{msk}', P)$, and output rsk_{-x} .

Additional helpful notations. In our PIR scheme later, we will only need to call set enumeration for final keys rsk . Therefore, our algorithm description above defines $\text{Set}(\text{rsk})$ only for final keys. However, in our proofs and narratives, it helps to define the set associated with an intermediate key sk as well — however, in this case we need not worry about the running time of $\text{Set}(\text{sk})$. This is defined in the most natural manner:

- If $\text{sk} = (\text{msk}, \perp)$ is the directly output of $\text{Gen}(1^\lambda, n)$, then $\text{Set}(\text{sk})$ is defined just like in Section 3.4 except that calls to $\text{RO}(\cdot)$ are replaced with $\text{PRF.Eval}(\text{msk}, \cdot)$;
- If $\text{sk} = (\text{msk}, x)$ is the output of an earlier **Add** operation, then $\text{Set}(\text{sk})$ is defined just like in Section 3.4 except that calls to $\text{RO}(\cdot)$ are replaced with the following outcomes: 1) we force the outcomes to be 1 at the input points $\{(0^B || x)[i :]\}_{i \in [\frac{\log n}{2} + B]}$; and 2) for all other inputs, we call $\text{PRF.Eval}(\text{msk}, \cdot)$ to obtain the outcome.

Performance bounds. $\text{Gen}(1^\lambda, n)$ takes $\text{poly}(\lambda, \log n)$ time. Due to Fact 3.9, $\text{Set}(\text{rsk})$ takes $\sqrt{n} \cdot \text{poly} \log(\lambda, n)$ time with $1 - 1/\log n$ probability. $\text{Member}(\text{sk}, x)$ takes $\text{poly}(\lambda, \log n)$ time. $\text{Add}(\text{sk}, x)$ takes constant time. $\text{ReSamp}(\text{sk}, x)$ takes $\text{poly}(\lambda, \log n)$ time.

Circuit for set enumeration. Later in our PIR scheme, during the offline phase, the server needs to perform set enumeration under fully homomorphic encryption. Therefore, we need to describe how to perform set enumeration in circuit. We will describe a circuit construction of size at most $\sqrt{n} \cdot \text{poly}(\lambda, \log n)$ which obtains as input a final key rsk , and outputs a set $S = \{(x, b)\}$ of size at most $2\sqrt{n} \log^2 n$ with distinct x 's, and a bit bSucc indicating success. We want to ensure that if $\text{bSucc} = \text{True}$, then the set generated is correct in the following sense:

- for every $(x, 1) \in S$, x is in the correct set defined by $\text{PRF.PEval}(\text{rsk}, \cdot)$; and
- for every element x in the set defined by $\text{PRF.PEval}(\text{rsk}, \cdot)$, the pair $(x, 1)$ appears in S .

Our circuit construction emulates the set enumeration algorithm of Section 3.4. Our circuit construction works as follows — henceforth we use the term “hash” to mean the computing outcome of $\text{PRF.PEval}(\text{rsk}, \cdot)$:

Circuit for set enumeration CSetEnum

1. Let $\text{bSucc} = \text{True}$.
2. For every $x \in \{0, 1\}^{\frac{1}{2} \log n + 1}$, let $b_x = \text{PRF.PEval}(\text{rsk}, x)$. Output an array containing $\{(x, b_x)\}_{x \in \{0, 1\}^{\frac{1}{2} \log n + 1}}$.
3. Obviously sort above array such that entries with $b_x = 1$ are moved to the front. Truncate the array at length $2\sqrt{n} \log^2 n$ elements, and if the truncation removes any string that hash to 1, set $\text{bSucc} = \text{False}$. Let $Z_{\frac{1}{2} \log n + 1}$ be the resulting truncated array, where each entry is of the form (x, b_x) .
4. For $\ell = \frac{1}{2} \log n + 2$ to $\log n$, do the following:
 - For each $(x, b_x) \in Z_{\ell-1}$, if $b_x = 1$, write down $(0 || x, \text{PRF.PEval}(\text{rsk}, 0 || x))$ and $(1 || x, \text{PRF.PEval}(\text{rsk}, 1 || x))$; else write down $(0 || x, 0)$ and $(1 || x, 0)$.

- Oblivious sort the resulting array such that all entries marked with 1 move to the front. Truncate the resulting array at length exactly $2\sqrt{n}\log^2 n$. If the truncation removes any string that hash to 1, set $\text{bSucc} = \text{False}$. Let Z_ℓ denote the resulting array where each entry is of the form (x, b_x) .

5. For every $(x, b_x) \in Z_{\log n}$, check if it is the case that for every $j \in [B]$, $\text{PRF.PEval}(\text{rsk}, 0^j || x) = 1$. If so, write down (x, b_x) , else, write down $(x, 0)$. Output the resulting array as well as bSucc .

Fact 4.2. *Using the AKS sorting network [AKS83] or the bitonic sorting network [Bat68] to realize the oblivious sort, the above algorithm can be implemented with a circuit of size $\sqrt{n} \cdot \text{poly}(\lambda, \log n)$.*

Proof. The proof is straightforward given the fact that the AKS sorting circuit has size $O(n' \log n')$ for sorting n' elements, and the bitonic sorting network has size $O(n' \log^2 n')$. Also, note that each $\text{PEval}(\text{rsk}, \cdot)$ consumes $\text{poly}(\lambda, \log n)$ gates to implement. \square

For correctness, we will imagine that the above algorithm is run where $\text{PRF.PEval}(\text{rsk}, \cdot)$ is replaced with calls to a random oracle RO — we denote the resulting algorithm as $\text{CSetEnum}^{\text{RO}}$. Note that we do not care about the computational model when stating the correctness probability.

Fact 4.3. *Suppose that $n \geq 4$. For any $x \in \{0, 1, \dots, n-1\}$,*

$$\Pr_{\text{RO} \leftarrow \mathcal{D}_n^{+x}} \left[\text{CSetEnum}^{\text{RO}} \text{ outputs } \text{bSucc} = \text{True} \right] \geq 1 - 1/\log n,$$

Moreover,

$$\Pr_{\text{RO} \leftarrow \mathcal{D}_n} \left[\text{CSetEnum}^{\text{RO}} \text{ outputs } \text{bSucc} = \text{True} \right] \geq 1 - 1/\log n$$

Proof. $\text{CSetEnum}^{\text{RO}}$ is a direct implementation of the set enumeration algorithm in Section 3.4 except that we truncate each Z_ℓ to size exactly $2\sqrt{n}\log^2 n$. Shi et al. [SACM21] proved that no matter whether RO is sampled from \mathcal{D}_n^{+x} or \mathcal{D}_n , with $1 - 1/\log n$ probability, the following good event holds: for all $\ell \in [\frac{\log n}{2} + 1, \log n]$, $|Z_\ell| \leq 2\sqrt{n}\log^2 n$ — see the proof of Lemma 6.4 in their paper. The algorithm outputs $\text{bSucc} = 1$ as long as the above good event holds. \square

5 PIR Scheme

We now describe a PIR scheme that supports a bounded number of queries denoted Q . Given this scheme, we can compile it to a scheme that supports unbounded number of queries by performing the offline setup phase every Q queries, and amortizing this cost over the Q queries.

Intuition. In the offline setup phase, the client chooses $\tilde{O}(Q)$ keys each of which defines a pseudorandom set of size roughly \sqrt{n} . It encrypts these keys under a fully homomorphic encryption (FHE) scheme, and sends the encrypted keys to the server. Through homomorphic evaluation, the server computes the encrypted parity of each of these sets, and returns the encrypted parities to client. The client decrypts the parities, and stores each set's key as well as its parity. These sets are divided into two parts: the last Q entries are called the *backup* sets or entries, and the remaining are called the *primary* sets or entries. The primary entries are used for answering queries, whereas the backup entries are later promoted to become primary entries as they get consumed. Henceforth,

we also use the terms primary table and backup table to refer to the tables that store all primary entries and backup entries, respectively.

In the online phase, whenever the client wants to make a query for the database's value at index $x \in \{0, 1, \dots, n-1\}$, it finds the first primary set (sk_i, p_i) such that $\mathbf{Set}(\text{sk}_i)$ contains the query x . It then resamples the decision whether x is in the set or not, and obtains a programmed key. It sends this programmed key to the server, which calls the set enumeration algorithm to enumerate the set S generated by the key. The server then returns the parity p of the set S to the client. The client computes $p_i \oplus p$ as the candidate answer to the query. Since the resampling operation removes the element x from the set with high probability, the candidate answer is correct with high probability. The correctness probability can be further boosted by repeating the same scheme k times and taking the majority vote among the k copies.

Detailed construction. We describe the detailed construction below.

PIR Scheme for $Q = \sqrt{n}$ queries

Run $k = \omega(\log \lambda)$ parallel copies of the single-copy scheme described below.

Offline phase:

- **Client:** // let $\text{lenT} := 6\sqrt{n} \cdot \log^3 n$
 - $\text{fsk} \leftarrow \text{FHE.Gen}(1^\lambda)$;
 - For $i \in [k \cdot (\text{lenT} + Q)]$ where $k = \omega(\log \lambda)$, $\text{sk}_i \leftarrow \text{PRSet.Gen}(1^\lambda, n)$, $\overline{\text{sk}}_i \leftarrow \text{FHE.Enc}(\text{fsk}, \text{sk}_i)$;
 - Send $(\overline{\text{sk}}_1, \dots, \overline{\text{sk}}_{k \cdot (\text{lenT} + Q)})$ to the server.
- **Server:**
 - For $i \in [k \cdot (\text{lenT} + Q)]$, $(\overline{S}_i, \overline{\text{bSucc}}_i) \leftarrow \text{FHE.Eval}(\text{CSetEnum}, \overline{\text{sk}}_i)$;
 - $\{\overline{p}_i\}_{i \in [k \cdot (\text{lenT} + Q)]} \leftarrow \text{FHE.Eval}(\text{CBatchParity}, \overline{S}_1, \dots, \overline{S}_{k \cdot (\text{lenT} + Q)})$, where the CBatchParity circuit is described below. Send $\{\overline{p}_i, \overline{\text{bSucc}}_i\}_{i \in [k \cdot (\text{lenT} + Q)]}$ to the client.
- **Client:**
 - for $i \in [k \cdot (\text{lenT} + Q)]$, $p_i \leftarrow \text{FHE.Dec}(\text{fsk}, \overline{p}_i)$; $\text{bSucc}_i \leftarrow \text{FHE.Dec}(\text{fsk}, \overline{\text{bSucc}}_i)$;
 - choose a subset $I \subseteq [k \cdot (\text{lenT} + Q)]$ of size exactly $\text{lenT} + Q$ such that for any $i \in I$, $\text{bSucc}_i = \text{True}$ — if not enough such entries are found, simply abort. Copy $\{(\text{sk}_i, p_i)\}_{i \in I}$ to a table.

We call the last Q entries of the above table the *backup* table, henceforth renamed to $T^* := \{(\text{sk}_i^*, p_i^*)\}_{i \in [Q]}$. We call the remaining lenT entries the *primary* table, henceforth renamed to $T := \{(\text{sk}_i, p_i)\}_{i \in [\text{lenT}]}$.

Online query for index $x \in \{0, \dots, n-1\}$:

- **Client:**
 - Sample $\text{sk} \leftarrow \text{PRSet.Gen}(1^\lambda, n)$ subject to $\text{PRSet.Member}(\text{sk}, x) = 1$ and append the entry $(\text{sk}, 0)$ to the table T of primary sets;
 - Find the first entry (sk_i, p_i) in T such that $\text{PRSet.Member}(\text{sk}_i, x) = 1$;

– Compute $\text{rsk} \leftarrow \text{PRSet.ReSamp}(\text{sk}_i, x)$ and send rsk to the server.

- **Server:** Compute $S \leftarrow \text{PRSet.Set}(\text{rsk})$, and return the parity bit p of the set S to the client. If the set enumeration algorithm has not completed even after making $6\sqrt{n} \log^5 n$ calls to the underlying PRF's $\text{PEval}(\text{rsk}, \cdot)$ function, then return $p = 0$ to the client.
- **Client:** let $\beta' := p \oplus p_i$ be the candidate answer of the current copy, and remove the last entry of T .

Recall that there are k parallel instances, and let β be the majority vote among the candidate answers of all k copies. Now, let (sk_j^*, p_j^*) denote the next available backup set and perform the following:

- let $\text{sk}' \leftarrow \text{PRSet.Add}(\text{sk}_j^*, x)$; let $p' := p_j^* \oplus \beta$ if $\text{Member}(\text{sk}_j^*, x) = 0$, else let $p' := p_j^*$;
- let $T_j := (\text{sk}', p')$, and mark the backup entry (sk_j^*, p_j^*) as unavailable.

The circuit CBatchParity. The circuit CBatchParity takes $S_1, S_2, \dots, S_{k \cdot (\text{lenT} + Q)}$ as input, where for $j \in [k \cdot (\text{lenT} + Q)]$, S_j contains exactly $2\sqrt{n} \log^2 n$ entries of the form (x, b_x) — specifically, $b_x = \text{True}$ implies that x is the j -th set and $b_x = \text{False}$ implies x is not in the j -th set. The circuit outputs $k \cdot (\text{lenT} + Q)$ parity bits $p_1, \dots, p_{k \cdot (\text{lenT} + Q)}$ of each of the $k \cdot (\text{lenT} + Q)$ sets.

The circuit can be constructed as follows using oblivious sort:

1. Let $\text{DB} \in \{0, 1\}^n$ be the server's database, let $\mathbf{D} := ((0, \text{DB}[0]), (1, \text{DB}[1]), \dots, (n-1, \text{DB}[n-1]))$, concatenated with all the sets $S_1, \dots, S_{k \cdot (\text{lenT} + Q)}$.
2. For $j \in [k \cdot (\text{lenT} + Q)]$, let $\mathbf{X}_j = \{(x, b_x, j)\}_{x \in S_j}$
3. Obviously sort the array $\mathbf{Y} := \mathbf{D} \parallel \mathbf{X}_1 \parallel \dots \parallel \mathbf{X}_{k \cdot (\text{lenT} + Q)}$, such that each entry of the form $(x, \text{DB}[x])$ is followed by all tuples of the form (x, b_x, j) . Henceforth, we call a tuple of the form (x, b_x, j) a consumer.
4. In a linear scan, all consumers receive the $\text{DB}[x]$ they are requesting. At this moment, each consumer entry is updated to $(x, b_x, j, \text{DB}[x])$.
5. Use a circuit that mirrors the oblivious sort circuit in Step 3, and reverse-routes the $\text{DB}[x]$ values back to the position where it came from. As a result, each consumer entry of the form $(x, b_x, j) \in \mathbf{Y}$ receives $\text{DB}[x]$.
6. At this moment, we have an array of the form $\mathbf{X}'_1 \parallel \dots \parallel \mathbf{X}'_{k \cdot (\text{lenT} + Q)}$, where each \mathbf{X}'_j contains exactly $2\sqrt{n} \log^2 n$ entries of the form $(x, b_x, j, \text{DB}[x])$. In a linear scan, we can compute for each $j \in [k \cdot (\text{lenT} + Q)]$, the parity bit

$$p_j = \bigoplus_{(x, b_x, j, \text{DB}[x]) \in \mathbf{X}'_j} (b_x \cdot \text{DB}[x])$$

It is not hard to see that if we instantiate the oblivious sort using either AKS [AKS83] or bitonic sort [Bat68], and given $\text{lenT} = 6\sqrt{n} \log^3 n$ and $Q = \sqrt{n}$, the above circuit has size $O(n \cdot \text{poly} \log n)$.

Performance bounds. We now analyze the performance bounds of our Q -bounded PIR construction. We may plug in $k = \log^{1.1} n$ since any super-logarithmic function will work. In the analysis below, the k parameter is absorbed in the $\text{poly} \log n$ term, so it does not show up explicitly.

- *Offline phase.* During the offline phase, the client's computation and bandwidth are upper bounded by $\sqrt{n} \cdot \text{poly}(\lambda, \log n)$. The server's computation is upper bounded by $n \cdot \text{poly}(\lambda, \log n)$.

- *Online phase.* The bandwidth is $\text{poly}(\lambda, \log n)$. The client’s computation is $\sqrt{n} \cdot \text{poly}(\lambda, \log n)$. The server’s computation is also $\sqrt{n} \cdot \text{poly}(\lambda, \log n)$.

Supporting unbounded number of queries and deamortization. To extend the scheme from Q -bounded to supporting an unbounded number of queries, we just need to rerun the offline phase every $Q = \sqrt{n}$ queries. For the scheme with unbounded queries, the amortized bandwidth per query is $\text{poly}(\lambda, \log n)$, the amortized client and server computation per query is $\sqrt{n} \cdot \text{poly}(\lambda, \log n)$.

This periodic offline setup can be deamortized very easily. Specially, upfront, we perform the offline setup for $2Q$ queries. During the i -th window of Q queries, we perform the offline setup for the $(i + 2)$ -th window of Q queries, and so on. This way, when the $(i + 2)$ -th window of Q queries starts, the corresponding offline setup will be ready.

6 Privacy Proof

Recall that privacy for a single-server PIR scheme was defined earlier in Section 3.2. We now prove that our PIR scheme in Section 5, when instantiated with the PRSet scheme in Section 4.2, satisfies privacy, as stated in the following theorem.

Theorem 6.1 (Privacy of our PIR scheme). *Suppose that the FHE scheme employed satisfies semantic security, and that the underlying programmable PRF scheme satisfies correctness, private programmability, and simulation security. Then, the PIR scheme in Section 5, when instantiated with the PRSet scheme in Section 4.2, satisfies privacy.*

In the remainder of this section, we will prove the above theorem.

6.1 Proof Roadmap

The privacy proof turns out to be challenging because the cryptography appears to be intertwined with the analysis of the underlying randomized process. Intuitively, during an online query for $x \in \{0, 1, \dots, n - 1\}$, the matched key from the primary table T should generate a set whose distribution is computationally indistinguishable from “sampling a random set from \mathcal{D}_n subject to containing x ”. Unfortunately, the technicality is that the distribution of the matched intermediate key itself does not admit a simple description, partly because there are two possible cases depending on whether the intermediate key is the outcome of a previous **Add** operation or not. If so, the key is of the form (msk, y) for some $y \in \{0, 1, \dots, n - 1\}$ that was queried before; else, the key is of the form (msk, \perp) .

Alternate view of the random process. To accomplish the proof, our high-level idea is to decompose the computational reductions and the probability analysis. To do so, we view the client’s random process of maintaining the hint table and sending PRSet keys to the server during online queries in the following light (defined more formally in **Hyb₃** later):

1. *Maintain constraints on each entry that defines the a-posteriori distribution.* Let $\mathbf{I} = \{i_1, i_2, \dots, i_q\}$ be the indices of the entries that are matched during each of the $q \leq Q$ queries so far. The client maintains the a-posteriori distribution of each entry of the primary table T conditioned on the local observation \mathbf{I} .

To maintain the a-posteriori distribution, the client maintains a set constraints of the form $\langle -x \rangle$, $\langle +x \rangle$, $\langle +y : -x \rangle$, or $\langle +y : +x \rangle$ on each entry. A negative constraint of the form $\langle -x \rangle$ means

that this entry was not promoted from the backup table, and we know that x is not in the set generated by the key. A negative constraint of the form $\langle +y : -x \rangle$ means that this entry was promoted from the backup table during a query for y , and we know that after forcing y to be in the set, x is not in the set generated by the key. The positive constraints $\langle +x \rangle$ and $\langle +y : +x \rangle$ are similarly defined but requiring x to be in the set.

During an online query for some x , the client sequentially scans through the current entries of T . For each entry j , it samples from the a-posteriori distribution to decide if j should be the match. Depending on the decision, it adds either a negative or positive constraint to the current entry.

2. *Lazy sampling from the a-posteriori distribution.* Whenever the client is about to send a key to the server, it performs lazy sampling of the key based on the a-posteriori distribution on the entry that the client has maintained. More specifically, there are two cases depending on whether the matched entry comes from the backup table or not : 1) it samples a key from the correct a-posteriori distribution, calls **ReSamp** and sends the resulting key to the server; 2) it samples a key from the correct a-posteriori distribution, calls both **Add** and **ReSamp**, and then sends the resulting key to the server.

Proof roadmap. Once we view the random process as the above, we can proceed with the proof as follows.

- First, for maintaining the a-posteriori distributions on keys locally, whenever promoting a backup key to the primary table, instead of storing the element x next to the key to force x to be in the set, we sample a random master key subject to containing x instead. This is described more formally in **Hyb₄** later.
- Next, recall that to compute the key to send to the server, the client needs to lazy-sample a key and perform **Add + ReSamp** or simply **ReSamp**. We now switch to a hybrid experiment in which the client need not call **Add** or **ReSamp** any more, but instead samples a simulated key on the fly subject to an appropriate set of constraints. This switch is possible due to the simulation security as well as the private programmability of the PRF. This step is formally described in **Hyb₅**.
- Next, we switch to a hybrid experiment in which during a query for some index x the key sent to the server is generated from the following process: sample a master secret key **msk** subject to an appropriate set of constraints, then call **ReSamp** to program the key to random values at all coins related to the decision whether x is in the set. In this hybrid, the keys sent to the server are the result of **ReSamp**, but we no longer need to call **Add**. This step is formally described in **Hyb₆**.
- At this moment, we can argue that all keys sent to the server are computationally indistinguishable from simulated keys generated freshly at random. The remainder of the argument is similar to (but slightly more involved than) the privacy proof of Shi et al. [**SACM21**] — note that in the scheme of Shi et al. [**SACM21**], the keys sent to the server are also the result of a **ReSamp** operation.

Technicalities. To make the proof work, there are a few technicalities. First, in one key hybrid (specifically, **Hyb₅**), we want to remove the **Add** and **ReSamp** operations; instead, we switch to sampling simulated keys subject to an appropriate set of constraints. In the next hybrid (i.e., **Hyb₆**),

we want to add **ReSamp** back to the keys sent to the server, but without performing an **Add** operation. The proof of these key steps rely on the simulation security and private programmability of the PRF, and the proof is non-trivial, partly because the security definitions of the programmable PRF are not stated in a form that match our needs. To facilitate this proof, we prove a technical lemma about the programmable PRF first in Section 6.2 first. This technical lemma states a property of the PRF that is closer to the form we want later in the sequence of hybrid experiments.

Second, in several steps of our proof, we must sample simulated keys subject to a set of constraints. To make the proof work, we need to make sure that the probability that a randomly sampled key satisfies these constraints is non-negligible. Therefore, we prove some useful facts about the distribution \mathcal{D}_n in Section 6.3.

Finally, to show indistinguishability of Hyb_6 and the ideal experiment in which the server receives random simulated keys, we need to view the random process from the perspective of the server again. Unlike the client, the server cannot observe the client's local variables \mathbf{I} . Our entire proof critically relies on being able to switch between the two alternative views of the same random process.

6.2 Technical Lemma for Privately Programmable PRF

We shall consider a programmable PRF whose output range is binary, i.e., $\{0, 1\}$. Henceforth, we use the notation $\text{pred}^X(\text{msk})$ to denote an event that looks at the outputs of $\text{PRF.Eval}(\text{msk}, \cdot)$ at all inputs in X , and outputs either 0 or 1. We say that $\text{pred}^X(\cdot)$ is an *admissible* event, iff 1) for a randomly sampled $\text{msk} \leftarrow \mathbf{Gen}(1^\lambda, L)$, it returns 1 with probability at least $1/p(\lambda)$ for some polynomial function $p(\cdot)$; and 2) pred is polynomial-time checkable.

Lemma 6.2 (Strong privacy of programmable PRF). *Let PRF be a programmable PRF with a binary output range, and suppose that $L = O(\log \lambda)$. Suppose that PRF satisfies private programmability and simulation security. Then, there exists a probabilistic polynomial-time simulator Sim such that the following two experiments are computationally indistinguishable to any probabilistic polynomial-time adversary.*

- $\text{RealPPRFStrong}(1^\lambda)$:
 - $X, X', \{v_x\}_{x \in X'}, \text{pred}^{X \cup X'} \leftarrow \mathcal{A}(1^\lambda, L)$ s.t. $|X| + |X'| \leq L$, $X \cap X' = \emptyset$, and $\text{pred}^{X \cup X'}(\cdot)$ is admissible;
 - for $x \in X$, let $v_x \stackrel{\$}{\leftarrow} \mathcal{V}$; let $P := \{(x, v_x)\}_{x \in X \cup X'}$;
 - sample $\text{msk} \leftarrow \mathbf{Gen}(1^\lambda, L)$ subject to $\text{pred}^{X \cup X'}(\text{msk}) = 1$, and let $\text{sk} \leftarrow \mathbf{Prog}(\text{msk}, P)$;
 - $\text{sk} \rightarrow \mathcal{A}$;
- $\text{IdealPPRFStrong}(1^\lambda)$:
 - $X, X', \{v_x\}_{x \in X'}, \text{pred}^{X \cup X'} \leftarrow \mathcal{A}(1^\lambda, L)$ s.t. $|X| + |X'| \leq L$, $X \cap X' = \emptyset$, and $\text{pred}^{X \cup X'}(\cdot)$ is admissible;
 - sample $\text{sk} \leftarrow \text{Sim}(1^\lambda, L)$ subject to the constraint that for any $x \in X'$, $\mathbf{PEval}(\text{sk}, x) = v_x$;
 - $\text{sk} \rightarrow \mathcal{A}$.

Proof. We consider the following intermediate hybrid experiment called **Hyb**. **Hyb** is almost the same as RealPPRFStrong except that when we sample the msk , it is sampled at random rather than subject to the constraint that $\text{pred}^{X \cup X'}(\text{msk}) = 1$.

Claim 6.3. *Suppose that the programmable PRF satisfies simulation security. Then, RealPPRFStrong is computationally indistinguishable from Hyb.*

Proof. Let Sim' be the simulator as in the definition of simulation security. Through a straightforward reduction to simulation security, Hyb is computationally indistinguishable from the following hybrid experiment denoted H:

- $X, X', \{v_x\}_{x \in X'}, \text{pred}^{X \cup X'} \leftarrow \mathcal{A}(1^\lambda, L)$ s.t. $|X| + |X'| \leq L$, $X \cap X' = \emptyset$, and $\text{pred}^{X \cup X'}(\cdot)$ is admissible;
- for $x \in X$, let $v_x \xleftarrow{\$} \mathcal{V}$; let $P := \{(x, v_x)\}_{x \in X \cup X'}$;
- $\text{sk} \leftarrow \text{Sim}'(1^\lambda, P, L)$;
- $\text{sk} \rightarrow \mathcal{A}$;

It suffices to show that H is computationally indistinguishable from RealPPRFStrong. We show that if there is an efficient adversary \mathcal{A} that can distinguish H and RealPPRFStrong with non-negligible probability, we can construct an efficient reduction \mathcal{B} that breaks the simulation security of the PRF scheme with non-negligible probability. Specifically, \mathcal{B} waits till \mathcal{A} submits $X, X', \{v_x\}_{x \in X'}, \text{pred}^{X \cup X'}$, it then chooses v_x at random for $x \in X$, and lets $P := \{(x, v_x)\}_{x \in X \cup X'}$. It gives P to its own challenger. It obtains a key sk from its own challenger. It then queries its challenger on the inputs $X \cup X'$, and checks if pred holds over the outcomes. If so, it gives sk to \mathcal{A} and outputs the same guess as \mathcal{A} . Otherwise, it outputs a random guess.

If \mathcal{B} is playing in the experiment RealPPRF with its own challenger, then, conditioned on $\text{pred}^{X \cup X'}$ being true, \mathcal{A} 's view is identically distributed as RealPPRFStrong. Else, if \mathcal{B} is playing in the experiment IdealPPRF with its own challenger, then, conditioned on $\text{pred}^{X \cup X'}$ being true, \mathcal{A} 's view is identically distributed as H.

Let p be the probability that the predicate $\text{pred}^{X \cup X'}$ holds if \mathcal{B} is playing in the experiment RealPPRF, and let p' be the probability that the predicate $\text{pred}^{X \cup X'}$ holds if \mathcal{B} is playing in the experiment IdealPPRF with its challenger. Since $\text{pred}^{X \cup X'}$ is admissible, p must be a non-negligible function. It must be that $|p - p'| \leq \text{negl}(\lambda)$ since otherwise we can easily construct an efficient adversary that distinguishes RealPPRF and IdealPPRF with non-negligible probability.

Therefore, if \mathcal{A} has a non-negligible advantage in distinguishing RealPPRFStrong and H, \mathcal{B} has a non-negligible advantage in distinguishing RealPPRF and IdealPPRF. \square

Claim 6.4. *Suppose that the programmable PRF satisfies private programmability. Then Hyb is computationally indistinguishable from IdealPPRFStrong, where the simulator Sim is the same as in the private programmability definition.*

Proof. We show that if there is an efficient adversary \mathcal{A} that can distinguish Hyb and IdealPPRFStrong with non-negligible probability, then we can construct an efficient reduction \mathcal{B} that can break private programmability with non-negligible probability. \mathcal{B} obtains $X, X', \{v_x\}_{x \in X'}, \text{pred}^{X \cup X'}$ from \mathcal{A} . It then forwards $X \cup X'$ to its own challenger, and obtains sk from its own challenger. It then checks if it is the case that for every $x \in X'$, $\text{PEval}(\text{sk}, x) = v_x$. If so, it forwards sk to \mathcal{A} , and outputs whatever \mathcal{A} outputs. Otherwise, \mathcal{B} outputs a random guess. If \mathcal{B} is playing the game RealPPRF with its own challenger, then conditioned on \mathcal{A} receiving sk from \mathcal{B} , \mathcal{A} 's view is identically distributed as in Hyb. Else, if \mathcal{B} is playing IdealPPRF with its challenger, then conditioned on \mathcal{A} receiving sk from \mathcal{B} , \mathcal{A} 's view is identically distributed as IdealPPRFStrong.

Let p be the probability that \mathcal{B} forwards sk to \mathcal{A} when \mathcal{B} is playing RealPPRF, and let p' be the corresponding probability when \mathcal{B} is playing IdealPPRF. Since $L \leq O(\log \lambda)$ and the PRF has a

binary output domain, we know that $p \geq 1/\text{poly}(\lambda)$. Moreover, $|p' - p| \leq \text{negl}(\lambda)$ since otherwise, we can easily construct an efficient adversary that distinguishes between RealPPRF and IdealPPRF with non-negligible probability. Therefore, if \mathcal{A} has non-negligible advantage in distinguishing Hyb and IdealPPRFStrong, then \mathcal{B} has non-negligible advantage in distinguishing RealPPRF and IdealPPRF. \square

6.3 Useful Facts about the Distribution \mathcal{D}_n

We define the following helpful notation where $x \in \{0, 1\}^{\log n}$:

$$\text{suffixes}(x) := \{(0^B || x)[i : \cdot]\}_{i \in [\frac{\log n}{2} + B]}$$

We first describe a couple useful facts which will later be used in our hybrid sequence.

Fact 6.5. *Consider two arbitrary elements $x, y \in \{0, 1, \dots, n-1\}$ which may be different or the same. There is a polynomial function $\text{poly}(\cdot)$ such that $\Pr_{S \leftarrow \mathcal{D}_n} [x, y \in S] \geq 1/\text{poly}(n)$. Or equivalently, let $\text{RO}(\cdot)$ denote a random oracle. Then, there is some polynomial function $\text{poly}(\cdot)$, such that the following event happens with at least $1/\text{poly}(n)$ probability: $\text{RO}(\cdot)$ outputs 1 on every input from $\text{suffixes}(x) \cup \text{suffixes}(y)$.*

Proof. The proof is straightforward. Since there are at most $2(\frac{\log n}{2} + B)$ points that we want to force to be 1, the probability that this happens is at least

$$\frac{1}{2^{2(\frac{\log n}{2} + B)}} \geq \frac{1}{n \log^5 n}$$

for sufficiently large n . \square

Intuitively, the following fact states that given the distribution \mathcal{D}_n , conditioned on one element x or two elements $x, y \in \{0, 1, \dots, n-1\}$ being in the set, the probability that up to \sqrt{n} other elements are not in the set must be at least inverse polynomial.

Fact 6.6. *Consider two arbitrary elements $x, y \in \{0, 1, \dots, n-1\}$ which may be different or the same, and $Q' \leq \sqrt{n}$ other elements $x_1, \dots, x_{Q'}$ such that $x_j \neq x$ and $x_j \neq y$ for any $j \in [Q']$. Then, there is some polynomial function $\text{poly}(\cdot)$, such that*

$$\Pr_{S \leftarrow \mathcal{D}_n} [\forall j \in [Q'] : x_j \notin S | x, y \in S] \geq 1/\text{poly}(n)$$

Or equivalently, let $\text{RO}(\cdot)$ denote a random oracle. Then, there is some polynomial function $\text{poly}(\cdot)$, such that the following event happens with at least $1/\text{poly}(n)$ probability over the choice of RO : for every $j \in [Q']$, $\text{RO}(\cdot)$ does not always output 1 over the input set $\text{suffixes}(x_j) \setminus (\text{suffixes}(x) \cup \text{suffixes}(y))$.

Proof. It suffices to prove the lemma for the worst case $Q' = Q = \sqrt{n}$. For convenience, define the following event Ev_j^- for $j \in [Q]$:

$$\text{Ev}_j^- : \text{RO}(\cdot) \text{ does not always output 1 over the input set } \text{suffixes}(x_j) \setminus (\text{suffixes}(x) \cup \text{suffixes}(y))$$

Let $\text{size}(x_j) = |\text{suffixes}(x_j) \setminus (\text{suffixes}(x) \cup \text{suffixes}(y))|$, then $\Pr[\text{Ev}_j^-] \geq 1 - 1/2^{\text{size}(x_j)}$. Note also that for any Ev_j^- , any set $I \subseteq [n]$ such that $j \notin I$, it holds that

$$\Pr[\text{Ev}_j^- | \{\text{Ev}_{j'}^-\}_{j' \in I}] \geq \Pr[\text{Ev}_j^-]$$

Therefore, we have that

$$\Pr[\text{Ev}_1^-, \dots, \text{Ev}_Q^-] \geq \prod_{j \in [Q]} \Pr[\text{Ev}_j^-]$$

Observe that as long as $x_j \neq x$ and $x_j \neq y$, $\text{size}(x_j) > B$. Further, there are at most 4 choices of x_j such that $\text{size}(x_j) = B + 1$, at most 8 choices of x_j such that $\text{size}(x_j) = B + 2$, at most 16 choices of x_j such that $\text{size}(x_j) = B + 3$, and so on. Therefore, we have the following where $\ell = \frac{\log n}{2}$:

$$\begin{aligned} \Pr[\text{Ev}_1^-, \dots, \text{Ev}_Q^-] &\geq \prod_{j \in [Q]} \Pr[\text{Ev}_j^-] \\ &\geq \left(1 - \frac{1}{2^{B+1}}\right)^4 \cdot \left(1 - \frac{1}{2^{B+2}}\right)^8 \cdot \left(1 - \frac{1}{2^{B+3}}\right)^{16} \cdot \dots \cdot \left(1 - \frac{1}{2^{B+\ell-1}}\right)^{2^\ell} \\ &\geq \left(1 - \frac{2}{\log^2 n}\right)^\ell \geq 1 - \frac{1}{\log n} \end{aligned}$$

□

6.4 Sequence of Hybrid Experiments

To prove Theorem 6.1, we define a sequence of hybrid experiments and show that the adversary \mathcal{A} 's views in every pair of adjacent hybrids are either identically distributed or computationally indistinguishable.

Experiment Real. Same as the real-world execution where an honest client interacts with \mathcal{A} acting as the server. Henceforth, we may assume that during the online queries, the client skips the steps of FHE decryption and computing the answer to the query. However, it still deletes the last entry of the table T (which was added earlier during the online query); further, it still promotes the next available backup entry to a primary entry. Note that locally skipping the FHE decryption and computation of the answer does not affect the distribution of the messages the client sends to the server \mathcal{A} .

Experiment Hyb_1 . Experiment Hyb_1 is almost the same as **Real** except that during the offline setup phase, the client replaces all FHE ciphertexts sent to the server with encryptions of 0.

Claim 6.7. *Suppose that the FHE scheme satisfies semantic security. Then, Hyb_1 is computationally indistinguishable from **Real**.*

Proof. Follows in a straightforward manner from the semantic security of FHE. □

Experiment Hyb_2 . Experiment Hyb_2 is the same as Hyb_1 except that we modify the client to record some extra information as it interacts with the server. Specifically, for each entry in the primary table, the client maintains a set of constraints. Initially, the constraint sets are all empty. During each online query for the index $x \in \{0, 1, \dots, n-1\}$, the client performs the following — below the text in blue denotes the additional actions taken by the client:

- Sample $\text{sk} \leftarrow \text{PRSet.Gen}(1^\lambda, n)$ subject to $\text{PRSet.Member}(\text{sk}, x) = 1$ and append the entry $(\text{sk}, 0)$ to the table T of primary sets; record the additional constraint $\langle +x \rangle$ for this entry;
- Find the first entry (sk_i, p_i) such that $\text{PRSet.Member}(\text{sk}_i, x) = 1$;
 - for every entry $j < i$ in T , if the entry (sk_j, p_j) was earlier promoted from a backup set during a query for y , record the additional constraint $\langle +y : -x \rangle$; else record the additional constraint $\langle -x \rangle$ for this entry.
 - For the i -th entry (sk_i, p_i) , if the entry (sk_i, p_i) was earlier promoted from a backup set during a query for y , record the additional constraint $\langle +y : +x \rangle$; else record the additional constraint $\langle +x \rangle$ for this entry.
- The rest of the client’s algorithm is the same as in Hyb_2 , except that when the client promotes a backup entry to the primary table T , it records the fact that the entry was promoted during a query for the index x , and empties the constraint set associated with the relevant entry.

Intuitively, the constraint $\langle +x \rangle$ (or $\langle -x \rangle$, resp.) mean that x should (or should not, resp.) be contained in the set; the constraint $\langle +x \rangle$ means that x should be contained in the set. The constraint $\langle +y : +x \rangle$ (or $\langle +y : -x \rangle$) means that after calling **Add** to force-add the element y , the element x should (or should not, resp.) be in the set.

Hyb_2 is clearly identically distributed as Hyb_1 since recording the extra information does not affect the distribution.

Experiment Hyb_3 . Experiment Hyb_3 is the same as Hyb_2 except that during the online phase, whenever the client is about to send the key of a pseudorandom set to the server, instead of sending the key computed by the honest algorithm, the client resamples a fresh key corresponding to the desired constraints, performs a corresponding **Add** operation on it if the key was promoted from a backup key, and then performs a **ReSamp** operation. More concretely, whenever the client is about to send a key for a pseudorandom set to the server during an online query, it instead sends the following “lazy-sampled” key:

- If the entry found in the table T was earlier promoted from a backup entry during a query for y , then, repeat: let $\text{sk} \leftarrow \text{Gen}(1^\lambda, n)$ and let $\text{sk}_{+y} \leftarrow \text{Add}(\text{sk}, y)$, until the following constraints are satisfied:
 - for every constraint of the form $\langle +y : -x \rangle$ recorded for this entry, it must be that $x \notin \text{Set}(\text{sk}_{+y})$;
 - for every constraint of the form $\langle +y : +x \rangle$ recorded for this entry, it must be that $x \in \text{Set}(\text{sk}_{+y})$;

Now, call $\text{rsk} \leftarrow \text{ReSamp}(\text{sk}_{+y}, x)$ where x is the current query, and send rsk to the server.

- Else, sample $\text{sk} \leftarrow \text{Gen}(1^\lambda, n)$ subject to the following constraints:
 - for every constraint of the form $\langle -x \rangle$ recorded for this entry, it must be that $x \notin \text{Set}(\text{sk})$;
 - for every constraint of the form $\langle +x \rangle$ recorded for this entry, it must be that $x \in \text{Set}(\text{sk})$;

Now, call $\text{rsk} \leftarrow \text{ReSamp}(\text{sk}, x)$ where x is the current query, and send rsk to the server.

Claim 6.8. Hyb_3 is identically distributed as Hyb_2 .

Proof. In Hyb_3 , we first sample an initial batch of pseudorandom set keys, and we use these keys to decide which primary entry is matched during each query. Let $\mathbf{I} := (i_1, i_2, \dots, i_Q)$ be the random variables that denote the index of the entry matched during each of the Q queries, and let $\mathbf{x} := x_1, x_2, \dots, x_Q$ be the Q queries. The experiment Hyb_3 is essentially maintaining the *a-posteriori* distribution of each primary entry conditioned on having observed the choices of \mathbf{I} and \mathbf{x} . Only when the client needs to send the key to the server (possibly after performing an **Add** operation on it and always after performing a **ReSamp** operation), we perform the actual sampling of the key from the *a-posteriori* distribution. Further, if the key was promoted from a backup key, we perform a corresponding **Add** operation; and we always perform a **ReSamp** operation on the lazily sampled key before sending it to the server. Therefore, the server \mathcal{A} 's views in Hyb_3 and Hyb_2 are identically distributed — in fact, this holds even if we allow \mathcal{A} to observe the choice of \mathbf{I} (which \mathcal{A} does not observe in the actual experiment). \square

Claim 6.9. *Suppose $Q = \sqrt{n}$. Except with negligible probability, Hyb_3 completes in polynomial-time. In other words, the lazy sampling of keys can be accomplished in polynomial time except with negligible probability.*

Proof. Observe that if the entry found in T was earlier promoted from a backup entry during a query for y , then there is exactly one positive constraint of the form $\langle +y : +x \rangle$, and at most Q negative constraints of the form $\langle +y : -x' \rangle$. Similarly, if the entry found in T was not promoted from a backup entry, then there is exactly one positive constraint of the form $\langle +x \rangle$, and at most Q negative constraints of the form $\langle -x' \rangle$.

The remainder of the proof follows in a straightforward fashion from the pseudorandomness of the underlying programmable PRF and due to Fact 6.5 and Fact 6.6. \square

Experiment Hyb_4 . Hyb_4 is otherwise the same as Hyb_3 except with the following modification. Recall that in Hyb_3 , upon receiving a new query x , we scan the primary table T , and for each key $\text{sk} \in T$, we check if $\text{Member}(\text{sk}, x) = 1$. If a key sk was promoted from the backup table during an earlier query y , the check is performed by calling $\text{PRF.Eval}(\text{sk}, \cdot)$ but forcing the outcomes on $\text{suffixes}(y)$ to be 1. In Hyb_4 , however, we make the following change:

- Whenever a backup key is promoted to become a primary key during a query for y , we replace the key with a resampled one, i.e., sample $\text{msk} \leftarrow \text{PRF.Gen}(1^\lambda, L)$ subject to $y \in \text{Set}(\text{msk})$. In other words, we are sampling a programmable PRF key subject to the constraint that its outcomes on any input from the set $\text{suffixes}(y)$ must be 1.

Claim 6.10. *Suppose that the programmable PRF satisfies private programmability, simulation security, and correctness. Then, Hyb_3 is computationally indistinguishable from Hyb_4 .*

Proof. We consider the following hybrid sequence. Experiment H is otherwise the same as Hyb_3 , except that when we promote a backup key to the primary table, we replace the consumed key with a key sk sampled from the following distribution: sample $\text{msk} \leftarrow \text{PRF.Gen}(1^\lambda, L)$, $\text{sk} \leftarrow \text{PRF.Prog}(\text{msk}, \{(z, 1)\}_{z \in \text{suffixes}(x)})$, where x is the current query. Due to the correctness of the PRF, Hyb_3 is statistically indistinguishable from H .

Experiment H' is otherwise the same as H , except that except that when we promote a backup key to the primary table, we replace the consumed key with a key sk sampled from the following distribution: sample $\text{sk} \leftarrow \text{PRF.Sim}(1^\lambda, L)$, subject to $x \in \text{Set}(\text{sk})$, where Sim is the simulator in the private programmability definition. H' is computationally indistinguishable from H due to Lemma 6.2.

Experiment H'' is otherwise the same as H' , except that when we promote a backup key to the primary table, we replace the consumed key with a key sk sampled from the following distribution: sample $msk \leftarrow \text{PRF.Gen}(1^\lambda, L)$, $sk \leftarrow \text{PRF.Prog}(msk, \emptyset)$ subject to $x \in \text{Set}(sk)$, where x is the current query. H'' is computationally indistinguishable from H' due to the private programmability of the PRF.

Finally, H'' is statistically indistinguishable from Hyb_4 , due to the correctness of the programmable PRF. \square

Experiment Hyb_5 . Hyb_5 is otherwise the same as Hyb_4 except with the following modification when performing lazy-sampling of the pseudorandom set key. Let x be the current query, and let Sim be the simulator in the private programmability definition.

- If the entry found in T was earlier promoted from the backup table during a query for y , then repeat: $sk \leftarrow \text{Sim}(1^\lambda, L)$ until sk satisfies the following constraints, and send sk to the server:
 - $\text{PRF.PEval}(sk, \cdot)$ outputs 1 on any input from the set $\text{suffixes}(y) \setminus \text{suffixes}(x)$;
 - for every negative constraint of the form $\langle +y : -x' \rangle$ that is recorded for this entry, $\text{PRF.PEval}(sk, \cdot)$ does not output all 1s on the input set $\text{suffixes}(x') \setminus (\text{suffixes}(x) \cup \text{suffixes}(y))$.
- Else, repeat: $sk \leftarrow \text{Sim}(1^\lambda, L)$ until sk satisfies the following constraints, and send sk to the server:
 - for every negative constraint of the form $\langle -x' \rangle$ that is recorded for this entry, $\text{PRF.PEval}(sk, \cdot)$ does not output all 1s on the input set $\text{suffixes}(x') \setminus \text{suffixes}(x)$.

Notably, in Hyb_5 , we no longer perform **Add** or **ReSamp** operations on the key sent to the server \mathcal{A} .

Claim 6.11. *Suppose that the underlying programmable PRF satisfies private programmability and simulation security. Then, Hyb_5 is computationally indistinguishable from Hyb_4 .*

Proof. Hyb_4 can be equivalently viewed as the following: during a query for index x , when we compute the lazy-sampled key, do the following depending on which case:

- *Case 1: the lazy-sampled key was promoted from the backup table during an earlier query for index y .* In this case, there is exactly one positive constraint of the form $\langle +y : +x \rangle$, and there can be at most Q negative constraints of the form $\langle +y : -x' \rangle$.

Therefore, the lazy-sampled key has the following distribution: sample a random $msk \leftarrow \text{PRF.Gen}(1^\lambda, L)$ subject to the following constraints:

- $\text{PRF.Eval}(msk, \cdot)$ evaluates to 1 on $\text{suffixes}(x) \setminus \text{suffixes}(y)$; and
- For each negative constraint of the form $\langle +y : -x' \rangle$, $\text{PRF.Eval}(msk, \cdot)$ does not evaluate to all 1s on the input set $\text{suffixes}(x') \setminus (\text{suffixes}(x) \cup \text{suffixes}(y))$.

Finally, we call **Prog** to program the resulting msk to force the outcomes at $\text{suffixes}(x)$ to be random, and the outcomes at $\text{suffixes}(y) \setminus \text{suffixes}(x)$ to be 1, and send the programmed key to the server \mathcal{A} .

- *Case 2: the lazy-sampled key was not promoted from the backup table.* In this case, there is exactly one positive constraint of the form $\langle +x \rangle$, and there can be at most Q negative constraints of the form $\langle -x' \rangle$. Therefore, the lazy-sampled key is sampled at random $msk \leftarrow \text{PRF.Gen}(1^\lambda, L)$, subject to the following constraints:

- $\text{PRF.Eval}(\text{msk}, \cdot)$ evaluates to 1 on $\text{suffixes}(x)$; and
- For each negative constraint of the form $\langle +y : -x' \rangle$, $\text{PRF.Eval}(\text{msk}, \cdot)$ does not evaluate to all 1s on the input set $\text{suffixes}(x') \setminus \text{suffixes}(x)$.

Finally, we call **Prog** to program the resulting msk to force the outcomes at $\text{suffixes}(x)$ to be random, and send the programmed key to the server \mathcal{A} .

To show that Hyb_4 is computationally indistinguishable from Hyb_5 , we can consider a sequence of hybrid experiments denoted $\text{H}_0, \dots, \text{H}_Q$. In H_i where $i \in \{0, 1, \dots, Q\}$, for the first i queries, we use the method of Hyb_5 to sample the key sent to the server, and for the remaining queries, we use the method of Hyb_4 to sample the key sent to the server. It suffices to show that every pair of adjacent hybrids are computationally indistinguishable.

We show that if there is an efficient adversary \mathcal{A} that can distinguish H_i and H_{i+1} where $i \in \{0, 1, \dots, Q-1\}$, we can construct an efficient reduction \mathcal{B} which can break the strong privacy of the underlying PRF (see Lemma 6.2) which is implied by private programmability and simulation security. Basically, \mathcal{B} acts as the client and interacts with \mathcal{A} like in H_i , except that for the i -th query, when it is about to send the key to \mathcal{A} , it performs the following instead:

- *Case 1: the i -th query wants to lazy-sample a key that was promoted from the backup table during an earlier query for the index y .* \mathcal{B} sends its challenger $X = \text{suffixes}(x)$, $X' = \text{suffixes}(y) \setminus \text{suffixes}(x)$, $\{v_x = 1\}_{x \in X'}$, and a predicate $\text{pred}^{X \cup X'}$ that wants the original unprogrammed PRF to output 1 at $\text{suffixes}(x) \setminus \text{suffixes}(y)$. \mathcal{B} obtains from its challenger some PRF key sk . If for every negative constraint of the form $\langle +y : -x' \rangle$, $\text{PRF.PEval}(\text{sk}, \cdot)$ does not evaluate to all 1s on $\text{suffixes}(x') \setminus (\text{suffixes}(x) \cup \text{suffixes}(y))$, then \mathcal{B} forwards sk to \mathcal{A} , and outputs the same guess as \mathcal{A} . Otherwise, \mathcal{B} aborts outputting a random guess. If \mathcal{B} is in the game RealPPRFStrong , and \mathcal{B} does not abort outputting a random guess, then \mathcal{A} 's view is identically distributed as in H_i . If \mathcal{B} is in the game IdealPPRFStrong , and \mathcal{B} does not abort outputting a random guess, then \mathcal{A} 's view is identically distributed as in H_{i+1} .

Let p be the probability that \mathcal{B} does not outputting a random guess when it is playing RealPPRFStrong , and let p' be the corresponding probability when it is playing IdealPPRFStrong . Due to Fact 6.6 and the pseudorandomness of the PRF, $p \geq 1/\text{poly}(\lambda, n)$. Further, it must be that $|p' - p| \leq \text{negl}(\lambda)$ since otherwise we can easily construct an adversary that can distinguish RealPPRFStrong and IdealPPRFStrong with non-negligible probability. Therefore, if \mathcal{A} has non-negligible advantage in distinguishing H_i and H_{i+1} , then \mathcal{B} has non-negligible advantage in distinguishing RealPPRFStrong and IdealPPRFStrong .

- *Case 2:* The proof of Case 2 is similar to Case 1 except that now, we replace $\text{suffixes}(y)$ with \emptyset .

□

Experiment Hyb_6 . Hyb_6 is otherwise the same as Hyb_5 , except the following modification when performing lazy-sampling of the pseudorandom set key. Let x be the current query, and let Sim be the simulator in the private programmability definition.

- If the entry found in T was earlier promoted from the backup table during a query for y , then repeat: $\text{msk} \leftarrow \text{Gen}(1^\lambda, L)$ until msk satisfies the following constraints:
 - $\text{PRF.Eval}(\text{msk}, \cdot)$ outputs 1 on any input from the set $\text{suffixes}(y) \cup \text{suffixes}(x)$;
 - for every negative constraint of the form $\langle +y : -x' \rangle$ that is recorded for this entry, $\text{PRF.Eval}(\text{msk}, \cdot)$ does not output all 1s on the input set $\text{suffixes}(x') \setminus (\text{suffixes}(x) \cup \text{suffixes}(y))$.

- Else, repeat: $\text{msk} \leftarrow \mathbf{Gen}(1^\lambda, L)$ until sk satisfies the following constraints, and send sk to the server:
 - for every negative constraint of the form $\langle -x' \rangle$ that is recorded for this entry, $\text{PRF.Eval}(\text{msk}, \cdot)$ does not output all 1s on the input set $\text{suffixes}(x') \setminus \text{suffixes}(x)$, and moreover it outputs 1 on $\text{suffixes}(x)$.

Now, call $\text{sk} \leftarrow \mathbf{Prog}(\text{msk}, \{(z, r_z)\}_{z \in \text{suffixes}(x)})$ where all r_z 's are sampled independently at random, and send sk to the server.

Claim 6.12. *Suppose that the PRF satisfies private programmability and simulation security. Then, Hyb_6 is computationally indistinguishable from Hyb_5 .*

Proof. The proof is very similar to that of Claim 6.11. To show that Hyb_5 is computationally indistinguishable from Hyb_6 , we can consider a sequence of hybrid experiments denoted H_0, \dots, H_Q . In H_i where $i \in \{0, 1, \dots, Q\}$, for the first i queries, we use the method of Hyb_5 to sample the key sent to the server, and for the remaining queries, we use the method of Hyb_6 to sample the key sent to the server. It suffices to show that every pair of adjacent hybrids are computationally indistinguishable.

We show that if there is an efficient adversary \mathcal{A} that can distinguish H_i and H_{i+1} where $i \in \{0, 1, \dots, Q-1\}$, we can construct an efficient reduction \mathcal{B} which can break the strong privacy of the underlying PRF (see Lemma 6.2) which is implied by private programmability and simulation security. Basically, \mathcal{B} acts as the client and interacts with \mathcal{A} like in H_i , except that for the i -th query, when it is about to send the key to \mathcal{A} , it performs the following instead:

- *Case 1: the i -th query wants to lazy-sample a key that was promoted from the backup table during an earlier query for the index y .* \mathcal{B} sends its challenger $X = \text{suffixes}(x)$, $X' = \emptyset$, and a predicate $\text{pred}^{X \cup X'}$ that wants the original unprogrammed PRF to output 1 at $\text{suffixes}(x)$. \mathcal{B} obtains from its challenger some PRF key sk . If for every negative constraint of the form $\langle +y : -x' \rangle$, $\text{PRF.PEval}(\text{sk}, \cdot)$ does not evaluate to all 1s on $\text{suffixes}(x') \setminus (\text{suffixes}(x) \cup \text{suffixes}(y))$, and moreover, it evaluates to all 1s on the input set $\text{suffixes}(y) \setminus \text{suffixes}(x)$, then \mathcal{B} forwards sk to \mathcal{A} , and outputs the same guess as \mathcal{A} . Otherwise, \mathcal{B} aborts outputting a random guess. If \mathcal{B} is in the game RealPPRFStrong , and \mathcal{B} does not abort outputting a random guess, then \mathcal{A} 's view is identically distributed as in H_i . If \mathcal{B} is in the game IdealPPRFStrong , and \mathcal{B} does not abort outputting a random guess, then \mathcal{A} 's view is identically distributed as in H_{i+1} .

Let p be the probability that \mathcal{B} does not outputting a random guess when it is playing RealPPRFStrong , and let p' be the corresponding probability when it is playing IdealPPRFStrong . Due to Fact 6.6 and the pseudorandomness of the PRF, $p \geq 1/\text{poly}(\lambda, n)$. Further, it must be that $|p' - p| \leq \text{negl}(\lambda)$ since otherwise we can easily construct an adversary that can distinguish RealPPRFStrong and IdealPPRFStrong with non-negligible probability. Therefore, if \mathcal{A} has non-negligible advantage in distinguishing H_i and H_{i+1} , then \mathcal{B} has non-negligible advantage in distinguishing RealPPRFStrong and IdealPPRFStrong .

- *Case 2: The proof of Case 2 is similar to Case 1 except that now, we replace $\text{suffixes}(y)$ with \emptyset .*

□

Fact 6.13. *Hyb_6 is identically distributed as the following process. During the offline phase, send FHE encryptions of 0 to the server. During each query, the client appends a key sampled at random from $\text{PRF.Gen}(1^\lambda, L)$ subject to containing x at the end of the table T . The client finds in T the first entry msk whose set contains x . The client sends $\text{PRF.Prog}(\text{msk}, \{(z, r_z)\}_{z \in \text{suffixes}(x)})$ to the*

server where r_z 's are independent random bits. The client then replaces the consumed entry with a fresh key sampled from $\text{PRF.Gen}(1^\lambda, L)$ subject to containing x , and deletes the last entry of the table.

Proof. By definition, the above process has the same local table distribution as in Hyb_6 . Recall that in Hyb_6 , the client samples another key subject to the corresponding constraints, programs the key and sends it to the server. In the above process, the client directly programs the key stored in the local table and sends it to the server. In Hyb_6 , conditioned on the matched indices \mathbf{I} , the matched entry in the table has the same a-posteriori distribution as the key we lazily sample which we then program and send to the server. Therefore, Hyb_6 can be equivalently rewritten as the randomized process in Fact 6.13. \square

Experiment Ideal. In the *Ideal* experiment, during the offline phase, the client sends FHE encryptions of 0 to the server. During each online query, whenever the client needs to send the server \mathcal{A} some key, it simply samples a key $\text{sk} \leftarrow \text{Sim}(1^\lambda, L)$ at random where Sim is the same simulator as in the private programmability definition, and sends sk to the server.

Claim 6.14. *Suppose that the PRF satisfies private programmability. Then Hyb_6 is computationally indistinguishable from *Ideal*.*

Proof. We can consider a sequence of hybrids denoted $\text{H}_0, \text{H}_1, \dots, \text{H}_Q$. In H_i , for the first i queries we do the following:

- The client sends the server a random simulated key sampled from $\text{sk} \leftarrow \text{Sim}(1^\lambda, L)$.

For the remaining queries, the client does the following just like in Hyb_6 (see Fact 6.13):

- Let x be the current query. The client appends a key sampled at random from $\text{PRF.Gen}(1^\lambda, L)$ subject to containing x at the end of the table T . The client finds the first entry msk whose set contains x . The client sends $\text{PRF.Prog}(\text{msk}, \{(z, r_z)\}_{z \in \text{suffixes}(x)})$ to the server where r_z 's are independent random bits. The client then replaces the consumed entry with a fresh key sampled from $\text{PRF.Gen}(1^\lambda, L)$ subject to containing x , and deletes the last entry of the table.

It suffices to prove that H_i and H_{i+1} are computationally indistinguishable for every $i \in \{0, 1, \dots, Q-1\}$. First, we prove the following fact.

Fact 6.15. *In H_i , conditioned on the server's view at the beginning of the $(i+1)$ -th query, the keys contained in the client's primary table T are identically distributed as sampling independent keys from $\text{msk} \leftarrow \text{Gen}(1^\lambda, L)$.*

Proof. From an information theoretic perspective, the server learns no information during the first i queries. Therefore, we can prove the claim by induction. The statement is true initially before any query is made. Now, suppose the statement is true at the end of the $(i'-1)$ -th query where $i' \leq i$, we prove that the statement is still true at the end of the i' -th query. It is easy to see this if we view the distribution of lenT randomly sampled simulated keys as the following distribution:

- First, sample the index $j^* \in [\text{lenT} + 1]$ which is the first entry that contains the queried element x . Note that since we appended to the table T a key sampled from $\text{PRF.Gen}(1^\lambda, L)$ subject to containing the current query x before searching through T , a satisfying key is guaranteed to be found.

- For any $j < j^*$, sample the j -th key at random from $\text{PRF.Gen}(1^\lambda, L)$ subject to not containing x ; sample the j^* -th key at random from $\text{PRF.Gen}(1^\lambda, L)$ subject to containing x ; and finally, for $j > j^*$, sample j -th key at random from $\text{PRF.Gen}(1^\lambda, L)$.

Now, during the i' -th query, we consume the j^* -th entry, and replace it with a key freshly sampled from $\text{PRF.Gen}(1^\lambda, L)$ subject to containing x . Therefore, this does not change the distribution. \square

Therefore, in H_i , during the $(i+1)$ -th query for some index x , the matched key has the following distribution: sample $\text{msk} \leftarrow \text{Gen}(1^\lambda, L)$ subject to containing x . Thus, the key returned to the server is identically distributed as: sample $\text{msk} \leftarrow \text{PRF.Gen}(1^\lambda, L)$ subject to $x \in \text{Set}(\text{msk})$ where x is the current query, call $\text{sk} \leftarrow \text{PRF.Prog}(\text{msk}, \{(z, r_z)\}_{z \in \text{suffixes}(x)})$ where the r_z 's are independently sampled random bits. Due to the private programmability of the PRF, we can replace the key sent to the server during the $(i+1)$ -th query with the outcome of $\text{Sim}(1^\lambda, L)$, which gives us H_{i+1} , and the adversary \mathcal{A} will not be able to distinguish the two except with negligible probability. \square

7 Correctness Proof

We now prove the correctness of our PIR scheme.

Offline phase. Due to Fact 4.3 and the pseudorandomness of the PRF, for each pseudorandom set key sk , the probability that CSetEnum returns $\text{bSucc} = \text{True}$ is at least $1 - 1/\log n$. Therefore, the probability that among $k = \omega(\log \lambda)$ copies, no copy returns $\text{bSucc} = \text{True}$ is negligibly small. Therefore, the probability that the offline phase cannot find $\text{lenT} + Q$ copies with $\text{bSucc} = \text{True}$ is negligibly small. Due to the correctness of the FHE, during the offline phase, the client obtains the correct parities for all $\text{lenT} + Q$ pseudorandom sets except with negligible probability.

Online phase. Given the above, to prove the correctness of our PIR scheme, it suffices to show the following: assume that to start with, the client is storing the correct parity bits for all of the pseudorandom sets. Then, each single copy of the PIR scheme is correct with probability at least $2/3$. If so, due to the standard Chernoff bound, when we do majority voting among $k = \omega(\log \lambda)$ copies, the majority vote is correct with all but $\text{negl}(\lambda)$ probability.

Experiment CReal. Same as the real-world experiment running a single copy of the PIR scheme, except that at the end of each query, we force the client's parity bits to be all correct (even if the client may have computed an incorrect parity bit).

Experiment CIdeal. Consider the following experiment where the client stores each set using a random oracle RO_j rather than a pseudorandom key.

Ideal correctness experiment CIdeal

Offline phase. Client generates $\text{lenT} + Q$ random oracles denoted $\{\text{RO}_j\}_{j \in [\text{lenT} + Q]}$, where each RO_j defines a random set. Let $\text{Label}_j = \perp$ for $j \in [\text{lenT}]$. The client obtains the correct parity bit p_j for each random set. The $\text{lenT} + Q$ random oracles are divided into lenT primary entries which form the primary table T , as well as Q backup entries.

Online query for $x \in \{0, 1, \dots, n - 1\}$.

- Client overwrites the $(\text{len}T + 1)$ -th entry of the primary table T to be the following random oracle: sample a fresh RO^* and force RO^* 's outcomes at $\text{suffixes}(x)$ to be 1.
- Client finds the first primary entry j such that the set defined by RO_j contains x . If the entry found is the last entry of T , return ErrNotFound .
- If $y := \text{Label}_j \neq \perp$ and the set generated by RO_j contains other elements related to y , return ErrParity .
- Client resamples RO_j 's outcomes at the points $\text{suffixes}(x)$. If this resampling ends up removing some element $x' \neq x$ from the set, return ErrReSampFail . If the resampled RO_j takes more than $6\sqrt{n} \log^5 n$ RO calls to enumerate the set, then return ErrTimeOut .
- Client takes the first unconsumed backup entry denoted RO^* , forces RO^* 's outputs at $\text{suffixes}(x)$ to be 1, and then uses the resulting random oracle to replace the j -th entry of T . Further, set $\text{Label}_j = x$. Return Success .

Intuitively, ErrNotFound characterizes the probability that the queried element is not in any of the primary sets; ErrParity represents the probability that when we promoted a backup entry to the primary table by forcefully adding some element y , it caused some element(s) related to y to be added to the set — in this case, the parity associated with this entry could be incorrect. Finally, ErrTimeOut represents the probability of an error caused by the set enumeration timing out.

Let $\text{Wrong}^{i, \text{CReal}}(x_1, \dots, x_Q)$ denote the event that upon the query sequence x_1, \dots, x_Q , the client computes the incorrect answer during the i -th query in experiment CReal . Let $\text{Wrong}^{i, \text{CIdeal}}(x_1, \dots, x_Q)$ denote the event that upon the query sequence x_1, \dots, x_Q , the client returns either ErrNotFound or ErrReSampFail or ErrTimeOut during the i -th query in experiment CIdeal .

Claim 7.1. *Assume that the programmable PRF satisfies pseudorandomness and correctness. Then, there exists a negligible function $\text{negl}(\cdot)$ such that for any $x_1, \dots, x_Q \in \{0, 1, \dots, n - 1\}$, for any $i \in [Q]$,*

$$\Pr[\text{Wrong}^{i, \text{CReal}}(x_1, \dots, x_Q)] \leq \Pr[\text{Wrong}^{i, \text{CIdeal}}(x_1, \dots, x_Q)] + \text{negl}(\lambda)$$

Proof. Experiment CIdeal makes the following modifications to CReal : 1) replaces PRF evaluations to RO calls, and 2) remove all instructions not related to correctness. Note that the event $\text{Wrong}^{i, \text{CReal}}$ depends only on the evaluation outcomes of the PRF and does not depend on the PRF key itself. Therefore, the claim follows in a straightforward fashion from the pseudorandomness and the correctness of the underlying programmable PRF. \square

Claim 7.2. *For any $x_1, \dots, x_Q \in \{0, 1, \dots, n - 1\}$, for any $i \in [Q]$, $\Pr[\text{Wrong}^{i, \text{CIdeal}}(x_1, \dots, x_Q)] \leq 1/3$.*

Proof. To prove the claim, we will make use of the following fact:

Fact 7.3. *In CIdeal , the RO_j found during each query has the following distribution: sample an RO at random subject to containing x . Moreover, at the end of each query, the table T (ignoring the $(\text{len}T + 1)$ -th entry) has the same distribution as $\text{len}T + Q$ independently sampled ROs.*

Proof. We can prove the fact using a similar argument as Fact 6.15, except that now, each entry of the table T is a random oracle instead of a key sampled from $\text{Sim}(1^\lambda, L)$. \square

We can now bound the probability of each type of error.

ErrNotFound. Due to Fact 7.3 and Fact 3.8, $\Pr[\text{ErrNotFound}] \leq 1/n$.

ErrReSampFail. Due to Fact 7.3 and Fact 3.7, we have $\Pr[\text{ErrReSampFail}] \leq 1/\log n$.

ErrTimeOut. Due to Fact 7.3 and Fact 3.9, $\Pr[\text{ErrTimeOut}] \leq 1/\log n$.

ErrParity. This is the most complicated error to bound. To bound the probability of ErrParity, we may equivalently consider the following experiment which is obtained from CIdeal, but removing all other errors we do not care about right now.

Experiment CIdealParity
// same experiment as the one in Lemma 7.7 in Shi et al. [SACM21]

Offline setup. For $j = 1$ to $\text{len}T$: sample a random oracle RO and let $T_j := \text{RO}$. Set $\text{Label}_j := \perp$.

Online query for index $x \in \{0, 1, \dots, n-1\}$.

- a) Sample a new RO^* such that the associated set contains x . Append RO^* to the table T as the last entry, and mark its label $\text{label}(T_{\text{len}T+1}) := \perp$.
- b) Let $T_j := \text{RO}_j$ be the smallest entry in the table T such that the set generated by RO_j contains x .
- c) If $y := \text{Label}_j \neq \perp$ and the set generated by RO_j contains other elements related to y , then return ErrParity.
- d) Sample a new RO' such that the generated set contains x . Overwrite $T_j := \text{RO}'$ and set $\text{Label}_j := x$.
- e) Remove the last entry from T and return Success.

We can show that in the above experiment, the probability that the i -th query returns ErrParity is upper bounded by $2/\log n$ using exactly the same approach as Shi et al. [SACM21]. Specifically, in the above experiment, for the i -th query to return ErrParity, there are two cases:

1. The i -th query for index x_i finds an entry with the $y = \text{Label}_j \neq \perp$, and x_i is related to y .
2. The i -th query for index x_i finds an entry with the $y = \text{Label}_j \neq \perp$, and x_i is not related to y .

In the proof of Lemma 7.7 of Shi et al. [SACM21], they argue that due to Fact 7.3 and Fact 3.7, the probability of the first case happening is upper bounded by the probability that a random RO subject to containing x_i also contains another element related to x_i , which is upper bounded by $1/\log n$. Through a more complicated argument, they also show that the probability of the second case happening is also upper bounded by $1/\log n$.

Therefore, we have that $\Pr[\text{ErrParity}] \leq 2/\log n$. □

Acknowledgment

This work is in part supported by an NSF award under the grant number CIF-1705007.

References

- [ACLS18] Sebastian Angel, Hao Chen, Kim Laine, and Srinath T. V. Setty. PIR with compressed queries and amortized query processing. In *S&P*, 2018.
- [AIVG20] Kinan Dak Albab, Rawane Issa, Mayank Varia, and Kalman Graffi. Batched differentially private information retrieval. Cryptology ePrint Archive, Report 2020/1596, 2020. <https://ia.cr/2020/1596>.
- [AKS83] M. Ajtai, J. Komlós, and E. Szemerédi. An $O(n \log n)$ sorting network. In *STOC*, 1983.
- [Bat68] Kenneth E. Batchier. Sorting networks and their applications. In *American Federation of Information Processing Societies: AFIPS Conference Proceedings: 1968 Spring Joint Computer Conference, Atlantic City, NJ, USA, 30 April - 2 May 1968*, pages 307–314, 1968.
- [BGI16] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In *CCS*, 2016.
- [BIM00] Amos Beimel, Yuval Ishai, and Tal Malkin. Reducing the servers computation in private information retrieval: Pir with preprocessing. In *CRYPTO*, pages 55–73, 2000.
- [BIPW17] Elette Boyle, Yuval Ishai, Rafael Pass, and Mary Wootters. Can we access a database both locally and privately? In *TCC*, 2017.
- [BKM17] Dan Boneh, Sam Kim, and Hart William Montgomery. Private puncturable PRFs from standard lattice assumptions. In *EUROCRYPT*, pages 415–445, 2017.
- [BLW17] Dan Boneh, Kevin Lewi, and David J. Wu. Constraining pseudorandom functions privately. In *PKC*, 2017.
- [BTVW17] Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained prfs (and more) from LWE. In *TCC*, 2017.
- [CC17] Ran Canetti and Yilei Chen. Constraint-hiding constrained PRFs for NC^1 from LWE. In *EUROCRYPT*, pages 446–476, 2017.
- [CG97] Benny Chor and Niv Gilboa. Computationally private information retrieval. In *STOC*, 1997.
- [CGKS95] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *FOCS*, 1995.
- [Cha04] Yan-Cheng Chang. Single database private information retrieval with logarithmic communication. In *ACISP*, 2004.
- [CHK22] Henry Corrigan-Gibbs, Alexandra Henzinger, and Dmitry Kogan. Single-server private information retrieval with sublinear amortized time. In *Eurocrypt*, 2022.
- [CHR17] Ran Canetti, Justin Holmgren, and Silas Richelson. Towards doubly efficient private information retrieval. In *TCC*, 2017.

- [CK20] Henry Corrigan-Gibbs and Dmitry Kogan. Private information retrieval with sublinear online time. In *EUROCRYPT*, 2020.
- [CKGS98] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *J. ACM*, 45(6):965–981, November 1998.
- [CMS99] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In *EUROCRYPT*, pages 402–414, 1999.
- [DCIO98] Giovanni Di-Crescenzo, Yuval Ishai, and Rafail Ostrovsky. Universal service-providers for database private information retrieval. In *PODC*, 1998.
- [dCP22] Leo de Castro and Antigoni Polychroniadou. Lightweight, maliciously secure verifiable function secret sharing. In *Eurocrypt*, 2022.
- [DG16] Zeev Dvir and Sivakanth Gopi. 2-server pir with subpolynomial communication. *J. ACM*, 63(4), 2016.
- [DHS14] Daniel Demmler, Amir Herzberg, and Thomas Schneider. Raid-pir: Practical multi-server pir. In *CCSW*, 2014.
- [Gas04] William I. Gasarch. A survey on private information retrieval. *Bulletin of the EATCS*, 82:72–107, 2004.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *ACM symposium on Theory of computing (STOC)*, 2009.
- [GI14] Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In *Advances in Cryptology – EUROCRYPT 2014*, pages 640–658, 2014.
- [GR05] Craig Gentry and Zulfikar Ramzan. Single-database private information retrieval with constant communication rate. In *ICALP*, 2005.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 75–92. Springer, 2013.
- [Hen16] Ryan Henry. Polynomial batch codes for efficient IT-PIR. *Proc. Priv. Enhancing Technol.*, 2016(4):202–218, 2016.
- [HH17] Syed Mahbub Hafiz and Ryan Henry. Querying for queries: Indexes of queries for efficient and expressive IT-PIR. In *CCS*, 2017.
- [HOWW19] Ariel Hamlin, Rafail Ostrovsky, Mor Weiss, and Daniel Wichs. Private anonymous data access. In *EUROCRYPT*, 2019.
- [IKOS04] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Batch codes and their applications. In *STOC*, 2004.
- [IKOS06] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography from anonymity. In *FOCS*, pages 239–248, 2006.

- [KCG21] Dmitry Kogan and Henry Corrigan-Gibbs. Private blocklist lookups with checklist. In *Usenix Security*, 2021.
- [KO97] E. Kushilevitz and R. Ostrovsky. Replication is not needed: single database, computationally-private information retrieval. In *FOCS*, 1997.
- [KW21] Sam Kim and David J. Wu. Watermarking cryptographic functionalities from standard lattice assumptions. *J. Cryptol.*, 34(3), jul 2021.
- [LG15] Wouter Lueks and Ian Goldberg. Sublinear scaling for multi-client private information retrieval. In *FC*, 2015.
- [Lip09] Helger Lipmaa. First CPIR protocol with data-dependent computation. In *ICISC*, 2009.
- [MCR21] Muhammad Haris Mughees, Hao Chen, and Ling Ren. Onionpir: Response efficient single-server pir. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, CCS '21*, page 2292–2306. Association for Computing Machinery, 2021.
- [OS07] Rafail Ostrovsky and William E. Skeith, III. A survey of single-database private information retrieval: techniques and applications. In *PKC*, pages 393–411, 2007.
- [PPY18] Sarvar Patel, Giuseppe Persiano, and Kevin Yeo. Private stateful information retrieval. In *CCS*, 2018.
- [PR93] P. Pudlák and V. Rödl. Modified ranks of tensors and the size of circuits. In *STOC*, 1993.
- [PS18] Chris Peikert and Sina Shiehian. Privately constraining and programming PRFs, the LWE way. In *Public Key Cryptography (2)*, volume 10770 of *Lecture Notes in Computer Science*, pages 675–701. Springer, 2018.
- [PY22] Giuseppe Persiano and Kevin Yeo. Limits of preprocessing for single-server PIR. In *SODA*, pages 2522–2548. SIAM, 2022.
- [SACM21] Elaine Shi, Waqar Aqeel, Balakrishnan Chandrasekaran, and Bruce Maggs. Puncturable pseudorandom sets and private information retrieval with near-optimal online bandwidth and time. In *CRYPTO*, 2021.
- [TDG16] Raphael R. Toledo, George Danezis, and Ian Goldberg. Lower-cost ϵ -private information retrieval. *PETS*, 2016.

A Smooth Tradeoff Between Space and Time

Throughout the paper, we focused on the special case when the client storage is $\tilde{O}_\lambda(\sqrt{n})$, and the server and client computation is also $\tilde{O}_\lambda(\sqrt{n})$ per query. Observe that the lower bound $S \cdot T \geq \Omega(n)$ by Corrigan-Gibbs et al. [CHK22] suggests a possible tradeoff between the client space S and the server/client computation per query T . Indeed, we can tune the parameters of our scheme to trade off the two parameters. The parameter choices are similar to Appendix A of Shi et al. [SACM21].

Suppose that we want the client’s storage to be $\tilde{O}_\lambda(f(n))$ for some function $f(n)$, and we want to guarantee $\tilde{O}_\lambda(n/f(n))$ server/client computation per query. Moreover, suppose that $f(n) \in$

$[\log^c n, \frac{n}{\log^c n}]$ for some suitable positive constant c . We can set the probability that any element $x \in \{0, 1, \dots, n-1\}$ is included in the set to be $\frac{1}{f(n)\log^2 n}$. This can be accomplished by applying the PRF to any suffix of $0^B || x$ of length at least $\log n - \log f(n) + 1$, and checking that the outcomes are all 1. We can set the $\text{lenT} = f(n)\log^3 n$ to make sure that Fact 3.8 still holds [SACM21]. As argued by Shi et al. [SACM21], the expected set enumeration time is now $O(\frac{n}{f(n)} \log n)$, and in the set enumeration algorithm, we can cap the number of calls to the PRF at $O(\frac{n}{f(n)} \cdot \log^5 n)$. Finally, we will set the batching parameter $Q = f(n)$.

With these parameters, the offline server time is $\tilde{O}_\lambda(n)$ and the offline client time and bandwidth are $\tilde{O}_\lambda(f(n))$. The online server and client time per query is $\tilde{O}_\lambda(n/f(n))$, and the per-query bandwidth is $\tilde{O}_\lambda(1)$. Since we need to perform the offline pre-processing every Q queries, we can amortize the cost of the offline phase over the Q queries. As a result, the amortized server and client time per query is $\tilde{O}_\lambda(n/f(n))$, and the per-query bandwidth is $\tilde{O}_\lambda(1)$.