

Two-server Distributed ORAM with Sublinear Computation and Constant Rounds

Ariel Hamlin¹ and Mayank Varia²

¹ Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA
hamlin.a@northeastern.edu

² Boston University, Boston, MA, USA varia@bu.edu

Abstract. Distributed ORAM (DORAM) is a multi-server variant of Oblivious RAM. Originally proposed to lower bandwidth, DORAM has recently been of great interest due to its applicability to secure computation in the RAM model, where circuit complexity and rounds of communication are equally important metrics of efficiency. All prior DORAM constructions either involve linear work per server (e.g., Floram) or logarithmic rounds of communication between servers (e.g., square root ORAM). In this work, we construct the first DORAM schemes in the 2-server, semi-honest setting that simultaneously achieve sublinear server computation and constant rounds of communication. We provide two constant-round constructions, one based on square root ORAM that has $O(\sqrt{N} \log N)$ local computation and another based on secure computation of a doubly efficient PIR that achieves local computation of $O(N^\epsilon)$ for any $\epsilon > 0$ but that allows the servers to distinguish between reads and writes. As a building block in the latter construction, we provide secure computation protocols for evaluation and interpolation of multivariate polynomials based on the Fast Fourier Transform, which may be of independent interest.

Keywords: Distributed Oblivious RAM, Square Root ORAM, Doubly Efficient PIR, Secure Multi-Party Computation, Fast Fourier Transform

1 Introduction

Oblivious RAM (ORAM) has been a vigorous area of study for the last three decades since it was introduced by Goldreich and Ostrovsky [17]. ORAM focuses on a client-server model where the server stores an outsourced database upon which the client wishes to execute a series of reads and writes. ORAM provides *privacy*, hiding the *contents* of the database, as well *obliviousness*, hiding the client's *access patterns*. In the traditional client-server model the client is assumed to be trusted. Recent efforts in the field have focused on lower bounds [35], optimal bandwidth [2, 31], and various different settings [15, 32].

Distributed ORAM (DORAM) is a variant of the basic client-server ORAM model in which there are *multiple* non-colluding servers. Data is duplicated across the servers and the client interacts with both as part of an access. The client again

remains the only trusted party. It was first introduced by Ostrovsky and Shoup [30], and later formally defined by Lu and Ostrovsky [26]. Lu and Ostrovsky were motivated by the desire to circumnavigate existing lower bounds in the single-server setting for bandwidth overhead, and their construction achieved $O(\log N)$ overhead by leveraging two non-communicating servers. Following their seminal paper there have been a number of works in the DORAM model that further reduce bandwidth [1, 7], reduce blocksize [25], or achieve practical efficiency [37].

Another advantage of the multi-server model of DORAM is its natural application to secure computation over databases in the RAM model. Traditional secure computation relies on a circuit representation that is at least linear in the size of the data over which it computes. This is prohibitive for any sublinear computation run on a database, such as binary search. Lu and Ostrovsky observe in [26] that the application of DORAM in this case is highly advantageous. The parties in the secure computation can simply emulate the DORAM client for any database access. In particular, they present a generic transformation from a 2-server DORAM scheme to a 2-party secure computation. It should be noted that works applying ORAM to secure computation are not limited to the DORAM setting, but also include adaptations of single server schemes. For example, there has been significant work on adapting tree-based ORAM schemes [33, 34] for secure computation. All of these DORAM constructions can be used in general-purpose secure computation like garbled RAM schemes [13, 14, 16, 27], or in special-purpose protocols like dynamic searchable encryption schemes [21].

There are two main approaches to constructing ORAM for secure computation: the first is to apply a generic MPC compiler, such as Garbled Circuits, to a ORAM or DORAM client [18, 19, 33, 34], and the second is to design a client specifically implemented by the two servers [5, 11, 23, 36]. Even if we start with an ORAM with our desired asymptotics (i.e. square-root ORAM [17]) applying a generic MPC compilers results in a server computation cost at least linear in the database size if we are to maintain constant rounds. There are a number of works that focus directly on the second model, which offers greater flexibility since the servers are typically afforded much more storage space than the client.

However, in both approaches, the multi-server setting introduces a new set of challenges apart from those found in the single-server ORAM setting. Wang et al. [33] also observe that the traditional efforts to optimize bandwidth overhead are ineffective in a setting where there are other controlling factors, such as the size of the circuit representation of the ORAM client. This is the motivation behind their Circuit ORAM construction, which focuses on optimizing circuit size. Doerner and Shelat [11] also show that in many cases, bandwidth is not the limiting factor but rather the latency between the two servers. This encouraged them to build a *constant round* DORAM for secure computation. Previous constructions had relied on recursive structures, which incurred a $O(\log N)$ rounds for each access, a prohibitive cost for latency dominated secure computation settings. Subsequent works in the constant round setting worked on improving on the $O(\sqrt{N})$ overhead of Floram, achieving $O(\log^3 N)$ overhead [23], or $O(\sqrt{N})$ in a black-box setting [5]. As with the original construction, these subsequent works

require *linear* local computation for each server. While for small N , latency costs may still dominate, for sufficiently large N this linear work is prohibitive. This leaves us with the following question:

Can we construct a 2-server Distributed ORAM for secure computation that achieves both constant round and sublinear server work?

In this work, we answer the above question in the affirmative.

1.1 Our contributions

We present the first DORAM constructions in the 2-server, semi-honest secure computation setting to achieve constant rounds *and* sublinear local computation on the servers.

- Our first sublinear DORAM construction achieves constant rounds and amortized local computation and bandwidth cost of $O(\sqrt{N} \log N)$ per access. It is based on square-root ORAM and has a modular build, allowing for subsequent improvements in the functionalities we rely on to be easily substituted.
- Our second sublinear DORAM construction is based on a secure computation of Doubly Efficient Private Information Retrieval (DEPIR) where the distinction between reads and writes is no longer hidden. In this setting, we achieve constant rounds with local computation and bandwidth of $O(N^\epsilon \cdot \text{poly}(\lambda))$ for any $\epsilon > 0$.

As an crucial building block toward the second construction, we present a secure two-party computation protocol for the Fast Fourier Transform (FFT) for multivariate polynomial evaluation and interpolation in quasilinear time and with only local computation; this may be of independent interest.

1.2 Technical Overview

In this section, we describe both of our DORAM constructions in more detail.

Sublinear DORAM. We start with describing the original square-root ORAM (introduced by Goldreich and Ostrovsky [17]) that our construction is based on. There is a single read-only array of size N , which we call the *store*, and a writable *stash* of \sqrt{N} size. Elements in the store are (address, value) pairs; at initialization, the elements are permuted with a permutation known only to the client, and all elements are encrypted. To perform a read at a particular address, the client checks the stash using a linear scan; if not present then it reads the permuted element from the read-only store, and if present then it is retrieved from the stash and a random ‘dummy’ element is read from the store instead. The newly-read element is placed in the stash, in order to maintain the invariant that each element is read only *once* from the store. In the case of a write, a dummy is read from the store and the element is written in the stash.

After enough queries have been made to fill the stash, a duration that we call an *epoch*, the elements from the stash are *reshuffled* back into the main store, with only the newest write at each location being kept.

While the basic square-root ORAM construction achieves constant rounds with sublinear communication and server computation, it is non-trivial to convert it to a two-party DORAM. There are two major issues incurred by shifting this to the two party case: (1) representing the permutation over the elements of the store and (2) merging the elements from the stash back into the store.

We first discuss how to represent the permutation that maps addresses to physical locations in the store. In [36], which is also based on square root ORAM, they choose to represent the permutation as a shared array in recursive ORAMs. This improves computation complexity but leads to $O(\log N)$ rounds of communication. To maintain constant rounds, we must instead find a compact representation of the permutation. We look for inspiration from the original square-root scheme. There, they generate a random ‘tag’ for each element in the store using a random oracle and then sort the elements according to the tag. A lookup then involves only a random oracle evaluation and a binary search across the sorted elements. However, because it is a single server scheme, they must use an oblivious sorting network in order to break the correlation between items in different epochs, which does not run efficiently in constant rounds. We leverage the fact that we have a two servers to break up the oblivious sort into its two components, ‘oblivious’ + ‘sort’. To prevent the server from mapping items between epochs, we use a simple constant round functionality to obliviously permute elements that allows each server to permute the elements in turn. As long as one server is honest, the data is permuted obliviously. This allows us to generate the tags using an oblivious pseudorandom function (OPRF), rather than a random oracle, on the newly obliviously permuted elements and then sort the tags locally. Lookup again is just an OPRF evaluation on the address shares and then a local binary search on the store.

The second challenge arises during the reshuffling phase of the protocol. In the original square-root ORAM, elements are simply moved back into their original locations (updated elements in the store, dummies back in the stash) by executing another oblivious shuffle. To solve this in constant rounds, we again exploit the ability to obliviously permute elements by using our two server architecture. In order to do that though, we must ensure that the elements that we are permuting do not contain any duplicates. For example, if a read was executed on index i , there would be two copies of element i , one in the stash and one in the store. To solve this issue, we note that the elements that have been read in the store is public knowledge to both servers. As long as we maintain the invariant *if an element has been read (or written to), it is in the stash, and each element only occurs in the stash once*, we can simply concatenate elements in the stash with the *unread* elements in the store at the end of an epoch. Once we have concatenated the elements we can obliviously permute them to get our new store. The stash can then just be filled with new dummy elements.

A more detailed discussion of our construction can be found in Section 3.

DORAM with Unlimited Reads. Thanks to the modularity of our base scheme, the components are easily extensible. In the second half of this work, we improve the performance of the read-only data store while keeping the rest of the construction (the stash, our periodic shuffling technique at the end of each epoch, etc) mostly intact.

The separation of our read-only store from a read-and-writable stash suggests an intriguing tradeoff: if we are willing to leak whether each operation is a read or a write operation, then it is beneficial to design an efficient read-only store that supports unlimited reads, and only pay for accessing the stash on (hopefully infrequent) write operations. This optimization allows us to increase the duration of each epoch, or in other words to amortize the cost of each shuffle over more reads. Concretely, in a scenario where the ratio of reads-to-writes is about N -to-1, then for any constant $\epsilon > 0$ we can construct a read-only store where whose amortized cost per query is just $O_\lambda(N^\epsilon)$. Here, the notation O_λ means that we suppress $\text{poly}(\lambda)$ terms in order to focus on the dependency on the database size. By reducing the size of the stash to $O_\lambda(N^\epsilon)$, we can support write operations with this performance as well.

Our strategy to construct a unlimited-reads store might seem counterintuitive at first: we start from a doubly efficient PIR [4, 6] that supports unlimited reads and convert it into a two-server distributed data store. A *doubly efficient private information retrieval* (DEPIR) scheme is a client-server protocol for oblivious access to a public dataset that only requires sublinear computation for both the client and server operations and constant rounds of communication between the two. At first glance, it may seem that a 1-server DEPIR is a strictly stronger primitive than a 2-server DORAM, so we might expect to construct the latter generically as a secure computation of the former. However, this intuition isn't true because there are three properties that we aim to satisfy with DORAM, but that (even a doubly efficient) PIR does not:

- Support for writes,
- Hiding the contents of the database, in addition to access patterns, and
- Ensuring that the secure computation is constant rounds when the two parties collectively emulate the (sublinear but not constant time) client, in addition to the client-server communication.

The main observation underlying this approach is that the SK-DEPIR protocol of Canetti et al. [6] is highly amenable to secure computation as operations mostly involve linear algebra in a finite field that can be done purely locally, plus bitstring and set operations that are easy to handle in constant rounds. SK-DEPIR constructions are based on a locally decodable code (LDC) in the style of a Reed-Muller code, which encodes a dataset as a multivariate polynomial. As a result, the most challenging part of our multiparty computation protocol involves securely emulating the client's procedures to evaluate or interpolate a multivariate polynomial at $O(N)$ points. The naive methods for polynomial evaluation (via application of the Vandermonde matrix) or polynomial interpolation (via the Lagrange interpolation polynomial) involve multiplication

of a public matrix by a secret-shared vector, which can be done non-interactively but requires $O(N^2)$ computation, which is too slow for our purposes.

Given a binary field $\mathbb{F} = \text{GF}(2^\ell)$ and a subspace $H^m \subset \mathbb{F}^m$, we construct secure computation protocols for evaluating or interpolating an m -variate polynomial $p \in \mathbb{F}[x_1, \dots, x_m]$ on all points in H^m in time that is quasilinear (rather than quadratic) in $|H^m|$. This protocol may be of independent interest, and in our protocol it is needed to achieve our goal of sublinear computation for the overall DORAM scheme. We construct this secure computation scheme in two stages: first we construct a secure computation protocol for the Additive Fast Fourier Transform protocol of Gao and Mateer [12] for univariate polynomials over a binary field, and then we bootstrap this protocol to handle multivariate polynomials by using recursion on the number of variables in the polynomial as previously shown by Kedlaya and Umans [24]. All operations in this protocol reduce to linear combinations of secret variables, so the entire computation can be done locally by each party on their own boolean secret shares without the need for any communication.

1.3 Related Work

We focus on schemes that are directly designed for secure computation. A direct comparison of their local computation, bandwidth, and number of rounds can be seen in Table 1. The construction of Zahur et al. [36] is very similar to our basic construction, but instead of implementing the permutation by OPRF evaluation, they use Waksman networks and a recursive position map. This allows for sublinear server work but that the cost of non-constant rounds. Doerner et al [11] uses function secret sharing to obtain a scheme with very good practical efficiency, but their need for linear server work limits scalability to large database lengths N . Gordon et al. [19] is in the more general DORAM model but uses PIR over tree-based ORAM. They are able to obtain $O(\log N)$ bandwidth but as with Doerner et al. they require linear local computation. Jarecki et al. focus on decreasing the round complexity and bandwidth of SC-DORAMs while still maintaining sublinear server computation. They are able to get the best combined set of parameters, but are still not able to achieve constant rounds of communication. Finally, Bunn et al. [5] achieve a 3-server DORAM scheme that achieves constant rounds and sublinear bandwidth, while providing a black-box construction. However, as with [11, 19] they require linear server work.

2 Preliminaries

In this section, we provide several definitions and constructions of existing cryptographic primitives that we leverage in this work. We begin with a brief summary of our notation.

Given a bitstring $x \in \{0, 1\}^\ell$, a 2-of-2 boolean secret sharing $\langle x \rangle$ denotes the uniform selection of two bitstrings x_1 for party 1 and x_2 for party 2 subject to the constraint that their boolean-xor $x_1 \oplus x_2 = x$. A binary field $\mathbb{F} = \text{GF}(2^\ell)$ is a

Table 1. Comparison of access in DORAM schemes for Secure Computation. Asterisks indicate schemes where the stash size is assumed to be $O(\log N)$ and $O(N^\epsilon)$, respectively, and the distinction between *read* and *write* is not hidden.

	No. Servers	Local Comp.	Bandwidth	Rounds
Zahur et al. [36]	2	$O(\sqrt{N \log^3 N})$	$O(\sqrt{N \log^3 N})$	$O(\log N)$
Floram [11]	2	$O(N)$	$O(\sqrt{N})$	$O(1)$
Floram* [11]	2	$O(N)$	$O(\log N)$	$O(1)$
Gordon et al. [19]	2	$O(N)$	$O(\log N)$	$O(1)$
Jarecki et al. [23]	3	$O(\log^3 N)$	$O(\log^3 N)$	$O(\log N)$
Bunn et al. [5]	3	$O(N)$	$O(\sqrt{N})$	$O(1)$
Sublinear DORAM	2	$O(\sqrt{N} \log N)$	$O(\sqrt{N} \log N)$	$O(1)$
Unlimited Reads DORAM*	2	$O_\lambda(N^\epsilon)$	$O_\lambda(N^\epsilon)$	$O(1)$

finite field of characteristic 2; there is a canonical bijection $\mathbb{F} \leftrightarrow \{0, 1\}^\ell$ such that field addition corresponds to boolean-xor. Hence, we overload the notation $\langle f \rangle$ so that it applies to field elements $f \in \mathbb{F}$. This secret sharing scheme commutes with linear algebra in the field, i.e., $\langle cf + c'f' \rangle = c\langle f \rangle + c'\langle f' \rangle$ can be computed locally by each server from public constants $c, c' \in \mathbb{F}$. and secret shares $\langle f \rangle, \langle f' \rangle$.

We use the convention of 0-indexing, with $[N] = \{0, 1, \dots, N-1\}$ as containing all whole numbers less than N . Additionally, $S \times S'$ denotes the Cartesian product of two sets. Bold letters \mathbf{v} denote vectors, subscripts \mathbf{v}_i indicate the i^{th} element of a vector, and $(w_i)_{i \in [N]}$ constructs a vector from an ordered list of items w_0, w_1, \dots, w_{N-1} . The notation $\|$ denotes concatenation of bitstrings, sets, or vectors into a single object of longer length containing the (ordered) union of all elements.

The notation $x \leftarrow \mathcal{D}$ indicates taking a sample from a probability distribution \mathcal{D} . By abuse of notation, $x \leftarrow S$ indicates sampling from the uniform distribution over set S ; we sometimes use $x \stackrel{\$}{\leftarrow} S$ for emphasis. We use \approx to indicate computational indistinguishability of two distributions; that is, $\mathcal{D} \approx \mathcal{D}'$ if no probabilistic polynomial time adversary \mathcal{A} has a noticeable difference in output when given a sample from \mathcal{D} or \mathcal{D}' .

2.1 Distributed Memory

First introduced by Bunn et al. [5], the ideal functionality \mathcal{F}_{mem} in Figure 1 captures the behavior achieved by a DORAM. The database is initialized on secret shares of the database, and subsequent accesses are also secret shared, as is their resulting output. This version of the definition deviates from the original in that the `Init` functionality returns shares of the database, and the

Fig. 1. Functionality \mathcal{F}_{mem}

1. On input of $(\text{Init}, \tilde{\text{DB}})$, set $\text{DB} = \tilde{\text{DB}}$, return random additive shares of DB^s to party s .
2. On input additive shares of $(\text{op}, \text{elem}, \text{DB})$ from two parties do:
 - (a) if $\text{op} = \text{read}$ then set $o = \text{DB}[\text{addr}]$
 - (b) if $\text{op} = \text{write}$ then set $o = \text{DB}[\text{addr}]$ and $\text{DB}[\text{addr}] = \text{val}$
 - (c) Let o^1, o^2 be random, additive shares of o , and DB^s be random additive shares of DB . Return (o^s, DB^s) to party s .

access protocol takes in those same shares. This syntactic difference is included only to make our own proofs cleaner and does not fundamentally change the definition. While Bunn et al. provide a viable 3-server construction that meets this ideal functionality and provides the necessary performance; we leverage the construction of Doerner et al. [11] that requires only 2-servers. From their construction, we obtain Lemma 1.

Lemma 1. *There exists a protocol Π_{DORAM} that implements the functionality \mathcal{F}_{mem} with the following complexity:*

- Access of $\text{op} = \text{read}$ or $\text{op} = \text{write}$ results in $O(1)$ rounds of communication, $O(n)$ local server computation, and $O(\sqrt{n})$ communication bandwidth.
- Initializing the functionality results in $O(1)$ rounds of communication, $O(n)$ local server computation, and $O(n)$ communication bandwidth.

2.2 Distributed Oblivious Pseudo-random Function

Distributed Oblivious Pseudo-random Function (DOPRF) achieves a distributed evaluation of a PRF between two parties. Typically one party hold the key, and the other the input, and only the second party learns the output. We require a variation of this ideal functionality, presented in Figure 2, in which both the key and the input are additively secret shared between two the two parties and both parties receive the output of the evaluation.

We introduce a construction in Figure 3 which meets our new ideal functionality that is effectively the semi-honest version of the DOPRF of Miao et al. [28], which itself is based on the work of Jarecki and Liu [22]. With only a small modification that allows the input and key to be secret-shared between the two servers. The construction leverages the Dodis-Yampolskiy pseudorandom function $F(k, x) = g^{1/(k+x)}$ [10], and it is secure under the q -Diffie Hellman inversion (q -DHI) assumption using a similar argument as in [28].

2.3 Constant-Round Equality Check

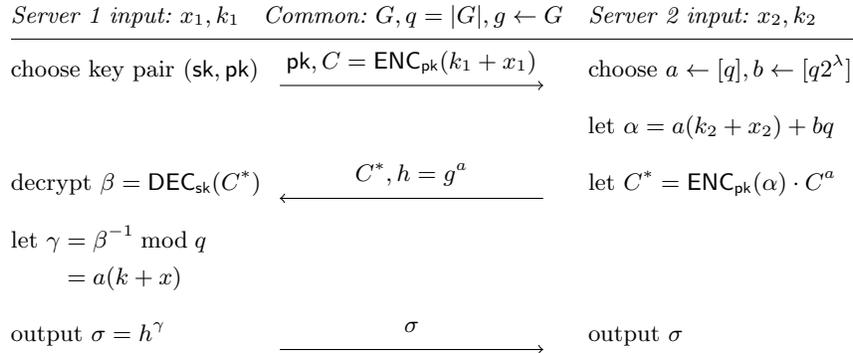
The functionality introduced in Figure 4 allows for two parties to check if the element for which they both hold shares is present in a database for which

Fig. 2. Functionality \mathcal{F}_{DOPRF}

The functionality is assumed to be initialized with PRF f .

1. Upon receiving (x_1, k_1) from party 1 and (x_2, k_2) from party 2, compute $\sigma = f_{k_1+k_2}(x_1 + x_2)$.
2. Returns σ to both party 1 and 2

Fig. 3. DOPRF Protocol, using ElGamal encryption (ENC, DEC)



they also hold shares. In particular it returns shares of a boolean b indicating the presence of a match, and if so the shares of that address. The database follows the invariant that there is only a single match within the database for the element. Both Damgard et al. and Nishide et al. [9, 29] construct solutions that achieve the computation³ with constant rounds.

2.4 Doubly Efficient Private Information Retrieval

First introduced by Canetti et al. and Boyle et al. [4, 6], Doubly Efficient Private Information Retrieval (DEPIR) is a variant of PIR achieving sub-linear server work by allowing pre-processing of the database. The major building block DEPIR is locally decodable codes (LDCs). Specifically, an application of Reed-Muller Codes, which allows for *smooth* LDCs.

Definition 1 (Smooth LDC). *A s -smooth, k -query locally decodable code with message length N , codeword size M , with alphabet Σ is denoted by $(s, k, N, M)_\Sigma$ -*

³ The exact functionality including the indicator bit is not included in their constructions, but they can be easily be extended with an additional round of a conditional computations.

Fig. 4. Functionality \mathcal{F}_{EQ-DB}

1. Upon receiving additive shares of $x \in \{0, 1\}^B$ and $DB \in (\{0, 1\}^B)^N$ from both parties 1 and 2, computes $DB_{\text{eq}} = \{x_i \stackrel{?}{=} x \mid x_i \in DB\}$.
2. Let $\mathbf{b} = \bigvee_{x_i \in DB_{\text{eq}}} x_i$ indicate if there was a match. If \mathbf{b} is non-zero, let addr^s be random, additive shares of addr such that $DB_{\text{eq}}[\text{addr}] = 1$ otherwise, let addr^s be random shares of zero. Return $(\mathbf{b}^s, \text{addr}^s)$ to party s .

smooth LDC and consists of a tuple of PPT algorithms (Enc, Query, Dec) with the following syntax:

- Enc takes a message $m \in \Sigma^N$ and outputs a codeword $c \in \Sigma^M$
- Query takes a index $i \in [N]$ and outputs a vector $\mathbf{x} = (x_1, \dots, x_k) \in [M]^N$
- Dec takes in vector codeword symbols $\mathbf{c} = (c_{x_1}, \dots, c_{x_k}) \in \Sigma^N$ and outputs a symbol $y \in \Sigma$

And has the following properties:

- **Local Decodability:** For all messages $m \in \Sigma^L$ and every index $i \in [N]$:

$$\Pr[\text{Dec}(\text{Enc}(m)_{\mathbf{x}}) = m_i : \mathbf{x} \leftarrow \text{Query}(i)] = 1$$

- **Smoothness:** For all indices $i \in [N]$, a LDC is s -smooth if when sampling $(x_1, \dots, x_k) \leftarrow \text{Query}(i)$, (x_1, \dots, x_k) is uniformly distributed on $[N]^s$ for every distinct subset of size s .

We now formally introduce DEPIR, in particular the secret key variant, called SK-DEPIR. Constructions rely on the *hidden permutation with noise* (HPN) assumption introduced by [6].⁴ While it is a new assumption, the validity of the class of permuted puzzles assumptions has been explored by Boyle et. al [3].

Assumption 1 (Hidden permutation with noise) Let $m < t < r < u < |\mathbb{F}|$ be functions of λ and N such that $|\mathbb{F}|^m = \text{poly}(\lambda)$ and $|\mathbb{F}|^t = \lambda^{\omega(1)}$. Define the distribution $\mathcal{D}(\pi, \text{addr}, T)$ that executes the Query protocol of \tilde{H}_{store} in the clear (without secret shares) to retrieve a set of vectors $\tilde{Y} = (\tilde{\mathbf{y}}_i)_{i \in [u]}$ and then outputs $Z = (\pi(\tilde{\mathbf{y}}_i))_{i \in [u]}$, when given a randomly-chosen permutation $\pi : \mathbb{F}^m \rightarrow \mathbb{F}^m$, integer $\text{addr} \in N$, and set $T \subset [u]$ as input. The hidden permutation with noise assumption states that the distribution $\mathcal{D}(\pi, \text{addr}, T)$ is computationally indistinguishable from the uniform distribution over $(\mathbb{F}^m)^u$.

Definition 2 (Doubly Efficient PIR). A Doubly Efficient PIR (DEPIR) for alphabet Σ consists of a tuple of PPT algorithms (KeyGen, Process, Query, Resp, Dec) with the following syntax:

⁴ A concurrent work by Boyle et. al [4] relies on an equivalent assumption called *Oblivious LDC*.

- **KeyGen** takes the security parameter 1^λ and outputs the key k
- **Process** takes a key k , database $\text{DB} \in \Sigma^N$ and outputs processed database $\tilde{\text{DB}}$
- **Query** takes a key k , database index $i \in [N]$ and outputs a query q and temporary state **State**
- **Resp** takes a query q , processed database $\tilde{\text{DB}}$ and outputs a server response c
- **Dec** takes a key k , server response c , temporary state **State** and outputs a database symbol $y \in \Sigma$

And has the following properties:

- **Correctness:** For all $\text{DB} \in \Sigma^N$ and $i \in [N]$:

$$\Pr \left[\begin{array}{l} k \leftarrow \text{KeyGen}(1^\lambda) \\ \tilde{\text{DB}} \leftarrow \text{Process}(k, \text{DB}) \\ (q, \text{State}) \leftarrow \text{Query}(k, i) \\ c \leftarrow \text{Resp}(\tilde{\text{DB}}, q) \end{array} : \text{Dec}(k, \text{State}, c) = \text{DB}_i \right] = 1$$

- **Double Efficiency:** The runtime of **KeyGen** is $\text{poly}(\lambda)$, the runtime of **Process** is $\text{poly}(N, \lambda)$, and the runtime of **Query**, **Dec** is $o(N) \cdot \text{poly}(\lambda)$, where N is the database size.
- **Security:** Any non-uniform PPT adversary \mathcal{A} has only $\text{negl}(\lambda)$ advantage in the following security game with a challenger \mathcal{C} :
 1. \mathcal{A} sends to \mathcal{C} a database $\text{DB} \in \Sigma^N$.
 2. \mathcal{C} picks a random bit $b \leftarrow \{0, 1\}$, and runs $k \leftarrow \text{KeyGen}(1^\lambda)$ to obtain a key k , and then runs $\tilde{\text{DB}} \leftarrow \text{Process}(k, \text{DB})$ to obtain a processed database $\tilde{\text{DB}}$, which it sends to \mathcal{A} .
 3. \mathcal{A} selects two addresses $i^0, i^1 \in [N]$, and sends (i^0, i^1) to \mathcal{C} .
 4. \mathcal{C} samples $(q, \text{State}) \leftarrow \text{Query}(k, i^b)$, and sends 1 to \mathcal{A} .
 5. Steps 3 and 4 are repeated an arbitrary (polynomial) number of times.
 6. \mathcal{A} outputs a bit b' , and his **advantage** in the game is defined to be $\Pr[b = b'] - \frac{1}{2}$.

As shown in [4, 6] we can achieve SK-DEPIR with sublinear or poly-log parameters. We will describe one such construction in Section 4.

Lemma 2. *There exists SK-DEPIR schemes with the following parameters, where N is the database size and λ is the security parameter:*

- **Sublinear SK-DEPIR:** For any $\epsilon > 0$, the running time of **Process** can be $N^{1+\epsilon} \cdot \text{poly}(\lambda)$, and the running time of **Query** and **Dec** can be $N^\epsilon \cdot \text{poly}(\lambda)$.
- **Polylog SK-DEPIR:** The running time of **Process** can be $\text{poly}(\lambda, N)$, and the running time of **Query** and **Dec** can be $\text{poly}(\lambda, \log N)$.

3 DORAM with Sublinear Computation

In this section we present our construction of \mathcal{F}_{mem} that achieves sublinear server work and communication with constant rounds.

3.1 Construction

In this section, we describe how we bootstrap from a linear-work \mathcal{F}_{mem} to a new protocol Π_{DORAM} that also instantiates \mathcal{F}_{mem} but with sublinear work and constant rounds, as desired. The overall architecture of the scheme can be found in Figure 5. We describe below our implementations of the store and stash.

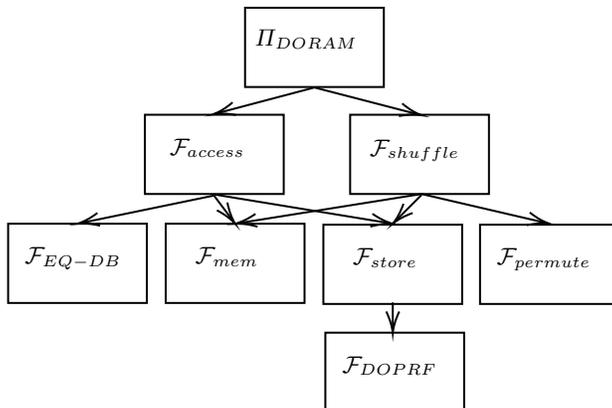
We implement the stash as an another two-party DORAM (matching the \mathcal{F}_{mem} functionality). We require a 2-party scheme with constant rounds, this can be instantiated by FLORAM [11] or Gordon et al. [19]. While they have linear server work for each access, because our stash is $t = \sqrt{N}$ records in size, this still results in sublinear server work within our protocol.

We implement the store in Figure 9 as a permuted array of elements sorted by PRF evaluation on the address of the element. Neither server knows the underlying permutation because it is created using our oblivious permutation protocol shown in Figure 8. We perform an \mathcal{F}_{DOPRF} evaluation across the shares of the addresses, which allows us to look up records in constant rounds by computing the OPRF based on the address of the element being searched for, and then each party performs a local binary search on their own store to find the shares of the element.

At the start of an epoch, the stash contains all the dummy elements and the store contains all the elements of the database concatenated with t dummy elements. The elements in the store are all permuted and indexed as above. Note that we consider dummy to be addressed from 1 to t , so valid elements are indexed started at t . We also have (in the clear) a counter, starting at 1. The access logic is encapsulated within our access protocol in Figure 7 and proceeds as follows. When we want to do a **read**, we check if the element is in the stash by calling \mathcal{F}_{EQ-DB} , which returns a secret shared boolean \mathbf{b} indicating if the element is present, as well as the shares of the address to each party if it is present. We then use \mathbf{b} as a selector bit in a shared conditional computation to see if we read the element (if the element is present in the stash) or the next dummy element (addressed at the counter) in the stash. Then we read an element from the store, using \mathcal{F}_{store} , again based on the selector bit. If the element is in the stash, we read the next dummy element at address counter, if it is not in the stash, we read the element itself. Finally, we write an element back to the stash, using \mathcal{F}_{mem} , either the dummy element we read (which is just overwriting the same element) if the element was in the stash, or the element read from the stash. The element is written back at the ‘counter’ location in the stash. The protocol then returns random additive shares of the element being read. If the operation is instead was a **write** the only variation in the above process is in the final step writing elements back to the stash, rather than writing to the ‘counter’ location automatically, if the element was previously in the stash, it is overwritten at that location.

At the end of an epoch (when the counter reaches t), the overarching Π_{DORAM} invokes Π_{shuffe} in Figure 10, which resets the state as mentioned above. In the original square-root ORAM scheme, removing duplicates required a costly oblivious sort operation, which is not constant round. By contrast, we achieve

Fig. 5. The overall architecture of the ideal functionalities used in the Π_{DORAM} construction.



a constant-round reshuffling algorithm by leveraging the following invariant of \mathcal{F}_{access} : *if an element has been read (or written to) it is in the stash, and each element only occurs in the stash once*. This invariant allows us to simply note which elements in the store have been read during the epoch and eliminate them, knowing that their most recent copy is represented in the stash. This claim applies to dummy elements as well: during the shuffle operation, we only need to insert new dummy elements to replace those that have been overwritten in the stash by real writes. Once the unread elements and the current stash have been permuted obliviously by $\mathcal{F}_{permute}$, the stash is reinitialized with the dummy values and counter is reset to 1. We also note that we leverage $\mathcal{F}_{shuffle}$ when we first initialize the DORAM. We call $\mathcal{F}_{shuffle}$ on the original shares of the secret shared database, concatenated with the necessary dummy elements. The set of read elements is empty as is the stash, resulting in a permutation of the original database and dummies after $\mathcal{F}_{permute}$ is called.

Our oblivious permutation protocol in Figure 8 does two things: it rerandomizes the shares held by each server and applies the same random permutation to each server's shares. Beginning with a vector of secret shares $\langle \mathbf{M} \rangle$, server 2 begins $\Pi_{permute}$ by encrypting her own shares \mathbf{M}^2 using an additively homomorphic encryption scheme and sending the result to server 1. Next, server 1 applies the same randomly-chosen permutation to her own shares \mathbf{M}^1 as well as the ciphertexts from server 2, and she then rerandomizes each pair of shares by adding a random value to her own share and subtracting the same value (homomorphically) from server 2's share. She sends encrypted versions of both shares to server 2, who performs the same permute-and-rerandomize operation and sends the result to server 1 to complete the constant-round oblivious permutation.

Fig. 6. Π_{DORAM} Protocol

On (Init, $\langle DB \rangle$):

1. Server 1 computes $V = \{(i, 0) \mid i \in [t]\}$
2. Each server calls $\mathcal{F}_{shuffle}$ on (shuffle, l , State) where $l = \{\}$ and State = $\{DB \parallel V, 0, \{\}\}$. Server s receives (S^s, M^s, k_s) as output.
3. Let $ctr = 1$ and $l = \{\}$. For Server s its current state is State ^{s} = (ctr, M^s, S^s, k_s) , returns State ^{s}

On additive shares of (op, elem, State):

1. If $ctr = t$, call $\mathcal{F}_{shuffle}$ on (shuffle, l , State) and update state for Server s with (S'^s, M'^s, k'^s) . Set $ctr = 1$ and $l = \{\}$.
2. Call \mathcal{F}_{access} on additive shares $(op_i, elem_i, State)$, recovering $(i_M, \langle elem \rangle, \langle State \rangle)$ and update state for Server s with S'^s . Set $l = l \cup i_M$ and $ctr = ctr + 1$
3. Return $(elem^s, State^s)$ to Server s .

3.2 Complexity Analysis

Now consider the asymptotic complexity of our scheme. We first evaluate the complexity of the underlying protocols, and then compute the amortized complexity of the overall Π_{DORAM} protocol. The overall complexity when $t = \sqrt{N}$ is shown in Table 2.

- $\Pi_{permute}$: Each server must perform $O(N + t)$ encryption, decryption and other local operations. The entire encrypted store is sent, again resulting in $O(N + t)$ bandwidth. The protocol runs in 3 rounds, or $O(1)$.
- Π_{store} : Here we consider two separate costs, one for initialization, and one for performing an access. During initialization, local computation is dominated by the sorting across the OPRF outputs, $O((N + t) \log(N + t))$, and bandwidth by the OPRF computation itself, $O(N + t)$. We obtain constant rounds in initialization by executing all of the OPRF evaluations in parallel. On access, local computation is dominated by searching for the tag, $O(\log(N + t))$, and the only round of interaction and bandwidth is the OPRF evaluation.
- Π_{access} : Finding the element in the stash only takes local computation and bandwidth linear in the stash size and constant rounds. The two other operations of cost are accessing stash and the store, each of which take $O(t)$ and $O(\log(N + t))$ local computation and $O(\sqrt{t})$ and $O(1)$ bandwidth respectively. This leaves access dominated by finding the element in the stash, $O(t)$ local computation and bandwidth⁵.

⁵ For any value of $t < \log(N + t)$ then the cost of Π_{store} controls, but in our setting we consider a t greater than that

Fig. 7. Π_{access} Protocol

1. On additive shares of $(op, elem_{in}, State)$, let $elem_{in} = (addr_{in}, val_{in})$ and $State = (ctr, M^s, k_s, S^s)$, where S is an array of $elem$.
2. Find element in stash or read next dummy address:
 - (a) Compute additive shares of index i by calling \mathcal{F}_{EQ-DB} in Figure 4 on additive shares of $(addr_{in}, S)$, receiving random additive shares (b, i) as output.
 - (b) Jointly compute random additive shares of i_s such that:

$$i_s = \begin{cases} i & b = 1, \text{ element in stash.} \\ ctr & b = 0, \text{ element not in stash.} \end{cases}$$

- (c) Then recover $elem_s$ by calling \mathcal{F}_{mem} on secret shares of $(read, (i_s, 0), S)$.
3. Look up either the next dummy element or the original element in the store:
 - (a) Jointly compute random additive shares of $addr_M$:

$$addr_M = \begin{cases} ctr & b = 1, \text{ element in stash.} \\ addr_{in} & b = 0, \text{ element not in stash.} \end{cases}$$

- (b) Call \mathcal{F}_{store} on the additive shares of $(read, addr_M, M, k_s)$, recovering $(i_M, \langle elem_M \rangle)$.
4. Write the read $elem_M$ or input $elem_{in}$ back to stash:
 - (a) If $op = read$, jointly compute random additive shares of:

$$(i_W, elem_W, elem) = \begin{cases} (ctr, elem_M, elem_s) & b = 1, \text{ element in stash.} \\ (ctr, elem_M, elem_M) & b = 0, \text{ element not in stash.} \end{cases}$$

- (b) If $op = write$, jointly compute random additive shares of:

$$(i_W, elem_W, elem) = \begin{cases} (i_s, elem_{in}, elem_{in}) & b = 1, \text{ element in stash.} \\ (ctr, elem_{in}, elem_{in}) & b = 0, \text{ element not in stash.} \end{cases}$$

- (c) Call \mathcal{F}_{mem} on additive shares of $(write, (i_W, elem_W), S)^a$.
5. Server s returns $(i_M, \langle elem \rangle, \langle State \rangle)$.

^a Any functionality that returns an updated share of S or M is assumed to update the held state $State$, but is elided for notational simplicity.

Fig. 8. $\Pi_{permute}$ Protocol

On input (permute, $\langle M \rangle$):

1. Each server runs $\text{pk}_s, \text{sk}_s \leftarrow \text{KeyGen}(1^\lambda)$ and sends pk_s to the other server.
2. For $s \in \{1, 2\}$, and $s' = 3 - s$
 - (a) Server s' encrypts their additive shares of the values $\mathbf{C}^{s'} = \{\text{ENC}_{\text{pk}_{s'}}(\text{elem}_i^{s'}) \mid \text{elem}_i^{s'} \in M^{s'}\}$ and sends $\mathbf{C}^{s'}$ to Server s .
 - (b) Server s chooses vector of random values $\{r_i \in \{0, 1\}^B\}_{i \in [N]}$, and a random permutation π and updates locally $\{\text{elem}_{\pi(i)}^s + r_i^s \mid \text{elem}_i^s \in M^s\}$. It then computes permutes and re-randomizes s' encrypted shares: $\mathbf{C}_r^{s'} = \{c_{\pi(i)s}^{s'} \cdot \text{ENC}_{\text{pk}_{s'}}(-r_i^s) \mid c_i^{s'} \in \mathbf{C}^{s'}\}$ and sends $\mathbf{C}_r^{s'}$ to Server s' .
 - (c) Server s' decrypts $\mathbf{C}_r^{s'}$ to get $M^{s'} = \{\text{elem}_{\pi(i)}^{s'} - r_i^s \mid i \in [N]\}$.
3. Return (M^s) to Server s .

Fig. 9. Π_{store} Protocol

On input (Init, $\langle M \rangle$):

1. Choose the new random PRF keys for k_1 and k_2 .
2. Server 1 and 2 call on \mathcal{F}_{DOPRF} on inputs (k_1, addr_i^1) and (k_2, addr_i^2) respectively for all $(\text{addr}_i, \text{val}_i) \in M$ in parallel. Let $\sigma_i = f_{k_1+k_2}(\text{addr}_i^1 + \text{addr}_i^2)$, and $\Sigma = \{\sigma_i \mid i \in [N]\}$. Both servers sort $M^{s'} = \{\sigma_i, \text{elem}_i\}_{i \in [N]}$ in lexicographic order by σ .
3. Return $(\Sigma, M^{s'}, k_s)$

On input (read, $\langle \text{addr}_i \rangle, \langle k \rangle, \langle M \rangle$):

1. Server 1 and Server 2 engage in \mathcal{F}_{DOPRF} on inputs (k_1, addr_i^1) and (k_2, addr_i^2) respectively. Both servers obtain the output $\text{addr} = f_{k_1+k_2}(\text{addr}_i^1 + \text{addr}_i^2)$.
2. Each Server s performs a local binary search in M^s for addr and recover its index i and additive shares of the element elem_i . Each server returns (i, elem_i^s) .

Fig. 10. $\Pi_{Shuffle}$ Protocol

On input (shuffle, State, l):

1. Let $\text{State} = (M^s, k_s, S^s)$.
2. Let M_r^s be all unread elements in M^1 and M^2 , i.e. $M_r \not\subseteq I$. Set $R^s = M_r^s \parallel S^s$.
3. Let $V = \{(i, 0) \mid i \in [t]\}$ each server calls \mathcal{F}_{mem} on additive shared input (Init, V). Server s receives S^s as output.
4. Servers 1 and 2 call $\mathcal{F}_{permute}$ on (permute, $\langle R \rangle$). Server s receives $(M^{s'})$ as output, which in turn it calls \mathcal{F}_{store} on (Init, $M^{s'})$ and receives (M^s, k_s) as output.
5. Server s returns (S^s, M^s, k_s) .

Table 2. A evaluation of each of the protocol’s server computation, bandwidth and rounds of communication where $t = \sqrt{N}$. Note that the numbers for Π_{mem} are taken from Lemma 1 and Π_{DORAM} has been amortized where appropriate.

	Local Computation	Bandwidth	Rounds
$\Pi_{DORAM}(\text{Init})$	$O(N \log N)$	$O(N \log N)$	$O(1)$
$\Pi_{DORAM}(\text{op})$	$O(\sqrt{N} \log N)$	$O(\sqrt{N} \log N)$	$O(1)$
Π_{access}	$O(\sqrt{N})$	$O(\sqrt{N})$	$O(1)$
$\Pi_{shuffle}$	$O(N \log N)$	$O(N \log N)$	$O(1)$
$\Pi_{mem}(\text{Init})$	$O(\sqrt{N})$	$O(\sqrt{N})$	$O(1)$
$\Pi_{mem}(\text{read})$	$O(\sqrt{N})$	$O(\sqrt[4]{N})$	$O(1)$
$\Pi_{store}(\text{Init})$	$O(N \log N)$	$O(N)$	$O(1)$
$\Pi_{store}(\text{read})$	$O(\log N)$	$O(1)$	$O(1)$
$\Pi_{permute}$	$O(N)$	$O(N)$	$O(1)$

- $\Pi_{shuffle}$: Shuffle is dominated by the initialization of the store, inheriting the performance and bandwidth complexity directly from Π_{store} .

We now consider the amortized complexity of the overall local computation of Π_{DORAM} during access. We consider the cost of shuffling averaged over an epoch of t accesses. The cost of accessing a single block, represented by Π_{access} , is $O(t)$. The cost of shuffle is $O((N + t) \log(N + t))$. We can consider the total cost of local computation during an epoch as:

$$D_{LC}(N, t) = t(t) + (N + t) \log(N + t)$$

Averaging over t -accesses we get:

$$D_{LC}(N, t) = t + \frac{N}{t} \log(N + t) + \log(N + t)$$

If we set $t = \sqrt{N}$ we get $D_{LC}(n) = O(\sqrt{N} \log N)$. For bandwidth, we do a similar computation and get $D_B(n) = O(\sqrt{N} \log N)$.

3.3 Security

In this section we provide the overall security statement and ideal functionalities. We refer the reader to the full paper [20] for the proof.

Notation and Valid Inputs.. We define a set of notations and valid inputs for our various protocols used in the following proofs. We assume $\text{op} \in \{0, 1\}$

where $\text{op} = 0$ represents a **read** operation and $\text{op} = 1$ is **write**. Valid elements are a tuple of an address and a value where $\text{addr} \in [N]$ and $\text{val} \in \{0, 1\}^B$. The input database DB is made up of N valid elements. The store M is represented as key-value store, where the keys are the output of PRF f with key k and the values consist of valid elements. The stash S is an array of t valid elements. We define the set of valid inputs for an DORAM of N elements of block size B and t dummies as $\text{Dom}_{N,B,t}$.

Theorem 1. Π_{DORAM} (Figure 6) implements functionality \mathcal{F}_{mem} and for each party there exists a PPT simulator for each Server $s \in \{1, 2\}$ Sim_D^s such that:

$$\left\langle \text{input}_{\mathcal{A}}, \text{output}^{\Pi_D}, \text{view}_{\mathcal{A}}^{\Pi_D} \right\rangle_{\text{input} \in \text{Dom}_{N,B,t}} \approx \left\langle \text{input}_{\mathcal{A}}, \mathcal{F}_{\text{mem}}(\text{input}), \text{Sim}_D^s(\text{input}_{\mathcal{A}}, \mathcal{F}_{\text{mem}}(\text{input})_{\mathcal{A}}) \right\rangle_{\text{input} \in \text{Dom}_{N,B,t}}$$

where $\text{input} = \{(\text{Init}, \text{DB}), (\text{op}_i, \text{elem}_i, \text{ctr}, \text{M}, \text{S}, k)\}$, and $\text{output} = \{(\text{ctr}, \text{M}, \text{S}, k), (\text{elem}, \text{ctr}, \text{M}, \text{S}, k)\}$.

Proof (Theorem 1). See full paper [20] for proof.

Fig. 11. Functionality $\mathcal{F}_{\text{access}}$

On additive shares of $(\text{op}, \text{elem}_{\text{in}}, \text{State})$ where $\text{State} = (\text{ctr}, \text{M}, k, \text{S})$ and $\text{elem}_{\text{in}} = (\text{addr}_{\text{in}}, \text{val}_{\text{in}})$ set $i_S, \text{addr}_M, i_W, \text{elem}_W$ and elem during the protocol according to the below table (as defined by op and if addr_e is found in the stash):

op	$\text{addr}_{\text{in}} \in \text{S}$	i_R	addr_M	i_W	elem_W	elem
read	yes	addr_{in}	ctr	ctr	elem_M	elem_S
read	no	ctr	addr_{in}	ctr	elem_M	elem_M
write	yes	addr_{in}	ctr	addr_{in}	elem_{in}	elem_{in}
write	no	ctr	addr_{in}	ctr	elem_{in}	elem_{in}

1. Recover elem_S by calling \mathcal{F}_{mem} on secret shares of $(\text{read}, (i_R, 0), \text{S})$.
2. Call $\mathcal{F}_{\text{store}}$ on the additive shares of $(\text{read}, \text{addr}_M, \text{M}, k_s)$, recovering $(i_M, (\text{elem}_M))$.
3. Call \mathcal{F}_{mem} on additive shares of $(\text{write}, (i_W, \text{elem}_W), \text{S})^a$.
4. Return $(i_M, (\text{elem}), (\text{State}))$

^a Any functionality that returns an updated share of S or M is assumed to update the held state State , but is elided for notational simplicity.

Fig. 12. Functionality $\mathcal{F}_{permute}$

On input of (permute, $\langle \tilde{M} \rangle$):

1. Choose random permutation π and set $M = \{\tilde{M}_{\pi(i)} \mid i \in [N]\}$.
2. Let M^s be a random additive share of M , and return M^s to Server s .

Fig. 13. Functionality \mathcal{F}_{store}

1. On input of (Init, $\langle M \rangle$): Choose PRF key k and for all $\text{elem}_i \in M$ compute $\sigma_i = f_k(\text{addr}_i)$ and let $\Sigma = \{\sigma_i \mid i \in [N]\}$. Set $M = \{(\sigma_i, \text{elem}_i) \mid \text{elem}_i \in M\}$ and sort M in lexicographic order by σ . Let elem_i^s and k_s be random additive shares of elem_i and k respectively, and $M^s = \{(\sigma_i, \text{elem}_i^s) \mid \text{elem}_i \in M\}$. Return (Σ, M^s, k_s) to Server s .
2. On input additive shares of (read, addr_i, k, M) from two parties return additive shares of $M[i_M]$ where $\sigma_{i_M} = f_k(\text{addr}_i)$ and i_M to each server.

Fig. 14. Functionality $\mathcal{F}_{shuffle}$

On input (shuffle, State, I):

1. Let State = (M, k, S) .
2. Let M_r be all unread elements in M , i.e. $M_r \not\subseteq I$. Set $R = M_r \parallel S$.
3. Let $V = \{(i, 0) \mid i \in [t]\}$ and call \mathcal{F}_{mem} on additive shared input (Init, V), receiving S' as output.
4. Call $\mathcal{F}_{permute}$ on (permute, $\langle R \rangle$) receiving M' as output, which in turn is passed into \mathcal{F}_{store} as (Init, M'). Finally, (M'', k') is received as output.
5. Server s returns random additive shares (S'^s, M''^s, k'^s) .

4 Sublinear DORAM with Unlimited Reads

In this section, we introduce an alternative DORAM construction $\tilde{\Pi}_{DORAM}$ that also implements the \mathcal{F}_{mem} functionality with constant rounds, sublinear server work, and sublinear communication. This protocol differs from the construction in Section 3 in that it does not attempt to hide whether a query is a read or write operation, and in exchange it achieves better performance.

The construction in this section only needs to invoke the shuffle functionality $\mathcal{F}_{shuffle}$ after t write operations, independent of the number of read operations. In scenarios where writes are infrequent, the amortized cost per read can have a small $O_\lambda(N^\epsilon)$ dependency on the database size N for any constant $\epsilon > 0$. To build the new DORAM protocol $\tilde{\Pi}_{DORAM}$, we start from a SK-DEPIR that supports unlimited reads while hiding access patterns, and we emulate the server using secure 2-party computation (which hides the database contents as well).

We first describe how we instantiate the new version of store that relies on SK-DEPIR in Section 4.1, then show how to use the new \mathcal{F}_{store} in the larger $\tilde{\Pi}_{DORAM}$ protocol in Section 4.2. Finally we show how to construct secure computation of multivariate polynomial evaluation and interpolation using FFT in Section 4.3.

4.1 Instantiating \mathcal{F}_{store} using secure computation of SK-DEPIR

In this section, we show that a secure 2-party computation (2PC) of the Canetti et al. construction leads to an instantiation $\tilde{\Pi}_{store}$ of the \mathcal{F}_{store} functionality. The construction we present in this section achieves sublinear communication with a constant number of rounds and quasilinear server work. To do so, we first construct a 2PC protocol for a locally decodable code, and then we construct $\tilde{\Pi}_{store}$ as a 2PC of a secret key doubly efficient private information retrieval (SK-DEPIR) protocol based on an LDC.

We focus on a block size $B = \ell$, so that each block can canonically be encoded as a field element in $\mathbb{F} = \text{GF}(2^\ell)$. Put another way, all references to the database size N are enumerated in terms of the number of blocks, but if one desires a lower block length like $B = 1$ then N should instead be interpreted in terms of the number of bits of the database.

2PC for a Locally Decodable Code. First, we construct a secure 2-party computation protocol $\tilde{\Pi}_{ldc}$ of the locally decodable code used by Canetti et al. [6], which is a Reed-Muller-based polynomial code. We depict our construction in Figure 15, in which the two parties maintain boolean secret shares of all input, intermediate, and output data from the LDC of Canetti et al.

Our 2PC protocol $\tilde{\Pi}_{ldc}$ operates over a binary field $\mathbb{F} = \mathbb{F}_2[z]/(\rho(z))$ of size $|\mathbb{F}| = 2^\ell$ defined using an irreducible polynomial ρ of degree ℓ . Elements of \mathbb{F} can be represented using bitstrings of length ℓ in the canonical way, such that the addition of two elements corresponds to the boolean-xor of their bitstring values. Furthermore, we consider $H \subset \mathbb{F}$ to be the subspace of \mathbb{F} of size $|H| = 2^h$ containing the span of basis elements $\mathcal{H} = \{z^{h-1}, z^{h-2}, \dots, z, 1\}$; this corresponds

Fig. 15. $\tilde{\Pi}_{ldc}$ protocol for secure 2-party computation of a Reed-Muller-style LDC

Given a binary field \mathbb{F} , a subspace $H \subset \mathbb{F}$, and an integer m , let $N = |H|^m$ and $M = |\mathbb{F}|^m$. Let $(x_i)_{i \in [r]}$ be arbitrary, distinct non-zero constants. Finally, define the bijections $\iota : [N] \rightarrow H^m$ and $\delta : B \rightarrow \mathbb{F}$ via lexicographic ordering.

On input (Enc, $\langle D \rangle$) for database $D \in B^N$:

1. For all $i \in [N]$, compute shares of $\mathbf{c}_i = \iota(i) \in H^m$ and $d_i = \delta(D[i]) \in \mathbb{F}$.
2. Securely interpolate polynomial $\psi : \mathbb{F}^m \rightarrow \mathbb{F}$ of degree t from $\{(\mathbf{c}_i, d_i)\}_{i \in [N]}$.
3. Securely evaluate $\mathbf{E} = (\mathbf{v}, \psi(\mathbf{v}))_{\mathbf{v} \in \mathbb{F}^m}$, the truth table of ψ on \mathbb{F}^m . Output $\langle \mathbf{E} \rangle$.

On input (Query, $\langle q \rangle$) for an index $q \in [N]$:

1. Randomly choose a degree- t polynomial $\phi : \mathbb{F} \rightarrow \mathbb{F}^m$ such that $\phi(0) = \iota(q)$.
2. Securely evaluate $\mathbf{y}_i = \phi(x_i)$ for all $i \in [r]$, and output shares $\langle (\mathbf{y}_i)_{i \in [r]} \rangle$.

On input (Dec, $\langle (a_i)_{i \in [r]} \rangle$):

1. Securely interpolate polynomial $\tilde{\phi} : \mathbb{F} \rightarrow \mathbb{F}$ of degree $r - 1$ with $\tilde{\phi}(x_i) = a_i \forall i$.
2. Output shares $\langle \delta^{-1}(\tilde{\phi}(0)) \rangle$ of the block corresponding to field element $\phi(0)$.

to bitstrings that have $\ell - h$ leading 0s. Also, protocol $\tilde{\Pi}_{ldc}$ performs operations in the vector spaces H^m and \mathbb{F}^m of sizes $N = |H|^m$ and $M = |\mathbb{F}|^m$, respectively.

We claim that this protocol can be securely evaluated efficiently and non-interactively. Throughout this section, we only consider boolean secret shares $\langle \cdot \rangle$, so that field addition and scalar multiplication can be performed locally by each server, without interaction. Hence, our claim amounts to the statement that all operations in $\tilde{\Pi}_{ldc}$ involve only linear algebra in the field along with concatenation/truncation of bitstrings, because all of these operations commute with boolean-xor.

Theorem 2. *Let $m < t < r < N < M$ be parameters of a Reed-Muller locally decodable code such that N and M are powers of 2. Then, protocol $\tilde{\Pi}_{ldc}$ in Fig. 15 is a secure two-party computation of an LDC with local decodability and smoothness. Furthermore, $\tilde{\Pi}_{ldc}$ requires no interaction between parties, and its computation cost is $O(M \log^2(M))$ for Enc and $O(r^2)$ for Query and Dec.*

Proof. Our 2PC protocol $\tilde{\Pi}_{ldc}$ contains methods for the servers to securely compute each of the 3 methods of an LDC on boolean secret-shared data. Ergo, the local decodability and smoothness of $\tilde{\Pi}_{ldc}$ follow immediately from the same properties of its non-secure-computation counterpart [6].

There are four types of operations used throughout $\tilde{\Pi}_{ldc}$, and we show below how to compute all of them non-interactively. The first two operations are used in Enc, and the last two in Query and Dec.

- Computing the lexicographic maps δ and ι : δ is the identity operation on bitstrings, and thanks to the specific basis we chose for H , computing $\iota(i)$

Fig. 16. $\tilde{\Pi}_{store}$ protocol, based on secure 2PC of the SK-DEPIR scheme of Canetti et al., given any integers $m < t < r < u < N < M$ satisfying the HPN assumption.

On input $(\text{Init}, \langle \mathbf{M} \rangle)$, run the following steps of the SK-DEPIR:

1. **KeyGen:** Randomly choose a subset $T \subset [u]$ of size r .
2. **Process:** Run $\tilde{\Pi}_{ldc}$ on input $(\text{Enc}, \langle \mathbf{M} \rangle)$ to obtain shares of encoded database $\langle \mathbf{E} \rangle$. Run u instances of $\Pi_{permute}$ on $\langle \mathbf{E} \rangle$ to form permuted $\{\langle \mathbf{E}^i \rangle\}_{i \in [u]}$. Run Π_{store} on input $(\text{Init}, \cup_{i \in [u], j \in [M]} (i \parallel j, \langle \mathbf{E}^i[j] \rangle))$ to obtain $(\Sigma, \langle \mathbf{M} \rangle, \langle \mathbf{k} \rangle)$.
3. **Output** $(\Sigma, \langle \mathbf{M} \rangle \cup \langle T \rangle, \langle \mathbf{k} \rangle)$, the state from KeyGen and Process.

On input $(\text{read}, \langle \text{addr} \rangle, \langle \mathbf{k} \rangle, \langle \mathbf{M} \rangle \cup \langle T \rangle)$, run the following steps of the SK-DEPIR:

1. **Query:** Run $\tilde{\Pi}_{ldc}$ on input $(\text{Query}, \langle \text{addr} \rangle)$ to obtain shares of r elements $Y = ((\langle \mathbf{y}_i \rangle)_{i \in [r]})$. Construct a longer vector $\tilde{Y} = ((\langle \tilde{\mathbf{y}}_i \rangle)_{i \in [u]})$ such that $\tilde{Y}|_T = Y$ and the remaining elements $\{\langle \tilde{\mathbf{y}}_i \rangle \mid i \in [u] \setminus T\}$ are chosen uniformly at random.
2. **Resp:** For $i \in [u]$, run Π_{store} on $(\text{read}, (i \parallel \langle \tilde{\mathbf{y}}_i \rangle), \langle \mathbf{k} \rangle, \langle \mathbf{M} \rangle)$. Construct a list $L = ((\langle \text{elem}^i \rangle)_{i \in [u]})$ of the shares of elements returned in response.
3. **Dec:** Truncate the list $\langle L|_T \rangle$ to responses of queries in Y . Run $\tilde{\Pi}_{ldc}$ on input $(\text{Dec}, \langle L|_T \rangle)$ to obtain shares of a field element $\langle \text{val} \rangle$. Output $(\langle \text{addr} \rangle, \langle \text{val} \rangle)$.

- merely involves partitioning the bits of $i \in [N]$ into m strings of length h , padding with 0s in the $\ell - h$ leftmost bits. These string operations can be performed independently in $O(N)$ time on each boolean secret share of i .
- Interpolation and evaluation of multivariate polynomial ψ : this task is challenging; we show in Section 4.3 a non-interactive secure 2-party protocol that performs these operations across all of \mathbb{F}^m in time $O(M \log^2(M))$.
 - Random sampling of multivariate polynomial ϕ : the parties already hold shares of the constant term $\phi_0 = \phi(0)$, and they can randomly choose all other t coefficients in $O(t)$ time.
 - Evaluation of ϕ at r points and interpolation of $\tilde{\phi}$ from r points: since the evaluation points $(x_i)_{i \in [r]}$ are publicly known, the coefficients for polynomial evaluation and Lagrange interpolation can also be publicly (pre-)computed. Ergo, evaluating or interpolating a polynomial of degree $\leq r$ only involves linear algebra and takes $O(r^2)$ time.

Constructing $\tilde{\Pi}_{store}$ as a secure computation of SK-DEPIR. Next, we construct a new protocol $\tilde{\Pi}_{store}$ that also instantiates \mathcal{F}_{store} . It is a secure two-party computation of the client-server SK-DEPIR protocol of Canetti et al. [6] in which the two parties jointly emulate the server. In Fig. 16, we show simultaneously a secure computation of the SK-DEPIR protocol and how its methods (along with $\tilde{\Pi}_{ldc}$, $\Pi_{permute}$, and Π_{store}) combine to instantiate a new read-only storage protocol $\tilde{\Pi}_{store}$.

At a high level, the protocol $\tilde{\Pi}_{store}$ operates as follows. During initialization, the parties collectively construct the LDC encoding of the database, permute it u

times, and store the concatenation of these u encoded databases E^0, E^1, \dots, E^{u-1} in an instance of Π_{store} ; the address corresponding to each element $E^i[j]$ is the concatenation of the instance number i and the location j within this instance. During a read operation, the parties look up $\tilde{\Pi}_{store}$ at one location within each permuted database E^i ; r of these lookup operations retrieve data that can collectively be used to decode the desired value, and the remaining $u - r$ lookups are “decoy” lookups that provide security under the HPN assumption.

Lemma 2 shows two settings of parameters that satisfy the HPN assumption. Using these parameters, we show that $\tilde{\Pi}_{store}$ is an efficient and secure read-only data store.

Theorem 3. *Under the HPN assumption, the protocol $\tilde{\Pi}_{store}$ in Fig. 16 securely implements functionality \mathcal{F}_{store} with constant rounds of communication. Furthermore, given any constant $\epsilon > 0$, there exist parameters m , t , u , and M such that the computation and communication cost of $\tilde{\Pi}_{store}$ is:*

- $O_\lambda(N^{1+\epsilon})$ for Init and $O_\lambda(N^\epsilon)$ for read, or
- $O_\lambda(\text{poly}(N))$ for Init and $O_\lambda(\log(N))$ for read.

Proof. Our 2PC protocol $\tilde{\Pi}_{store}$ computes all methods of the Canetti et al. SK-DEPIR protocol over boolean secret-shared data. In particular, there exist constant-round secure computation protocols for all set operations in $\tilde{\Pi}_{store}$.

- Within KeyGen: to choose a subset $T \subset [u]$, form a set of r 1s and $(u - r)$ 0s, then permute this set using $\Pi_{permute}$. The result is a secret-shared indicator vector $\langle \mathbf{T} \rangle$ of length u indicating which elements are in T .
- Within Query: form the set \tilde{Y} by oversampling. Run the LDC Query operation on u values rather than r values (using an LDC protocol with u constants x_i) to compute $Y = (\langle \mathbf{y}_i \rangle)_{i \in [u]}$, and let $\tilde{Y} = (\langle \tilde{\mathbf{y}} \rangle)_{i \in [u]}$ be a secret-shared set of u random values. Compute \tilde{Y} by multiplexing: in parallel, set each element $\langle \tilde{\mathbf{y}}_i \rangle = \langle \mathbf{T}[i] \wedge \mathbf{y}_i \oplus (\mathbf{T}[i] \oplus 1) \wedge \tilde{\mathbf{y}}_i \rangle$. (Since all values are boolean secret-shared, the bitwise-AND should be performed in 1 round between the $\mathbf{T}[i]$ and each bit of \mathbf{y}_i in turn, and similarly for the second term.)
- To truncate the list $\langle L \rangle|_T$ within Dec: first form a secret sharing of the index vector $\langle \mathbf{I} \rangle$ that equals 0 at decoy values and where $\mathbf{I}|_T = \{1, 2, \dots, r\}$ at real values by bit composing $\langle \mathbf{T} \rangle$ to an additive secret sharing $[[\mathbf{T}]]$ [9, 29], computing $[[\mathbf{I}_i]] = [[\mathbf{T}_i \cdot \sum_{j=0}^i \mathbf{T}_j]] \forall i \in [u]$, and bit decomposing $[[\mathbf{I}]]$ into a boolean secret sharing $\langle \mathbf{I} \rangle$. Then, concatenate componentwise the elements of $\langle \mathbf{I} \rangle$ and $\langle L \rangle$, permute this set using $\Pi_{permute}$, open all shares of indices \mathbf{I} in parallel, and locally sort the values of $\langle L \rangle$ using the indices.

The computational cost of Query is $O(u^2)$ due to oversampling, and the cost of the set truncation within Dec is $O(u \log(u))$ as shown by Damgard et al. [9].

As a secure two-party computation of an existing SK-DEPIR scheme, $\tilde{\Pi}_{store}$ inherits the correctness property from Def. 2, which states that the read operation always returns the correct decoded database entry. Additionally, the use of $\Pi_{permute}$ within the protocol provides the random permutation π as required

for use of the HPN assumption, so we also inherit the indistinguishability-style security property from Def. 2. Using these properties, it is straightforward to prove that $\tilde{\Pi}_{store}$ instantiates \mathcal{F}_{store} using a similar sequence of hybrids as in the original proof; we omit the details for brevity.

The claims about computational costs follow from Lemma 2 plus the following two observations. First, the cost of `Init` is dominated by the cost of the SK-DEPIR `Process` method, since the $O(M)$ cost of $\Pi_{permute}$ and the $O(u \log(u))$ cost of the oblivious search in `KeyGen` are smaller than the parameters in Lemma 2. Second, the cost of `read` is dominated by the cost of the LDC `Query` and `Dec` since the call to Π_{store} within `Resp` costs $O(\log(u \cdot M))$ as per Table 2, which is $O_\lambda(\log(N))$ since $u < N$ and $M = O_\lambda(N)$.

4.2 The new DORAM construction $\tilde{\Pi}_{DORAM}$

In this section, we show how to construct the new DORAM construction $\tilde{\Pi}_{DORAM}$ using this new instantiation $\tilde{\Pi}_{store}$ of the \mathcal{F}_{store} functionality, which only needs to be shuffled and reconstructed after a specified bound t of `write` operations have been performed, irrespective of the number of `read` operations.

The updated $\tilde{\Pi}_{DORAM}$ protocol is shown in the full version. The protocol now initializes two versions of the store, one to keep track of which elements are written and leaks access patterns and the other, $\tilde{\Pi}_{access}$, that supports unlimited reads and does not leak access patterns. This first store is critical to maintain invariant used for reshuffling. In order to know what written elements that are found in the stash, we use this store to keep track of the items ‘written’ into the store. With the distinction between reads and writes no longer hidden, $\tilde{\Pi}_{DORAM}$ only increments the epoch counter when a write is performed. Reads do not count towards the contents of the stash.

The most significant change is within the access protocol, shown in the full version [20]. It now differentiates between `read` and `write` operations. For a `read` operation it calls $\tilde{\Pi}_{store}$ and does not write anything back to the stash. The `write` operation continues to be unchanged from the original protocol.

Reshuffling is shown in the full version [20]. When it comes time to reshuffle after t writes, the protocol is similar except in one key difference. Though we now support two different stores, we perform the concatenation with elements in the stash *only* with the original store, not the augmented SK-DEPIR store. The latter does not keep track of the elements read (the indices i_M returned by $\tilde{\Pi}_{store}$ are simply random values and do not allow for the recovery of the elements). Instead we have to rely on the unread elements in the original store; namely the elements that were not written to as part of the store. The unread elements are identical to the elements found in the augmented store, so concatenation will result in the correct operation of $\tilde{\Pi}_{shuffle}$. The only other change is to instantiate these two stores, rather than the one store used within the original protocol.

We also consider the complexity of these new schemes in Table 3. Recall that we use O_λ to indicate complexities that only depend on N , ignoring any $\text{poly}(\lambda)$ terms. The main difference between the complexity of our two schemes is the blowup incurred by the new implementation of $\tilde{\Pi}_{store}$. The LDC encoding incurs

Table 3. A evaluation of each of the protocol’s server computation, bandwidth and rounds of communication where stash size, $t = N^\epsilon$. Note that we assume N reads for every 1 write and \tilde{I}_{DORAM} has been amortized where appropriate.

	Local Computation	Bandwidth	Rounds
$\tilde{I}_{DORAM}(\text{Init})$	$O_\lambda(N^{1+\epsilon} \log(N^{1+\epsilon}))$	$O_\lambda(N^{1+\epsilon} \log(N^{1+\epsilon}))$	$O(1)$
$\tilde{I}_{DORAM}(\text{op})$	$O_\lambda(N^\epsilon)$	$O_\lambda(N^\epsilon)$	$O(1)$
\tilde{I}_{access}	$O_\lambda(N^\epsilon)$	$O_\lambda(N^\epsilon)$	$O(1)$
$\tilde{I}_{shufffle}$	$O_\lambda(N^{1+\epsilon} \log(N^{1+\epsilon}))$	$O_\lambda(N^{1+\epsilon} \log(N^{1+\epsilon}))$	$O(1)$
$\tilde{I}_{store}(\text{Init})$	$O_\lambda(N^{1+\epsilon})$	$O_\lambda(N^{1+\epsilon})$	$O(1)$
$\tilde{I}_{store}(\text{read})$	$O_\lambda(N^\epsilon)$	$O_\lambda(N^\epsilon)$	$O(1)$

a $O_\lambda(N^{1+\epsilon})$ overhead for any choice of $\epsilon > 0$. This means any protocols that were original dominated by the computation or bandwidth of \tilde{I}_{store} initialization inherit this new cost. Recall in the original scheme the dominate cost of \tilde{I}_{access} was the linear scan of the stash. In this setting, with the disparity of reads vs writes, we consider a smaller stash size. If we assume one write for every N reads and our smaller stash size, \tilde{I}_{DORAM} accesses amortize to be $O_\lambda(N^\epsilon)$.

4.3 2PC for Multivariate FFT over Binary Fields

The one remaining task in the specification of protocol \tilde{I}_{lde} is to construct a secure computation of multivariate polynomial evaluation and interpolation. One effective, but slow, technique is to use Lagrange interpolation. For a univariate polynomial $p = \sum_i p \cdot x^i$, we can transform secret shares of a vector $\mathbf{p} = (p)$ of coefficients into shares of the vector $\hat{\mathbf{p}} = (p(i))$ of its evaluation at all points (or vice-versa) via multiplication by the Vandermonde matrix $\hat{\mathbf{p}} = A \cdot \mathbf{p}$, or its inverse $\mathbf{p} = A^{-1} \cdot \hat{\mathbf{p}}$, and shares of this matrix-vector multiplication can be computed locally by each party since the Vandermonde matrix A is public. However, the computational cost for matrix-vector multiplication is $\Omega(N^2)$.

The Fast Fourier Transform (FFT) [8] is a well-known algorithm for computing polynomial evaluation in quasilinear time, and the Inverse FFT similarly calculates polynomial interpolation efficiently. The fastest known FFT for binary fields is the additive FFT algorithm by Gao and Mateer [12]. As its name suggests, this algorithm solely involves linear operations. In this section, we design a secure computation protocol \tilde{I}_{mFFT} of FFT for multilinear polynomials over binary fields that can be performed locally (i.e., without interaction) with quasilinear computational cost. While this contribution may be of independent interest, in this work it completes the task from Section 4.1 of constructing a non-interactive \tilde{I}_{lde} protocol with quasilinear (rather than quadratic) computation cost. For example, in the Enc protocol within \tilde{I}_{lde} , it allows for securely com-

Fig. 17. $\tilde{\Pi}_{1FFT}$ protocol for secure 2-party computation of the Additive Fast Fourier Transform of a univariate polynomial p in a binary field \mathbb{F} .

Input: public integer h and basis $\mathcal{H} = \{v_0, v_1, \dots, v_{h-1}\}$ of a subspace of \mathbb{F} of size 2^h , plus shares $\langle \mathbf{p} \rangle$ of coefficients of a polynomial $p \in \mathbb{F}[x]$ of degree $2^h - 1$.

Output: shares $\langle \hat{\mathbf{p}} \rangle = \text{FFT}(h, \mathcal{H}, \langle \mathbf{p} \rangle)$ of the evaluations of p on all points spanned by \mathcal{H} , in the ordering specified by $\mathcal{H}[i] = \sum_{j=0}^{h-1} i_j v_j$, where i_j is the j^{th} bit of i .

1. As the base case: if $h = 1$, then return $(\langle p(0) \rangle, \langle p(v_0) \rangle)$. For a degree-1 polynomial p , we compute $\langle p(0) \rangle = \langle p_0 \rangle$ and $\langle p(v_0) \rangle = \langle p_0 \rangle + v_0 \langle p_1 \rangle$.
2. Compute the new bases $\tilde{\mathcal{H}} = \{\tilde{v}_i\}$ and $\tilde{\mathcal{H}} = \{\tilde{v}_i\}$ of size $h-1$ containing basis elements $\tilde{v}_i = v_i \cdot v_{h-1}^{-1}$ and $\tilde{v}_i = \tilde{v}_i^2 - \tilde{v}_i$ for all $i \in [h-1]$.
3. Compute coefficients $\langle q_i \rangle = v_{h-1}^i \cdot \langle p_i \rangle$ of the polynomial $q = p(v_{h-1} \cdot x)$.
4. Execute the Taylor expansion algorithm $T(h, \langle \mathbf{q} \rangle)$ in Fig. 18. Let $\langle \mathbf{f} \rangle$ and $\langle \mathbf{g} \rangle$ denote the shares of the resulting polynomials, each of degree $2^{h-1} - 1$.
5. Recursively compute $\langle \hat{\mathbf{f}} \rangle = \text{FFT}(h-1, \tilde{\mathcal{H}}, \langle \mathbf{f} \rangle)$ and $\langle \hat{\mathbf{g}} \rangle = \text{FFT}(h-1, \tilde{\mathcal{H}}, \langle \mathbf{g} \rangle)$.
6. Set $\langle \hat{p}_i \rangle = \langle \hat{f}_i \rangle + \mathcal{H}[i] \cdot \langle \hat{g}_i \rangle$ and $\langle \hat{p}_{i+2^{h-1}} \rangle = \langle \hat{p}_i \rangle + \langle \hat{g}_i \rangle \forall i \in [2^{h-1}]$. Return $\hat{\mathbf{p}}$.

putting the coefficients of the polynomial $\psi : H^m \rightarrow \mathbb{F}$ in time $O(N \log^2 N)$ and securely evaluating the polynomial ψ at all locations in \mathbb{F}^m in time $O(M \log^2 M)$.

We describe this protocol in two steps. First, we show how to securely evaluate FFT for univariate polynomials (building a secure computation of Taylor series expansion as a building block). Second, we bootstrap to a secure evaluation of FFT for multivariate polynomials. For brevity, we show these FFT protocols only in the forward (polynomial evaluation) direction. It is straightforward to validate that the same techniques apply to construct a secure computation protocol of inverse FFT (i.e., polynomial interpolation) in quasilinear time.

2PC protocol $\tilde{\Pi}_{1FFT}$ for univariate FFT. In this section, we present a secure two-party computation protocol $\tilde{\Pi}_{1FFT}$. Let $H \subset \mathbb{F}$ be a subspace (possibly the entire field) of size $|H| = 2^h$ defined by a basis \mathcal{H} , and let $p = \sum_{i=0}^{2^h-1} p_i \cdot x^i$ be a univariate polynomial of degree less than 2^h . This protocol begins with shares of the 2^h coefficients $\langle \mathbf{p} \rangle = (\langle p_i \rangle)_{i \in [2^h]}$ of the polynomial, and it returns the shares $\langle \hat{\mathbf{p}} \rangle = (\langle p(i) \rangle)_{i \in [2^h]}$ of its evaluation at all 2^h points in H .

The protocol $\tilde{\Pi}_{1FFT}$ is shown in Fig. 17, and it uses the Taylor series expansion algorithm in Fig. 18 as a building block. Each step of these algorithms only involves addition and scalar multiplication of secret-shared values, so the secure computation $\tilde{\Pi}_{1FFT}$ can be performed locally. These algorithms are precisely the secret-shared versions of their counterparts in Gao and Mateer [12].

We provide a high-level intuition of $\tilde{\Pi}_{1FFT}$ when considering the basis $\mathcal{H} = \{z^{h-1}, z^{h-2}, \dots, z, 1\}$, in which case $q = p$; full details are given in [12]. The core idea of the Fast Fourier Transform is to reduce the evaluation of one polynomial q into the evaluation of two polynomials f and g of half the degree, plus quasilinear

Fig. 18. Protocol for Taylor expansion of a polynomial $q(x) \in \mathbb{F}[x]$ at $x^2 - x$.

Input: public integer h , shares of coefficients $\langle \mathbf{q} \rangle$ of a polynomial of degree $2^h - 1$.

Output: $\langle \mathbf{f} \rangle, \langle \mathbf{g} \rangle = \text{T}(h, \langle \mathbf{q} \rangle)$ such that each vector is of length $\leq 2^{h-1}$ and they collectively form the Taylor series expansion $q(x) = \sum_{i=0}^{2^h-1} (f_i + g_i x) \cdot (x^2 - x)^i$.

1. As the base case: if $h = 1$ so $\deg(q) = 1$, return $\langle f_0 \rangle = \langle g_0 \rangle$ and $\langle g_0 \rangle = \langle q_1 \rangle$.
2. Partition the vector $\langle \mathbf{q} \rangle$ into $\langle \mathbf{t}^0 \rangle$ containing the first 2^{h-1} elements, $\langle \mathbf{t}^1 \rangle$ containing the next 2^{h-2} elements, and $\langle \mathbf{t}^2 \rangle$ containing the last 2^{h-2} elements.
3. Compute the vectors $\langle \mathbf{t} \rangle = \langle \mathbf{t}^1 \rangle + \langle \mathbf{t}^2 \rangle$ of length 2^{h-2} , $\langle \mathbf{q}^0 \rangle = \langle \mathbf{t}^0 \rangle + (\mathbf{0} \parallel \langle \mathbf{t} \rangle)$ of length 2^{h-1} , and $\langle \mathbf{q}^1 \rangle = (\langle \mathbf{t} \rangle \parallel \langle \mathbf{t}^2 \rangle)$ of length 2^{h-1} . Here, $\mathbf{0}$ denotes the vector containing 2^{h-2} zero elements, and \parallel denotes vector concatenation.
4. Recursively, find $\langle \mathbf{f}^0 \rangle, \langle \mathbf{g}^0 \rangle = \text{T}(h-1, \langle \mathbf{q}^0 \rangle)$ and $\langle \mathbf{f}^1 \rangle, \langle \mathbf{g}^1 \rangle = \text{T}(h-1, \langle \mathbf{q}^1 \rangle)$.
5. Return the concatenated vectors $\langle \mathbf{f} \rangle = \langle \mathbf{f}^0 \rangle \parallel \langle \mathbf{f}^1 \rangle$ and $\langle \mathbf{g} \rangle = \langle \mathbf{g}^0 \rangle \parallel \langle \mathbf{g}^1 \rangle$.

work to “stitch” the results together into an evaluation of q . Gao and Mateer [12] show how this can be done over binary fields, based on these observations:

- The Taylor expansion $q(x) = \sum_{i=0}^{2^h-1} (f_i + g_i x) \cdot (x^2 - x)^i$ leads to an equation $q(x) = f(x^2 - x) + x \cdot g(x^2 - x)$ involving polynomials $f(z) \triangleq \sum_{i=0}^{2^h-1} f_i \cdot z^i$ and $g(z) \triangleq \sum_{i=0}^{2^h-1} g_i \cdot z^i$ of lower degree $2^{h-1} - 1$.
- The function $x \mapsto x^2 - x$ is 2-to-1, and specifically it maps the 2^h -sized space spanned by \mathcal{H} into the smaller 2^{h-1} space spanned by the basis $\tilde{\mathcal{H}}$.

Ergo, in order to evaluate the polynomial q at all points spanned by \mathcal{H} , it suffices to evaluate polynomials f and g at all points spanned by the smaller basis $\tilde{\mathcal{H}}$ and combine the results using the Taylor expansion $q(x) = f(x^2 - x) + x \cdot g(x^2 - x)$.

We provide a secure 2-party computation of Gao and Mateer’s method of computing the Taylor expansion of q in Fig. 18, and we provide a 2PC of polynomial evaluation in Fig. 17. The only operations that involve secret-shared data are linear combinations and splitting/joining vectors, all of which can be performed locally. Note that step 2 of Fig. 17 involves more complicated algebra, but it only involves public (non-secret-shared) values, so it can be performed locally and pre-computed before parties receive their input shares.

2PC protocol $\tilde{\Pi}_{mFFT}$ for multivariate FFT. Recall that the locally decodable code used in $\tilde{\Pi}_{ldc}$ is based on Reed-Muller codes, and as a result it uses multivariate polynomials. Here, we show how to bootstrap from an FFT for univariate polynomials into one for multivariate polynomials. The full protocol is shown in Fig. 19, and it is based on a technique used by Kedlaya and Umans [24].

Protocol $\tilde{\Pi}_{mFFT}$ operates via recursion over many evaluations of univariate polynomials. Given an m -variate polynomial p of total degree $< 2^h - 1$ for which the parties have shares of all coefficients, we rewrite the polynomial by

Fig. 19. $\tilde{\Pi}_{mFFT}$ protocol for secure 2-party evaluation of a multivariate polynomial $p(x_0, \dots, x_{m-1})$ at all points in a subspace $H^m \in \mathbb{F}^m$.

Input: shares of coefficients $\langle \mathbf{p} \rangle$ of an m -variate polynomial p of total degree $2^h - 1$.

Output: shares $\langle \hat{\mathbf{p}} \rangle$ of evaluations $p(\mathbf{x})$ at all points $\mathbf{x} \in H^m$.

1. As the base case: if $m = 1$, then run protocol $\tilde{\Pi}_{1FFT}$ as shown in Fig. 17.
2. By rearranging terms, write $p(\mathbf{x}) = \sum_{i=0}^{2^h-1} x_0^i \cdot p_i(x_1, x_2, \dots, x_{m-1})$. Observe that the parties collectively hold shares of the coefficients of each $\langle \mathbf{p}_i \rangle$.
3. Recursively, get shares $\langle \hat{\mathbf{p}}_i \rangle$ of evaluations of each p_i at all points in H^{m-1} .
4. For each vector $\mathbf{c} \in H^{m-1}$, compute shares of the evaluation of the univariate polynomial $\langle p(x_0, \mathbf{c}) \rangle = \sum_{i=0}^{2^h-1} \langle \hat{\mathbf{p}}_i(\mathbf{c}) \rangle \cdot x_0^i$ on all points in H using $\tilde{\Pi}_{1FFT}$.

conditioning on the power of the first variable: $p(x_0, \dots, x_{m-1}) = \sum_{i=0}^{2^h-1} x_0^i \cdot p_i(x_1, x_2, \dots, x_{m-1})$. We can evaluate the $(m-1)$ -variate polynomials p_i recursively, and use the results to evaluate the univariate polynomial over x_0 . Since each univariate polynomial evaluation takes time quasilinear in $|H|$, a simple recurrence relation shows that the entire evaluation is quasilinear in $|H|^m = N$. This completes the construction, and it is the necessary building block to complete the proof of Theorem 2 and achieve quasilinear server computation for our LDC protocol.

Acknowledgments We gratefully acknowledge conversations with Daniel Wichs and Jack Doerner for their valuable insights. This material is based upon work supported by the National Science Foundation under Grants 1414119, 1718135, 1750795, and 1931714.

References

1. Abraham, I., Fletcher, C.W., Nayak, K., Pinkas, B., Ren, L.: Asymptotically tight bounds for composing ORAM with PIR. In: Public Key Cryptography (1). Lecture Notes in Computer Science, vol. 10174, pp. 91–120. Springer (2017)
2. Asharov, G., Komargodski, I., Lin, W., Nayak, K., Peserico, E., Shi, E.: Optorama: Optimal oblivious RAM. In: EUROCRYPT (2). Lecture Notes in Computer Science, vol. 12106, pp. 403–432. Springer (2020)
3. Boyle, E., Holmgren, J., Weiss, M.: Permuted puzzles and cryptographic hardness. In: TCC (2). Lecture Notes in Computer Science, vol. 11892, pp. 465–493. Springer (2019)
4. Boyle, E., Ishai, Y., Pass, R., Wootters, M.: Can we access a database both locally and privately? In: TCC (2). Lecture Notes in Computer Science, vol. 10678, pp. 662–693. Springer (2017)
5. Bunn, P., Katz, J., Kushilevitz, E., Ostrovsky, R.: Efficient 3-party distributed ORAM. In: SCN. Lecture Notes in Computer Science, vol. 12238, pp. 215–232. Springer (2020)

6. Canetti, R., Holmgren, J., Richelson, S.: Towards doubly efficient private information retrieval. In: TCC (2). Lecture Notes in Computer Science, vol. 10678, pp. 694–726. Springer (2017)
7. Chan, T.H., Katz, J., Nayak, K., Polychroniadou, A., Shi, E.: More is less: Perfectly secure oblivious algorithms in the multi-server setting. In: ASIACRYPT (3). Lecture Notes in Computer Science, vol. 11274, pp. 158–188. Springer (2018)
8. Cooley, J.W., Tukey, J.W.: An algorithm for the machine calculation of complex fourier series. *Mathematics of computation* **19**(90), 297–301 (1965)
9. Damgård, I., Fitzi, M., Kiltz, E., Nielsen, J.B., Toft, T.: Unconditionally secure constant-rounds multi-party computation for equality, comparison, bits and exponentiation. In: TCC. Lecture Notes in Computer Science, vol. 3876, pp. 285–304. Springer (2006)
10. Dodis, Y., Yampolskiy, A.: A verifiable random function with short proofs and keys. In: Public Key Cryptography. Lecture Notes in Computer Science, vol. 3386, pp. 416–431. Springer (2005)
11. Doerner, J., Shelat, A.: Scaling ORAM for secure computation. In: ACM Conference on Computer and Communications Security. pp. 523–535. ACM (2017)
12. Gao, S., Mateer, T.D.: Additive fast fourier transforms over finite fields. *IEEE Trans. Inf. Theory* **56**(12), 6265–6272 (2010)
13. Garg, S., Lu, S., Ostrovsky, R.: Black-box garbled RAM. In: FOCS. pp. 210–229. IEEE Computer Society (2015)
14. Garg, S., Lu, S., Ostrovsky, R., Scafuro, A.: Garbled RAM from one-way functions. In: STOC. pp. 449–458. ACM (2015)
15. Garg, S., Mohassel, P., Papamanthou, C.: TWORAM: round-optimal oblivious RAM with applications to searchable encryption. *IACR Cryptol. ePrint Arch.* **2015**, 1010 (2015)
16. Gentry, C., Halevi, S., Lu, S., Ostrovsky, R., Raykova, M., Wichs, D.: Garbled RAM revisited. In: EUROCRYPT. Lecture Notes in Computer Science, vol. 8441, pp. 405–422. Springer (2014)
17. Goldreich, O., Ostrovsky, R.: Software protection and simulation on oblivious rams. *J. ACM* **43**(3), 431–473 (1996)
18. Gordon, S.D., Katz, J., Kolesnikov, V., Krell, F., Malkin, T., Raykova, M., Vahlis, Y.: Secure two-party computation in sublinear (amortized) time. In: ACM Conference on Computer and Communications Security. pp. 513–524. ACM (2012)
19. Gordon, S.D., Katz, J., Wang, X.: Simple and efficient two-server ORAM. In: ASIACRYPT (3). Lecture Notes in Computer Science, vol. 11274, pp. 141–157. Springer (2018)
20. Hamlin, A., Varia, M.: Two-server distributed ORAM with sublinear computation and constant rounds. *IACR Cryptol. ePrint Arch.* **2020**, 1547 (2020)
21. Hoang, T., Yavuz, A.A., Durak, F.B., Guajardo, J.: Oblivious dynamic searchable encryption via distributed PIR and ORAM. *IACR Cryptol. ePrint Arch.* **2017**, 1158 (2017)
22. Jarecki, S., Liu, X.: Efficient oblivious pseudorandom function with applications to adaptive OT and secure computation of set intersection. In: TCC. Lecture Notes in Computer Science, vol. 5444, pp. 577–594. Springer (2009)
23. Jarecki, S., Wei, B.: 3pc ORAM with low latency, low bandwidth, and fast batch retrieval. In: ACNS. Lecture Notes in Computer Science, vol. 10892, pp. 360–378. Springer (2018)
24. Kedlaya, K.S., Umans, C.: Fast modular composition in any characteristic. In: FOCS. pp. 146–155. IEEE Computer Society (2008)

25. Kushilevitz, E., Mour, T.: Sub-logarithmic distributed oblivious RAM with small block size. In: *Public Key Cryptography (1)*. Lecture Notes in Computer Science, vol. 11442, pp. 3–33. Springer (2019)
26. Lu, S., Ostrovsky, R.: Distributed oblivious RAM for secure two-party computation. In: *TCC*. Lecture Notes in Computer Science, vol. 7785, pp. 377–396. Springer (2013)
27. Lu, S., Ostrovsky, R.: How to garble RAM programs. In: *EUROCRYPT*. Lecture Notes in Computer Science, vol. 7881, pp. 719–734. Springer (2013)
28. Miao, P., Patel, S., Raykova, M., Seth, K., Yung, M.: Two-sided malicious security for private intersection-sum with cardinality. In: *CRYPTO (3)*. Lecture Notes in Computer Science, vol. 12172, pp. 3–33. Springer (2020)
29. Nishide, T., Ohta, K.: Multiparty computation for interval, equality, and comparison without bit-decomposition protocol. In: *Public Key Cryptography*. Lecture Notes in Computer Science, vol. 4450, pp. 343–360. Springer (2007)
30. Ostrovsky, R., Shoup, V.: Private information storage (extended abstract). In: *STOC*. pp. 294–303. ACM (1997)
31. Patel, S., Persiano, G., Raykova, M., Yeo, K.: Panorama: Oblivious RAM with logarithmic overhead. In: *FOCS*. pp. 871–882. IEEE Computer Society (2018)
32. Roche, D.S., Aviv, A.J., Choi, S.G.: A practical oblivious map data structure with secure deletion and history independence. In: *IEEE Symposium on Security and Privacy*. pp. 178–197. IEEE Computer Society (2016)
33. Wang, X., Chan, T.H., Shi, E.: Circuit ORAM: on tightness of the goldreich-ostrovsky lower bound. In: *ACM Conference on Computer and Communications Security*. pp. 850–861. ACM (2015)
34. Wang, X.S., Huang, Y., Chan, T.H., Shelat, A., Shi, E.: SCORAM: oblivious RAM for secure computation. In: *ACM Conference on Computer and Communications Security*. pp. 191–202. ACM (2014)
35. Weiss, M., Wichs, D.: Is there an oblivious RAM lower bound for online reads? In: *TCC (2)*. Lecture Notes in Computer Science, vol. 11240, pp. 603–635. Springer (2018)
36. Zahur, S., Wang, X., Raykova, M., Gascón, A., Doerner, J., Evans, D., Katz, J.: Revisiting square-root ORAM: efficient random access in multi-party computation. In: *IEEE Symposium on Security and Privacy*. pp. 218–234. IEEE Computer Society (2016)
37. Zhang, J., Ma, Q., Zhang, W., Qiao, D.: MSKT-ORAM: A constant bandwidth ORAM without homomorphic encryption. *IACR Cryptol. ePrint Arch.* **2016**, 882 (2016)