# Laconic Oblivious Transfer and its Applications 

Chongwon Cho<br>HRL Laboratories

Nico Döttling*<br>UC Berkeley

Sanjam Garg ${ }^{\dagger}$<br>UC Berkeley

Divya Gupta ${ }^{\dagger, \ddagger}$

Microsoft Research India

Peihan Miao ${ }^{\dagger}$<br>UC Berkeley

Antigoni Polychroniadou ${ }^{\S}$<br>Cornell University


#### Abstract

In this work, we introduce a novel technique for secure computation over large inputs. Specifically, we provide a new oblivious transfer (OT) protocol with a laconic receiver. Laconic OT allows a receiver to commit to a large input $D$ (of length $M$ ) via a short message. Subsequently, a single short message by a sender allows the receiver to learn $m_{D[L]}$, where the messages $m_{0}, m_{1}$ and the location $L \in[M]$ are dynamically chosen by the sender. All prior constructions of OT required the receiver's outgoing message to grow with $D$.

Our key contribution is an instantiation of this primitive based on the Decisional DiffieHellman (DDH) assumption in the common reference string (CRS) model. The technical core of this construction is a novel use of somewhere statistically binding (SSB) hashing in conjunction with hash proof systems. Next, we show applications of laconic OT to non-interactive secure computation on large inputs and multi-hop homomorphic encryption for RAM programs.


[^0]
## Contents

1 Introduction ..... 4
1.1 Laconic OT ..... 4
1.2 Warm-Up Application: Non-Interactive Secure Computation on Large Inputs ..... 5
1.3 Main Application: Muti-Hop Homomorphic Encryption for RAM Programs ..... 6
1.4 Roadmap ..... 7
2 Technical Overview ..... 8
2.1 Laconic OT ..... 8
2.1.1 Laconic OT with Factor-2 Compression ..... 8
2.1.2 Bootstrapping Laconic OT ..... 10
2.2 Non-interactive Secure Computation on Large Inputs ..... 12
2.3 Multi-Hop Homomorphic Encryption for RAM Programs ..... 13
3 Laconic Oblivious Transfer ..... 14
3.1 Laconic OT ..... 15
3.2 Updatable Laconic OT ..... 17
4 Laconic Oblivious Transfer with Factor-2 Compression ..... 18
4.1 Somewhere Statistically Binding Hash Functions and Hash Proof Systems ..... 18
4.2 HPS-friendly SSB Hashing ..... 19
4.3 A Hash Proof System for Knowledge of Preimage Bits ..... 22
4.4 The Laconic OT Scheme ..... 23
5 Construction of Updatable Laconic OT ..... 26
5.1 Background ..... 26
5.1.1 Garbled Circuits ..... 26
5.1.2 Merkle Tree ..... 27
5.2 Construction ..... 27
5.3 Security ..... 32
6 Warm-Up Application: Non-Interactive Secure Computation (NISC) on Large Inputs in RAM Setting ..... 36
6.1 Background ..... 36
6.1.1 Random Access Machine (RAM) Model of Computation ..... 36
6.1.2 Oblivious Transfer ..... 37
6.2 Formal Model for NISC in RAM Setting ..... 38
6.3 Construction ..... 40
6.4 Correctness ..... 42
6.5 Security Proof ..... 44
6.6 Extension ..... 46
7 Main Application: Multi-Hop Homomorphic Encryption for RAM Programs ..... 47
7.1 Our Model ..... 47
7.2 Building Blocks Needed ..... 50
7.2.1 2-message Secure Function Evaluation based on Garbled Circuits ..... 50
7.2.2 Re-Randomizable Secure Function Evaluation based on Garbled Circuits ..... 50
7.3 Our Construction of Multi-hop RAM Scheme ..... 52
7.3.1 UMA Secure Construction ..... 52
7.3.2 Correctness ..... 58
7.3.3 Extending to Multiple Executions ..... 61
7.3.4 Security Proof ..... 61
7.3.5 UMA to Full Security for Multi-hop RAM Scheme ..... 66

## 1 Introduction

Big data poses serious challenges for the current cryptographic technology. In particular, cryptographic protocols for secure computation are typically based on Boolean circuits, where both the computational complexity and communication complexity scale with the size of the input dataset, which makes it generally unsuitable for even moderate dataset sizes. Over the past few decades, substantial effort has been devoted towards realizing cryptographic primitives that overcome these challenges. This includes works on fully-homomorphic encryption (FHE) [Gen09, BV11b, BV11a, GSW13] and on the RAM setting of oblivious RAM [Gol87, Ost90] and secure RAM computation [OS97, GKK ${ }^{+}$12, LO13, GHL+14, GGMP16]. Protocols based on FHE generally have a favorable communication complexity and are basically non-interactive, yet incur a prohibitively large computational overhead (dependent on the dataset size). On the other hand, protocols for the RAM model generally have a favorable computational overhead, but lack in terms of communication efficiency (that grows with the program running time), especially in the multi-party setting. Can we achieve the best of both worlds? In this work we make positive progress on this question. Specifically, we introduce a new tool called laconic oblivious transfer that helps to strike a balance between the two seemingly opposing goals.

Oblivious transfer (or OT for short) is a fundamental and powerful primitive in cryptography [Kil88, IPS08]. Since its first introduction by Rabin [Rab81], OT has been a foundational building block for realizing secure computation protocols [Yao82, GMW87, IPS08]. However, typical secure computation protocols involve executions of multiple instances of an oblivious transfer protocol. In fact, the number of needed oblivious transfers grows with the input size of one of the parties, which is the receiver of the oblivious transfer. ${ }^{1}$ In this work, we observe that a two-message OT protocol, with a short message from the receiver, can be a key tool towards the goal of obtaining simultaneous improvements in computational and communication cost for secure computation.

### 1.1 Laconic OT

In this paper, we introduce the notion of laconic oblivious transfer (or laconic OT for short). Laconic OT allows an OT receiver to commit to a large input $D \in\{0,1\}^{M}$ via a short message. Subsequently, the sender responds with a single short message to the receiver depending on dynamically chosen two messages $m_{0}, m_{1}$ and a location $L \in[M]$. The sender's response message allows the receiver to recover $m_{D[L]}$ (while $m_{1-D[L]}$ remains computationally hidden). Furthermore, without any additional communication with the receiver, the sender could repeat this process for multiple choices of $L$. The construction we give is secure against semi-honest adversaries, but it can be upgraded to the malicious setting in a similar way as we will discuss in Section 1.2 for the first application.

Our construction of laconic OT is obtained by first realizing a "mildly compressing" laconic OT protocol for which the receiver's message is factor-2 compressing, i.e., half the size of its input. We base this construction on the Decisional Diffie-Hellman (DDH) assumption. We note that, subsequent to our work, the factor-2 compression construction has been simplified by Döttling and Garg [?] (another alternative simplification can be obtained using [?]). Next we show that such a "mildly compressing" laconic OT can be bootstrapped, via the usage of a Merkle Hash Tree

[^1]and Yao's Garbled Circuits [Yao82], to obtain a "fully compressing" laconic OT, where the size of the receiver's message is independent of its input size. The laconic OT scheme with a Merkle Tree structure allows for good properties like local verification and local updates, which makes it a powerful tool in secure computation with large inputs.

We will show new applications of laconic OT to non-interactive secure computation and homomorphic encryption for RAM programs, as briefly described below in Sections 1.2 and 1.3.

### 1.2 Warm-Up Application: Non-Interactive Secure Computation on Large Inputs

Can a receiver publish a (small) encoding of her large confidential database $D$ so that any sender, who holds a secret input $x$, can reveal the output $f(x, D)$ (where $f$ is a circuit) to the receiver by sending her a single message? For security, we want the receiver's encoding to hide $D$ and the sender's message to hide $x$. Using laconic OT, we present the first solution to this problem. In our construction, the receiver's published encoding is independent of the size of her database, but we do not restrict the size of the sender's message. ${ }^{2}$

RAM Setting. Consider the scenario where $f$ can be computed using a RAM program $P$ of running time $t$. We use the notation $P^{D}(x)$ to denote the execution of the program $P$ on input $x$ with random access to the database $D$. We provide a construction where as before the size of the receiver's published message is independent of the size of the database $D$. Moreover, the size of the sender's message (and computational cost of the sender and the receiver) grows only with $t$ and the receiver learns nothing more than the output $P^{D}(x)$ and the locations in $D$ touched during the computation. Note that in all prior works on general secure RAM computation [OS97, GKK ${ }^{+} 12$, LO13, $\mathrm{WHC}^{+} 14, \mathrm{GHL}^{+} 14$, GLOS15, GLO15] the size of the receiver's message grew at least with its input size. ${ }^{3}$

Against Malicious Adversaries. The results above are obtained in the semi-honest setting. We can upgrade to security against a malicious sender by use of (i) non-interactive zero knowledge proofs (NIZKs) [FLS90] at the cost of additionally assuming doubly enhanced trapdoor permutations or bilinear maps [CHK04, GOS06], (ii) the techniques of Ishai et al. [IKO ${ }^{+}$11] while obtaining slightly weaker security, ${ }^{4}$ or (iii) interactive zero-knowledge proofs but at the cost of additional interaction.

Upgrading to security against a malicious receiver is tricky. This is because the receiver's public encoding is short and hence, it is not possible to recover the receiver's entire database just given the encoding. Standard simulation-based security can be obtained by using (i) universal arguments

[^2]as done by [CV12, COV15] at the cost of additional interaction, or (ii) using SNARKs at the cost of making extractability assumptions $\left[B C C T 12, \mathrm{BSCG}^{+} 13\right] .{ }^{5}$

Other Related Work. Prior works consider secure computation which hides the input size of one [MRK03,IP07,ADT11,LNO13] or both parties [LNO13]. Our notion only requires the receiver's communication cost to be independent of the its input size, and is therefore weaker. However, these results are largely restricted to special functionalities, such as zero-knowledge sets and computing certain branching programs (which imply input-size hiding private set intersection). The general result of [LNO13] uses FHE and as mentioned earlier needs more rounds of interaction. ${ }^{6}$

### 1.3 Main Application: Muti-Hop Homomorphic Encryption for RAM Programs

Consider a scenario where $S$ (a server), holding an input $x$, publishes an encryption $\mathrm{ct}_{0}$ of her private input $x$ under her public key. Now this ciphertext is passed on to a client $Q_{1}$ that homomorphically computes a (possibly private) program $P_{1}$ accessing (private) memory $D_{1}$ on the value encrypted in $\mathrm{ct}_{0}$, obtaining another ciphertext $\mathrm{ct}_{1}$. More generally, the computation could be performed by multiple clients. In other words, clients $Q_{2}, Q_{3}, \cdots$ could sequentially compute private programs $P_{2}, P_{3}, \cdots$ accessing their own private databases $D_{2}, D_{3}, \cdots$. Finally, we want $S$ to be able to use her secret key to decrypt the final ciphertext and recover the output of the computation. For security, we require simulation based security for a client $Q_{i}$ against a collusion of the server and any subset of the clients, and IND-CPA security for the server's ciphertext.

Though we described the simple case above, we are interested in the general case when computation is performed in different sequences of the clients. Examples of two such computation paths are shown in Figure 1. Furthermore, we consider the setting of persistent databases, where each client is able to execute dynamically chosen programs on the encrypted ciphertexts while using the same database that gets updated as these programs are executed.

FHE-Based Solution. Gentry's [Gen09] fully homomorphic encryption (FHE) scheme offers a solution to the above problem when circuit representations of the desired programs $P_{1}, P_{2}, \ldots$ are considered. Specifically, $S$ could encrypt her input $x$ using an FHE

[^3]scheme. Now, the clients can publicly com-
pute arbitrary programs on the encrypted value using a public evaluation procedure. This procedure can be adapted to preserve the privacy of the computed circuit [OPP14, DS16, BPMW16] as well. However, this construction only works for circuits. Realizing the scheme for RAM programs involves first converting the RAM program into a circuit of size at least linear in the size of the database. This linear effort can be exponential in the running time of the program for several applications of interest such as binary search.

Our Relaxation. In obtaining homomorphic encryption for RAM programs, we start by relaxing the compactness requirement in FHE. ${ }^{7}$ Compactness in FHE requires that the size of the ciphertexts does not grow with computation. In particular, in our scheme, we allow the evaluated ciphertexts to be bigger than the original ciphertext. Gentry, Halevi and Vaikuntanathan [GHV10] considered an analogous setting for the case of circuits. As in Gentry et al. [GHV10], in our setting computation itself will happen at the time of decryption. Therefore, we additionally require that clients $Q_{1}, Q_{2}, \cdots$ first ship pre-processed versions of their databases to $S$ for the decryption, and security will additionally require that $S$ does not learn the access pattern of the programs on client databases. This brings us to the following question:

Can we realize multi-hop encryption schemes for RAM programs where the ciphertext grows linearly only in the running time of the computation performed on it?

We show that laconic OT can be used to realize such a multi-hop homomorphic encryption scheme for RAM programs. Our result bridges the gap between growth in ciphertext size and computational complexity of homomorphic encryption for RAM programs.

Our work also leaves open the problem of realizing (fully or somewhat) homomorphic encryption for RAM programs with (somewhat) compact ciphertexts and for which computational cost grows with the running time of the computation, based on traditional computational assumptions. Our solution for multi-hop RAM homomorphic encryption is for the semi-honest (or, semi-malicious) setting only. We leave open the problem of obtaining a solution in the malicious setting. ${ }^{8}$

### 1.4 Roadmap

We now lay out a roadmap for the remainder of the paper. In Section 2 we give a technical overview of this work. We introduce the notion of laconic OT formally in Section 3, and give a construction with factor-2 compression in Section 4, which can be bootstrapped to a fully compressing updatable laconic OT in Section 5. Finally we present our two applications in Sections 6 and 7.

[^4]
## 2 Technical Overview

### 2.1 Laconic OT

We will now provide an overview of laconic OT and our constructions of this new primitive. Laconic OT consists of two major components: a hash function and an encryption scheme. We will call the hash function Hash and the encryption scheme (Send, Receive). In a nutshell, laconic OT allows a receiver $R$ to compute a succinct digest digest of a large database $D$ and a private state $\hat{\mathrm{D}}$ using the hash function Hash. After digest is made public, anyone can non-interactively send OT messages to $R$ w.r.t. a location $L$ of the database such that the receiver's choice bit is $D[L]$. Here, $D[L]$ is the database-entry at location $L$. In more detail, given digest, a database location $L$, and two messages $m_{0}$ and $m_{1}$, the algorithm Send computes a ciphertext e such that $R$, who owns $\hat{\mathrm{D}}$, can use the decryption algorithm Receive to decrypt e to obtain the message $m_{D[L]}$.

For security, we require sender privacy against semi-honest receiver. In particular, given an honest receiver's view, which includes the database $D$, the message $m_{1-D[L]}$ is computationally hidden. We formalize this using a simulation based definition. On the other hand, we do not require receiver privacy as opposed to standard oblivious transfer, namely, no security guarantee is provided against a cheating (semi-honest) sender. This is mostly for ease of exposition. Nevertheless, adding receiver privacy to laconic OT can be done in a straightforward manner via the usage of garbled circuits and two-message OT (see Section 3.1 for a detailed discussion).

For efficiency, we have the following requirement: First, the size of digest only depends on the security parameter and is independent of the size of the database $D$. Moreover, after digest and $\hat{D}$ are computed by Hash, the workload of both the sender and receiver (that is, the runtime of both Send and Receive) becomes essentially independent of the size of the database (i.e., depending at most polynomially on $\log (|D|))$.

Notice that our security definition and efficiency requirement immediately imply that the Hash algorithm used to compute the succinct digest must be collision resistant. Thus, it is clear that the hash function must be keyed and in our case it is keyed by a common reference string.

Construction at a high level. We first construct a laconic OT scheme with factor-2 compression, which compresses a $2 \lambda$-bit database to a $\lambda$-bit digest. Next, to get laconic OT for databases of arbitrary size, we bootstrap this construction using an interesting combination of Merkle hashing and garbled circuits. Below, we give an overview of each of these steps.

### 2.1.1 Laconic OT with Factor-2 Compression

We start with a construction of a laconic OT scheme with factor-2 compression, i.e., a scheme that hashes a $2 \lambda$-bit database to a $\lambda$-bit digest. This construction is inspired by the notion of witness encryption [GGSW13]. We will first explain the scheme based on witness encryption. Then, we show how this specific witness encryption scheme can be realized with the more standard notion of hash proof systems (HPS) [CS02]. Our overall scheme will be based on the security of Decisional Diffie-Hellman (DDH) assumption.

Construction Using Witness Encryption. Recall that a witness encryption scheme is defined for an NP-language $\mathcal{L}$ (with corresponding witness relation $\mathcal{R}$ ). It consists of two algorithms Enc and Dec. The algorithm Enc takes as input a problem instance $x$ and a message $m$, and
produces a ciphertext. A recipient of the ciphertext can use Dec to decrypt the message if $x \in \mathcal{L}$ and the recipient knows a witness $w$ such that $\mathcal{R}(x, w)$ holds. There are two requirements for a witness encryption scheme, correctness and security. Correctness requires that if $\mathcal{R}(x, w)$ holds, then $\operatorname{Dec}(x, w, \operatorname{Enc}(x, m))=m$. Security requires that if $x \notin \mathcal{L}$, then $\operatorname{Enc}(x, m)$ computationally hides $m$.

We will now discuss how to construct a laconic OT with factor-2 compression using a two-to-one hash function and witness encryption. Let $\mathrm{H}: \mathcal{K} \times\{0,1\}^{2 \lambda} \rightarrow\{0,1\}^{\lambda}$ be a keyed hash function, where $\mathcal{K}$ is the key space. Consider the language $\mathcal{L}=\left\{(K, L, y, b) \in \mathcal{K} \times[2 \lambda] \times\{0,1\}^{\lambda} \times\{0,1\} \mid \exists D \in\right.$ $\{0,1\}^{2 \lambda}$ such that $\mathrm{H}(K, D)=y$ and $\left.D[L]=b\right\}$. Let (Enc, Dec) be a witness encryption scheme for the language $\mathcal{L}$.

The laconic OT scheme is as follows: The Hash algorithm computes $y=\mathrm{H}(K, D)$ where $K$ is the common reference string and $D \in\{0,1\}^{2 \lambda}$ is the database. Then $y$ is published as the digest of the database. The Send algorithm takes as input $K, y$, a location $L$, and two messages $\left(m_{0}, m_{1}\right)$ and proceeds as follows. It computes two ciphertexts $\mathrm{e}_{0} \leftarrow \operatorname{Enc}\left((K, L, y, 0), m_{0}\right)$ and $\mathrm{e}_{1} \leftarrow \operatorname{Enc}\left((K, L, y, 1), m_{1}\right)$ and outputs $\mathrm{e}=\left(\mathrm{e}_{0}, \mathrm{e}_{1}\right)$. The Receive algorithm takes as input $K, L, y, D$, and the ciphertext $\mathrm{e}=\left(\mathrm{e}_{0}, \mathrm{e}_{1}\right)$ and proceeds as follows. It sets $b=D[L]$, computes $m \leftarrow \operatorname{Dec}\left((K, L, y, b), D, \mathrm{e}_{b}\right)$ and outputs $m$.

It is easy to check that the above scheme satisfies correctness. However, we run into trouble when trying to prove sender privacy. Since H compresses $2 \lambda$ bits to $\lambda$ bits, most hash values have exponentially many pre-images. This implies that for most values of ( $K, L, y$ ), it holds that both $(K, L, y, 0) \in \mathcal{L}$ and $(K, L, y, 1) \in \mathcal{L}$, that is, most problem instances are yes-instances. However, to reduce sender privacy of our scheme to the security of witness encryption, we ideally want that if $y=\mathrm{H}(K, D)$, then $(K, L, y, D[L]) \in \mathcal{L}$ while $(K, L, y, 1-D[L]) \notin \mathcal{L}$. To overcome this problem, we will use a somewhere statistically binding hash function that allows us to artificially introduce no-instances as described below.

Somewhere Statistically Binding Hash to the Rescue. Somewhere statistically binding (SSB) hash functions [HW15, KLW15, OPWW15] support a special key generation procedure such that the hash value information theoretically fixes certain bit(s) of the pre-image. In particular, the special key generation procedure takes as input a location $L$ and generates a key $K^{(L)}$. Then the hash function keyed by $K^{(L)}$ will bind the $L$-th bit of the pre-image. That is, $K^{(L)}$ and $y=\mathrm{H}\left(K^{(L)}, D\right)$ uniquely determines $D[L]$. The security requirement for SSB hashing is the indexhiding property, i.e., keys $K^{(L)}$ and $K^{\left(L^{\prime}\right)}$ should be computationally indistinguishable for any $L \neq L^{\prime}$.

We can now establish security of the above laconic OT scheme when instantiated with SSB hash functions. To prove security, we will first replace the key $K$ by a key $K^{(L)}$ that statistically binds the $L$-th bit of the pre-image. The index hiding property guarantees that this change goes unnoticed. Now for every hash value $y=\mathrm{H}\left(K^{(L)}, D\right)$, it holds that $(K, L, y, D[L]) \in \mathcal{L}$ while $(K, L, y, 1-D[L]) \notin \mathcal{L}$. We can now rely on the security of witness encryption to argue that $\operatorname{Enc}\left(\left(K^{(L)}, L, y, 1-D[L]\right), m_{1-D[L]}\right)$ computationally hides the message $m_{1-D[L]}$.

Working with DDH. The above described scheme relies on a witness encryption scheme for the language $\mathcal{L}$. We note that witness encryption for general NP languages is only known under strong assumptions such as graded encodings [GGSW13] or indistinguishability obfuscation [GGH $\left.{ }^{+} 13 \mathrm{~b}\right]$. Nevertheless, the aforementioned laconic OT scheme does not need full power of general witness
encryption. In particular, we will leverage the fact that hash proof systems [CS02] can be used to construct statistical witness encryption schemes for specific languages [GGSW13]. Towards this end, we will carefully craft an SSB hash function that is hash proof system friendly, that is, allows for a hash proof system (or statistical witness encryption) for the language $\mathcal{L}$ required above. Our construction of the HPS-friendly SSB hash is based on the Decisional Diffie-Hellman assumption and is inspired from a construction by Okamoto et al. [OPWW15].

We will briefly outline our HPS-friendly SSB hash below. We strongly encourage the reader to see Section 4.2 for the full construction or see [?] for a simplified construction.

Let $\mathbb{G}$ be a (multiplicative) cyclic group of order $p$ generated by a generator $g$. A hashing key is of the form $\hat{\mathbf{H}}=g^{\mathbf{H}}$ (the exponentiation is done component-wisely), where the matrix $\mathbf{H} \in \mathbb{Z}_{p}^{2 \times 2 \lambda}$ is chosen uniformly at random. The hash function of $\mathbf{x} \in \mathbb{Z}_{p}^{2 \lambda}$ is computed as $\mathbf{H}(\hat{\mathbf{H}}, \mathbf{x})=\hat{\mathbf{H}}^{\mathbf{x}} \in \mathbb{G}^{2}$ (where $\left(\hat{\mathbf{H}}^{\mathbf{x}}\right)_{i}=\prod_{k=1}^{2 \lambda} \hat{\mathbf{H}}_{i, k}^{x_{k}}$, hence $\hat{\mathbf{H}}^{\mathbf{x}}=g^{\mathbf{H x}}$ ). The binding key $\hat{\mathbf{H}}^{(i)}$ is of the form $\hat{\mathbf{H}}^{(i)}=g^{\mathbf{A}+\mathbf{T}}$, where $\mathbf{A} \in \mathbb{Z}_{p}^{2 \times 2 \lambda}$ is a random rank 1 matrix, and $\mathbf{T} \in \mathbb{Z}_{p}^{2 \times 2 \lambda}$ is a matrix with zero entries everywhere, except that $\mathbf{T}_{2, i}=1$.

Now we describe a witness encryption scheme (Enc, Dec) for the language $\mathcal{L}=\{(\hat{\mathbf{H}}, i, \hat{\mathbf{y}}, b) \mid \exists \mathbf{x} \in$ $\mathbb{Z}_{p}^{2 \lambda}$ s.t. $\hat{\mathbf{H}}^{\mathbf{x}}=\hat{\mathbf{y}}$ and $\left.x_{i}=b\right\} . \operatorname{Enc}((\hat{\mathbf{H}}, i, \hat{\mathbf{y}}, b), m)$ first sets

$$
\hat{\mathbf{H}}^{\prime}=\binom{\hat{\mathbf{H}}}{g^{\mathbf{e}_{i}^{\top}}} \in \mathbb{G}^{3 \times 2 \lambda}, \hat{\mathbf{y}}^{\prime}=\binom{\hat{\mathbf{y}}}{g^{b}} \in \mathbb{G}^{3},
$$

where $\mathbf{e}_{i} \in \mathbb{Z}_{p}^{2 \lambda}$ is the $i$-th unit vector. It then picks a random $\mathbf{r} \in \mathbb{Z}_{p}^{3}$ and computes a ciphertext $c=\left(\left(\left(\hat{\mathbf{H}}^{\prime}\right)^{\top}\right)^{\mathbf{r}},\left(\left(\hat{\mathbf{y}}^{\prime}\right)^{\top}\right)^{\mathbf{r}} \oplus m\right)$. To decrypt a ciphertext $c=(\hat{\mathbf{h}}, z)$ given a witness $\mathbf{x} \in \mathbb{Z}_{p}^{2 \lambda}$, we compute $m=z \oplus \hat{\mathbf{h}}^{\mathbf{x}}$. It is easy to check correctness. For the security proof, see Section 4.3.

### 2.1.2 Bootstrapping Laconic OT

We will now provide a bootstrapping technique that constructs a laconic OT scheme with arbitrary compression factor from one with factor- 2 compression. Let $\ell O T_{\text {const }}$ denote a laconic OT scheme with factor-2 compression.

Bootstrapping the Hash Function via a Merkle Tree. A binary Merkle tree is a natural way to construct hash functions with an arbitrary compression factor from two-to-one hash functions, and this is exactly the route we pursue. A binary Merkle tree is constructed as follows: The database is split into blocks of $\lambda$ bits, each of which forms the leaf of the tree. An interior node is computed as the hash value of its two children via a two-to-one hash function. This structure is defined recursively from the leaves to the root. When we reach the root node (of $\lambda$ bits), its value is defined to be the (succinct) hash value or digest of the entire database. This procedure defines the hash function.

The next step is to define the laconic OT algorithms Send and Receive for the above hash function. Our first observation is that given the digest, the sender can transfer specific messages corresponding to the values of the left and right children of the root (via $2 \lambda$ executions of $\ell O T_{\text {const }}$. Send). Hence, a naive approach for the sender is to output $\ell O T_{\text {const }}$ encryptions for the path of nodes from the root to the leaf of interest. This approach runs into an immediate issue because to compute $\ell O T_{\text {const }}$ encryptions at any layer other than the root, the sender needs to know the value at that internal node. However, in the scheme a sender only knows the value of the root and nothing else.

Traversing the Merkle Tree via Garbled Circuits. Our main idea to make the above naive idea work is via an interesting usage of garbled circuits. At a high level, the sender will output a sequence of garbled circuits (one per layer of the tree) to transfer messages corresponding to the path from the root to the leaf containing the $L$-th bit, so that the receiver can traverse the Merkle tree from the root to the leaf as illustrated in Figure 2.


Above GCircuit is a circuit garbling procedure, which garbles the circuit $\ell O T_{\text {const }}$.Send $\left(\right.$ crs,$\cdot$, Keys $\left.{ }^{2}\right)$ using input keys Keys ${ }^{1}$ (see Section 5.1.1for the definition of garbled circuits).

Figure 2: The Bootstrapping Step
In more detail, the construction works as follows: The Send algorithm outputs $\ell O T_{\text {const }}$ encryptions using the root digest and a collection of garbled circuits, one per layer of the Merkle tree. The $i$-th circuit has a bit $b$ hardwired in it, which specifies whether the path should go to the left or right child at the $i$-th layer. It takes as input a pair of sibling nodes (node ${ }_{0}$, node ${ }_{1}$ ) along the path at layer $i$ and outputs $\ell O T_{\text {const }}$ encryptions corresponding to nodes on the path at layer $i+1$ w.r.t. node $_{b}$ as the hash value. Conceptually, the circuit computes $\ell O T_{\text {const }}$ encryptions for the next layer.

The $\ell O T_{\text {const }}$ encryptions at the root encrypt the input keys of the first garbled circuit. In the garbled circuit at layer $i$, the messages being encrypted/sent correspond to the input keys of the garbled circuit at layer $i+1$. The last circuit takes two sibling leaves as input which contains $D[L]$, and outputs $\ell O T_{\text {const }}$ encryptions of $m_{0}$ and $m_{1}$ corresponding to location $L$ (among the $2 \lambda$ locations).

Given a laconic OT ciphertext, which consists of $\ell O T_{\text {const }}$ ciphertexts w.r.t. the root digest and a sequence of garbled circuits, the receiver can traverse the Merkle tree as follows. First he runs $\ell O T_{\text {const }}$. Receive for the $\ell O T_{\text {const }}$ ciphertexts using as witness the children of the root, obtaining the input labels corresponding to these to be fed into the first garbled circuit. Next, he uses the input labels to evaluate the first garbled circuit, obtaining $\ell O T_{\text {const }}$ ciphertexts for the second layer. He then runs $\ell O T_{\text {const }}$.Receive again for these ciphertexts using as witness the children of the second node on the path. This procedure continues till the last layer.

Security of the construction can be established using the sender security of $\ell O T_{\text {const }}$.Receive and simulation based security of the circuit garbling scheme.

Extension. Finally, for our RAM applications we need a slightly stronger primitive which we call updatable laconic OT that additionally allows for modifications/writes to the database while ensuring that the digest is updated in a consistent manner. The construction sketched in this paragraph can be modified to support this stronger notion. For a detailed description of this notion refer to Section 3.2.

### 2.2 Non-interactive Secure Computation on Large Inputs

The Circuit Setting. This is the most straightforward application of laconic OT. We will provide a non-interactive secure computation protocol where the receiver $R$, holding a large database $D$, publishes a short encoding of it such that any sender $S$, with private input $x$, can send a single message to reveal $C(x, D)$ to $R$. Here, $C$ is the circuit being evaluated.

Recall the garbled circuit based approach to non-interactive secure computation, where $R$ can publish the first message of a two-message oblivious transfer (OT) for his input $D$, and the sender responds with a garbled circuit for $C[x, \cdot]$ (with hardcoded input $x$ ) and sends the input labels corresponding to $D$ via the second OT message. The downside of this protocol is that $R$ 's public message grows with the size of $D$, which could be substantially large.

We resolve this issue via our new primitive laconic OT. In our protocol, $R$ 's first message is the digest digest of his large database $D$. Next, the sender generates the garbled circuit for $C[x, \cdot]$ as before. It also transfers the labels for each location of $D$ via laconic OT Send messages. Hence, by efficiency requirements of laconic OT, the length of $R$ 's public message is independent of the size of $D$. Moreover, sender privacy against a semi-honest receiver follows directly from the sender privacy of laconic OT and security of garbled circuits. To achieve receiver privacy, we can enhance the laconic OT with receiver privacy (discussed in Section 3.1).

The RAM Setting. This is the RAM version of the above application where $S$ holds a RAM program $P$ and $R$ holds a large database $D$. As before, we want that (1) the length of $R$ 's first message is independent of $|D|$, (2) $R$ 's first message can be published and used by multiple senders, (3) the database is persistent for a sequence of programs for every sender, and (4) the computational complexity of both $S$ and $R$ per program execution grows only with running time of the corresponding program. For this application, we only achieve unprotected memory access (UMA) security against a corrupt receiver, i.e., the memory access pattern in the execution of $P^{D}(x)$ is leaked to the receiver. We achieve full security against a corrupt sender.

For simplicity, consider a read-only program such that each CPU step outputs the next location to be read based on the value read from last location. At a high level, since we want the sender's complexity to grow only with the running time $t$ of the program, we cannot create a garbled circuit that takes $D$ as input. Instead, we would go via the garbled RAM based approaches where we have a sequence of $t$ garbled circuits where each circuit executes one CPU step. A CPU step circuit takes the current CPU state and the last bit read from the database $D$ as input and outputs an updated state and a new location to be read. The new location would be read from the database and fed into the next CPU step. The most non-trivial part in all garbled RAM constructions is being able to compute the correct labels for the next circuit based on the value of $D[L]$, where $L$ is the location being read. Since we are working with garbled circuits, it is crucial for security that the receiver does not learn two labels for any input wire. We solve this issue via laconic OT as follows.

For the simpler case of sender security, $R$ publishes the short digest of $D$, which is fed into the first garbled circuit and this digest is passed along the sequence of garbled circuits. When a circuit wants to read a location $L$, it outputs the laconic OT ciphertexts which encrypt the input keys for the next circuit and use digest of $D$ as the hash value. ${ }^{9}$ Security against a corrupt receiver follows from the sender security of laconic OT and security of garbled circuits. To achieve security against a corrupt sender, $R$ does not publishes digest in the clear. Instead, the labels for digest for the first circuit are transferred to $R$ via regular OT.

Note that the garbling time of the sender as well as execution time of the receiver will grow only with the running time of the program. This follows from the efficiency requirements of laconic OT.

Above, we did not describe how we deal with general programs that also write to the database or memory. We achieve this via updatable laconic OT (for definition see Section 3.2), This allows for transferring the labels for updated digest (corresponding to the updated database) to the next circuit. For a formal description of our scheme for general RAM programs, see Section 6.

### 2.3 Multi-Hop Homomorphic Encryption for RAM Programs

Our model and problem - a bit more formally. We consider a scenario where a server $S$, holding an input $x$, publishes a public key pk and an encryption ct of $x$ under pk. Now this ciphertext is passed on to a client $Q$ that will compute a (possibly private) program $P$ accessing memory $D$ on the value encrypted in ct, obtaining another ciphertext ct'. Finally, we want that the server can use its secret key to recover $P^{D}(x)$ from the ciphertext ct' and $\widetilde{D}$, where $\widetilde{D}$ is an encrypted form of $D$ that has been previously provided to $S$ in a one-time setup phase. More generally, the computation could be performed by multiple clients $Q_{1}, \ldots, Q_{n}$. In this case, each client is required to place a pre-processed version of its database $\widetilde{D}_{i}$ with the server during setup. The computation itself could be performed in different sequences of the clients (for different extensions of the model, see Section 7.1). Examples of two such computation paths are shown in Figure 1.

For security, we want IND-CPA security for server's input $x$. For honest clients, we want program-privacy as well as data-privacy, i.e., the evaluation does not leak anything beyond the output of the computation even when the adversary corrupts the server and any subset of the clients. We note that data-privacy is rather easy to achieve via encryption and ORAM. Hence we focus on the challenges of achieving UMA security for honest clients, i.e., the adversary is allowed to learn the database $D$ as well as memory access pattern of $P$ on $D$.

UMA secure multi-hop scheme. We first build on the ideas from non-interactive secure computation for RAM programs. Every client first passes its database to the server. Then in every round, the server sends an OT message for input $x$. We assume for simplicity that every client has an up-to-date digest of its own database. Next, the first client $Q_{1}$ generates a garbled program for $P_{1}$, say $\mathrm{ct}_{1}$ and sends it to $Q_{2}$. Here, the garbled program consists of $t_{1}\left(t_{1}\right.$ is the running time of $P_{1}$ ) garbled circuits accessing $D_{1}$ via laconic OT as described in the previous application. Now, $Q_{2}$ appends its garbled program for $P_{2}$ to the end of $\mathrm{ct}_{1}$ and generates $\mathrm{ct}_{2}$ consisting of $\mathrm{ct}_{1}$ and new garbled program. Note that $P_{2}$ takes the output of $P_{1}$ as input and hence, the output keys of the

[^5]last garbled circuit of $P_{1}$ have to be compatible with the input keys of the first garbled circuit of $P_{2}$ and so on. If we continue this procedure, after the last client $Q_{n}$, we get a sequence of garbled circuits where the first $t_{1}$ circuits access $D_{1}$, the next set accesses from $D_{2}$ and so on. Finally, the server $S$ can evaluate the sequence of garbled circuits given $D_{1}, \ldots, D_{n}$. It is easy to see that correctness holds. But we have no security for clients.

The issue is similar to the issue pointed out by [GHV10] for the case of multi-hop garbled circuits. If the client $Q_{i-1}$ colludes with the server, then they can learn both input labels for the garbled program of $Q_{i}$. To resolve this issue it is crucial that $Q_{i}$ re-randomizes the garbled circuits provided by $Q_{i-1}$. For this we rely on re-randomizable garbled circuits provided by [GHV10], where given a garbled circuit anyone can re-garble it such that functionality of the original circuit is preserved while the re-randomized garbled circuit is unrecognizable even to the party who generated it. In our protocol we use re-randomizable garbled circuits but we stumble upon the following issue.

Recall that in the RAM application above, a garbled circuit outputs the laconic OT ciphertexts corresponding to the input keys of the next circuit. Hence, the input keys of the $(\tau+1)$-th circuit have to be hardwired inside the $\tau$-th circuit. Since all of these circuits will be re-randomized for security, for correctness we require that we transform the hardwired keys in a manner consistent with the future re-randomization. But for security, $Q_{i-1}$ does not know the randomness that will be used by $Q_{i}$.

Our first idea to resolve this issue is as follows: The circuits generated by $Q_{i-1}$ will take additional inputs $s_{i}, \ldots s_{n}$ which are the randomness used by future parties for their re-randomization procedure. Since we are in the non-interactive setting, we cannot run an OT protocol between clients $Q_{i-1}$ and later clients. We resolve this issue by putting the first message of OT for $s_{j}$ in the public key of client $Q_{j}$ and client $Q_{i-1}$ will send the OT second messages along with $\mathrm{ct}_{i-1}$. We do not want the clients' public keys to grow with the running time of the programs, hence, we think of $s_{j}$ as PRF keys and each circuit re-randomization will invoke the PRF on a unique input.

The above approach causes a subtle issue in the security proof. Suppose, for simplicity, that client $Q_{i}$ is the only honest client. When arguing security, we want to simulate all the garbled circuits in $\mathrm{ct}_{i}$. To rely on the security of re-randomization, we need to replace the output of the PRF with key $s_{i}$ with uniform random values but this key is fed as input to the circuits of the previous clients. We note that this is not a circularity issue but makes arguing security hard. We solve this issue as follows: Instead of feeding in PRF keys directly to the garbled circuits, we feed in corresponding outputs of the PRF. We generate the PRF output via a bunch of PRF circuits that take the PRF keys as input (see Figure 3). Now during simulation, we will first simulate these PRF circuits, followed by the simulation of the main circuits. We describe the scheme formally in Section 7.3.

## 3 Laconic Oblivious Transfer

In this section, we will introduce a primitive we call Laconic OT (or, $\ell O T$ for short). We will start by describing laconic OT and then provide an extension of it to the notion of updatable laconic OT.


Figure 3: One step circuit for $P_{i}$ along with the attached PRF circuits generated by $Q_{i}$.

### 3.1 Laconic OT

Definition 3.1 (Laconic OT). A laconic OT ( (OOT) scheme syntactically consists of four algorithms crsGen, Hash, Send and Receive.

- crs $\leftarrow \operatorname{crsGen}\left(1^{\lambda}\right)$. It takes as input the security parameter $1^{\lambda}$ and outputs a common reference string crs.
- (digest, $\hat{\mathrm{D}}) \leftarrow$ Hash $(\mathrm{crs}, D)$. It takes as input a common reference string crs and a database $D \in\{0,1\}^{*}$ and outputs a digest digest of the database and a state $\hat{\mathrm{D}}$.
- e $\leftarrow$ Send(crs, digest, $L, m_{0}, m_{1}$ ). It takes as input a common reference string crs, a digest digest, a database location $L \in \mathbb{N}$ and two messages $m_{0}$ and $m_{1}$ of length $\lambda$, and outputs a ciphertext e.
- $m \leftarrow \operatorname{Receive}^{\hat{\mathrm{D}}}(\mathrm{crs}, \mathrm{e}, L)$. This is a RAM algorithm with random read access to $\hat{\mathrm{D}}$. It takes as input a common reference string crs, a ciphertext e , and a database location $L \in \mathbb{N}$. It outputs a message $m$.

We require the following properties of an $\ell O T$ scheme (crsGen, Hash, Send, Receive).

- Correctness: We require that it holds for any database $D$ of size at most $M=\operatorname{poly}(\lambda)$ for any polynomial function poly $(\cdot)$, any memory location $L \in[M]$, and any pair of messages

$$
\left(m_{0}, m_{1}\right) \in\{0,1\}^{\lambda} \times\{0,1\}^{\lambda} \text { that }
$$

$$
\operatorname{Pr}\left[m=m_{D[L]} \left\lvert\, \begin{array}{c|cl}
\text { crs } & \leftarrow \operatorname{crsGen}\left(1^{\lambda}\right) \\
& \text { digest, }, \mathrm{D}) & \leftarrow \operatorname{Hash}(\operatorname{crs}, D) \\
\mathrm{e} & \leftarrow \operatorname{Send}\left(\operatorname{crs}, \operatorname{digest}, L, m_{0}, m_{1}\right) \\
m & \leftarrow \operatorname{Receive}^{\mathrm{D}}(\operatorname{crs}, \mathrm{e}, L)
\end{array}\right.\right]=1,
$$

where the probability is taken over the random choices made by crsGen and Send.

- Sender Privacy Against Semi-Honest Receivers: There exists a PPT simulator lOTSim such that the following holds. For any database $D$ of size at most $M=\operatorname{poly}(\lambda)$ for any polynomial function poly $(\cdot)$, any memory location $L \in[M]$, and any pair of messages $\left(m_{0}, m_{1}\right) \in$ $\{0,1\}^{\lambda} \times\{0,1\}^{\lambda}$, let crs $\leftarrow \operatorname{crsGen}\left(1^{\lambda}\right)$ and digest $\leftarrow \operatorname{Hash}(\operatorname{crs}, D)$. Then it holds that

$$
\left(\mathrm{crs}, \operatorname{Send}\left(\mathrm{crs}, \text { digest, } L, m_{0}, m_{1}\right)\right) \stackrel{\mathrm{c}}{\approx}\left(\mathrm{crs}, \ell \mathrm{OTSim}\left(D, L, m_{D[L]}\right)\right) .
$$

- Efficiency Requirement: The length of digest is a fixed polynomial in $\lambda$ independent of the size of the database; we will assume for simplicity that $\mid$ digest $\mid=\lambda$. Moreover, the algorithm Hash runs in time $|D| \cdot \operatorname{poly}(\log |D|, \lambda)$, Send and Receive run in time poly $(\log |D|, \lambda)$.
Receiver Privacy. In the above definition, we do not require receiver privacy as opposed to standard oblivious transfer, namely, no security guarantee is provided against a cheating (semihonest) sender. This is mostly for ease of exposition. We would like to point out that adding receiver privacy (i.e., standard simulation based security against a semi-honest sender) to laconic OT can be done in a straightforward way. Instead of sending digest directly from the receiver to the sender and sending e back to the receiver, the two parties compute Send together via a two-round secure 2 PC protocol, where the input of the receiver is digest and the input of the sender is $\left(L, m_{0}, m_{1}\right)$, and only the receiver obtains the output e. This can be done using standard two-message OT and garbled circuits.

Multiple executions of Send that share the same digest. Notice that since the common reference string is public (i.e., not chosen by the simulator), the sender can involve Send function multiple times while still ensuring that security can be argued from the above definition (for the case of single execution) via a standard hybrid argument.

It will be convenient to use the following shorthand notations (generalizing the above notions) to run laconic OT for every single element in a database. Let Keys $=\left(\left(\mathrm{Key}_{1,0}, \mathrm{Key}_{1,1}\right), \ldots,\left(\mathrm{Key}_{M, 0}, \mathrm{Key}_{M, 1}\right)\right)$ be a list of $M=|D|$ key-pairs, where each key is of length $\lambda$. Then we will define

$$
\operatorname{Send}(\text { crs, digest, } \mathrm{Keys})=\left(\text { Send }\left(\mathrm{crs}, \operatorname{digest}, 1, \operatorname{Key}_{1,0}, \operatorname{Key}_{1,1}\right), \ldots, \text { Send }\left(\mathrm{crs}, \operatorname{digest}, M, \operatorname{Key}_{M, 0}, \operatorname{Key}_{M, 1}\right)\right) .
$$

Likewise, for a vector $\mathrm{e}=\left(\mathrm{e}_{1}, \ldots, \mathrm{e}_{M}\right)$ of ciphertexts define

$$
\operatorname{Receive}^{\hat{\mathrm{D}}}(\text { crs }, \mathrm{e})=\left(\operatorname{Receive}^{\hat{\mathrm{D}}}\left(\mathrm{crs}, \mathrm{e}_{1}, 1\right), \ldots, \operatorname{Receive}^{\hat{\mathrm{D}}}\left(\text { crs }, \mathrm{e}_{M}, M\right)\right) .
$$

Similarly, let Labels $=\operatorname{Keys}_{D}=\left(\operatorname{Key}_{1, D[1]}, \ldots, \operatorname{Key}_{M, D[M]}\right)$, and define

$$
\ell O T S i m(\operatorname{crs}, D, \text { Labels })=\left(\ell \mathrm{OTSim}\left(\operatorname{crs}, D, 1, \operatorname{Key}_{1, D[1]}\right), \ldots, \ell \mathrm{OTSim}\left(\operatorname{crs}, D, M, \operatorname{Key}_{M, D[M]}\right) .\right.
$$

By the sender security for multiple executions, we have that

$$
(\text { crs }, \text { Send }(\text { crs }, \text { digest }, \text { Keys })) \stackrel{c}{\approx}(\text { crs, } \ell \mathrm{OTSim}(\text { crs }, D, \text { Labels })) .
$$

### 3.2 Updatable Laconic OT

For our applications, we will need a version of laconic OT for which the receiver's short commitment digest to his database can be updated quickly (in time much smaller than the size of the database) when a bit of the database changes. We call this primitive supporting this functionality updatable laconic OT and define more formally below. At a high level, updatable laconic OT comes with an additional pair of algorithms SendWrite and ReceiveWrite which transfer the keys for an updated digest digest* to the receiver. For convenience, we will define ReceiveWrite such that it also performs the write in D.

Definition 3.2 (Updatable Laconic OT). An updatable laconic OT (updatable $\ell O T$ ) scheme consists of algorithms crsGen, Hash, Send, Receive as per Definition 3.1 and additionally two algorithms SendWrite and ReceiveWrite with the following syntax.

- $\mathrm{e}_{\mathrm{w}} \leftarrow$ SendWrite $\left(\mathrm{crs}\right.$, digest, $\left.L, b,\left\{m_{j, 0}, m_{j, 1}\right\}_{j=1}^{\mid \text {digest } \mid}\right)$. It takes as input the common reference string crs, a digest digest, a location $L \in \mathbb{N}$, a bit $b \in\{0,1\}$ to be written, and |digest| pairs of messages $\left\{m_{j, 0}, m_{j, 1}\right\}_{j=1}^{\mid \text {digest } \mid}$, where each $m_{j, c}$ is of length $\lambda$. And it outputs a ciphertext $\mathrm{e}_{\mathrm{w}}$.
- $\left\{m_{j}\right\}_{j=1}^{\mid \text {digest } \mid} \leftarrow$ ReceiveWrite $^{\hat{\mathrm{D}}}\left(\mathrm{crs}, L, b, \mathrm{e}_{\mathrm{w}}\right)$. This is a RAM algorithm with random read/write access to $\hat{\mathrm{D}}$. It takes as input the common reference string crs, a location $L$, a bit $b \in\{0,1\}$ and a ciphertext $\mathrm{e}_{\mathrm{w}}$. It updates the state $\hat{\mathrm{D}}$ (such that $D[L]=b$ ) and outputs messages $\left\{m_{j}\right\}_{j=1}^{\mid \text {digest } \mid}$.

We require the following properties on top of properties of a laconic OT scheme.

- Correctness With Regard To Writes: For any database $D$ of size at most $M=\operatorname{poly}(\lambda)$ for any polynomial function poly $(\cdot)$, any memory location $L \in[M]$, any bit $b \in\{0,1\}$, and any messages $\left\{m_{j, 0}, m_{j, 1}\right\}_{j=1}^{\mid \text {digest } \mid}$ of length $\lambda$, the following holds. Let $D^{*}$ be identical to $D$, except that $D^{*}[L]=b$,

where the probability is taken over the random choices made by crsGen and SendWrite. Furthermore, we require that the execution of ReceiveWrite ${ }^{\hat{D}}$ above updates $\hat{\mathrm{D}}$ to $\hat{\mathrm{D}}^{*}$. (Note that digest is included in $\hat{\mathrm{D}}$, hence digest is also updated to digest*.)
- Sender Privacy Against Semi-Honest Receivers With Regard To Writes: There exists a PPT simulator $\ell$ OTSimWrite such that the following holds. For any database $D$ of size at most $M=\operatorname{poly}(\lambda)$ for any polynomial function poly $(\cdot)$, any memory location $L \in[M]$, any bit $b \in\{0,1\}$, and any messages $\left\{m_{j, 0}, m_{j, 1}\right\}_{j=1}^{\text {|digest } \mid}$ of length $\lambda$, let $\operatorname{crs} \leftarrow \operatorname{crsGen}\left(1^{\lambda}\right)$,
(digest, $\hat{\mathrm{D}}) \leftarrow \operatorname{Hash}(\mathrm{crs}, D)$, and $\left(\right.$ digest $\left.^{*}, \hat{\mathrm{D}}^{*}\right) \leftarrow \operatorname{Hash}\left(\mathrm{crs}, D^{*}\right)$, where $D^{*}$ is identical to $D$ except that $D^{*}[L]=b$. Then it holds that

$$
\left.\stackrel{c}{\approx}\left(\begin{array}{l}
\left.\mathrm{crs}, \text { SendWrite }\left(\mathrm{crs}, \text { digest, } L, b,\left\{m_{j, 0}, m_{j, 1}\right\}_{j=1}^{\mid \text {digest } \mid}\right)\right) \\
\left(\mathrm{crs}, \ell \mathrm{OTSimWrite}\left(\operatorname{crs}, D, L, b,\left\{m_{j, \text { digest }}^{j}\right.\right.\right. \\
\}_{j \in[\mid \text { digest } \mid]}
\end{array}\right)\right) .
$$

- Efficiency Requirements: We require that both SendWrite and ReceiveWrite run in time poly $(\log |D|, \lambda)$.


## 4 Laconic Oblivious Transfer with Factor-2 Compression

In this section, based on the DDH assumption we will construct a laconic OT scheme for which the hash function Hash compresses a database of length $2 \lambda$ into a digest of length $\lambda$. We would refer to this primitive as laconic OT with factor-2 compression. We note that, subsequent to our work, the factor-2 compression construction has been simplified by Döttling and Garg [?] (another alternative simplification can be obtained using [?]). We refer the reader to [?] for the simpler construction and preserve the older construction here.

We will first construct the following two primitives as building blocks: (1) a somewhere statistically binding (SSB) hash function, and (2) a hash proof system that allows for proving knowledge of preimage bits for this SSB hash function. We will then present the $\ell O T$ scheme with factor- 2 compression in Section 4.4.

### 4.1 Somewhere Statistically Binding Hash Functions and Hash Proof Systems

In this section, we give definitions of somewhere statistically binding (SSB) hash functions [HW15] and hash proof systems [CS98]. For simplicity, we will only define SSB hash functions that compress $2 \lambda$ values in the domain into $\lambda$ bits. The more general definition works analogously.

Definition 4.1 (Somewhere Statistically Binding Hashing). An SSB hash function SSBH consists of three algorithms crsGen, bindingCrsGen and Hash with the following syntax.

- crs $\leftarrow \operatorname{crsGen}\left(1^{\lambda}\right)$. It takes the security parameter $\lambda$ as input and outputs a common reference string crs.
- crs $\leftarrow$ bindingCrsGen $\left(1^{\lambda}, i\right)$. It takes as input the security parameter $\lambda$ and an index $i \in[2 \lambda]$, and outputs a common reference string crs.
- $y \leftarrow$ Hash (crs, $x$ ). For some domain $\mathfrak{D}$, it takes as input a common reference string crs and a string $x \in \mathfrak{D}^{2 \lambda}$, and outputs a string $y \in\{0,1\}^{\lambda}$.

We require the following properties of an SSB hash function.

- Statistically Binding at Position $i$ : For every $i \in[2 \lambda]$ and an overwhelming fraction of crs in the support of bindingCrsGen $\left(1^{\lambda}, i\right)$ and every $x \in \mathfrak{D}^{2 \lambda}$, we have that (crs, Hash(crs, $x$ )) uniquely determines $x_{i}$. More formally, for all $x^{\prime} \in \mathfrak{D}^{2 \lambda}$ such that $x_{i} \neq x_{i}^{\prime}$ we have that Hash $\left(\right.$ crs,$\left.x^{\prime}\right) \neq \operatorname{Hash}(\operatorname{crs}, x)$.
- Index Hiding: It holds for all $i \in[2 \lambda]$ that $\operatorname{crsGen}\left(1^{\lambda}\right) \stackrel{c}{\approx} \operatorname{bindingCrsGen}\left(1^{\lambda}, i\right)$, i.e., common reference strings generated by crsGen and bindingCrsGen are computationally indistinguishable.

Next, we define hash proof systems [CS98] that are designated verifier proof systems that allow for proving that the given problem instance in some language. We give the formal definition as follows.

Definition 4.2 (Hash Proof System). Let $\mathcal{L}_{z} \subseteq \mathcal{M}_{z}$ be an NP-language residing in a universe $\mathcal{M}_{z}$, both parametrized by some parameter $z$. Moreover, let $\mathcal{L}_{z}$ be characterized by an efficiently computable witness-relation $\mathcal{R}$, namely, for all $x \in M_{z}$ it holds that $x \in \mathcal{L}_{z} \Leftrightarrow \exists w: \mathcal{R}(x, w)=1$. A hash proof system HPS for $\mathcal{L}_{z}$ consists of three algorithms KeyGen, $\mathrm{H}_{\text {public }}$ and $\mathrm{H}_{\text {secret }}$ with the following syntax.

- ( $\mathrm{pk}, \mathrm{sk}) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}, z\right)$ : Takes as input the security parameter $\lambda$ and a parameter $z$, and outputs a public-key and secret key pair (pk, sk).
- $y \leftarrow \mathrm{H}_{\text {public }}(\mathrm{pk}, x, w)$ : Takes as input a public key pk , an instance $x \in \mathcal{L}_{z}$, and a witness $w$, and outputs a value $y$.
- $y \leftarrow \mathrm{H}_{\text {secret }}(\mathrm{sk}, x)$ : Takes as input a secret key sk and an instance $x \in M_{z}$, and outputs a value $y$.

We require the following properties of a hash proof system.

- Perfect Completeness: For every $z$, every (pk, sk) in the support of $\operatorname{KeyGen}\left(1^{\lambda}, z\right)$, and every $x \in \mathcal{L}_{z}$ with witness $w$ (i.e., $\mathcal{R}(x, w)=1$ ), it holds that

$$
\mathrm{H}_{\text {public }}(\mathrm{pk}, x, w)=\mathrm{H}_{\text {secret }}(\mathrm{sk}, x) .
$$

- Perfect Soundness: For every $z$ and every $x \in M_{z} \backslash \mathcal{L}_{z}$, let $(\mathrm{pk}, \mathrm{sk}) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}, z\right)$, then it holds that

$$
\left(z, \mathrm{pk}, \mathrm{H}_{\text {secret }}(\mathrm{sk}, x)\right) \equiv(z, \mathrm{pk}, u),
$$

where $u$ is distributed uniformly random in the range of $\mathrm{H}_{\text {secret }}$. Here, $\equiv$ denotes distributional equivalence.

### 4.2 HPS-friendly SSB Hashing

In this section, we will construct an HPS-friendly SSB hash function that supports a hash proof system. In particular, there is a hash proof system that enables proving that a certain bit of the pre-image of a hash-value has a certain fixed value (in our case, either 0 or 1 ).

We start with some notations. Let $(\mathbb{G}, \cdot)$ be a cyclic group of order $p$ with generator $g$. Let $\mathbf{M} \in \mathbb{Z}_{p}^{m \times n}$ be a matrix. We will denote by $\hat{\mathbf{M}}=g^{\mathbf{M}} \in \mathbb{G}^{m \times n}$ the element-wise exponentiation of $g$ with the elements of $\mathbf{M}$. We also define $\hat{\mathbf{L}}=\hat{\mathbf{H}}^{\mathbf{M}} \in \mathbb{G}^{m \times k}$, where $\hat{\mathbf{H}} \in \mathbb{G}^{m \times n}$ and $\mathbf{M} \in \mathbb{Z}_{p}^{n \times k}$ as follows: Each element $\hat{\mathbf{L}}_{i, j}=\prod_{k=1}^{n} \hat{\mathbf{H}}_{i, k}^{\mathbf{M}_{k, j}}$ (intuitively this operation corresponds to matrix multiplication in the exponent). This is well-defined and efficiently computable.

Computational Assumptions. In the following, we first define the computational problems on which we will base the security of our HPS-friendly SSB hash function.

Definition 4.3 (The Decisional Diffie-Hellman (DDH) Problem). Let $(\mathbb{G}, \cdot)$ be a cyclic group of prime order $p$ and with generator $g$. Let $a, b, c$ be sampled uniformly at random from $\mathbb{Z}_{p}$ (i.e., $a, b, c \stackrel{\$}{\leftarrow} \mathbb{Z}_{p}$ ). The DDH problem asks to distinguish the distributions $\left(g, g^{a}, g^{b}, g^{a b}\right)$ and $\left(g, g^{a}, g^{b}, g^{c}\right)$.

Definition 4.4 (Matrix Rank Problem). Let $m, n$ be integers and let $\mathbb{Z}_{p}^{m \times n ; r}$ be the set of all $m \times n$ matrices over $\mathbb{Z}_{p}$ with rank $r$. Further, let $1 \leq r_{1}<r_{2} \leq \min (m, n)$. The goal of the matrix rank
 where $\mathbf{M}_{1} \stackrel{\$}{\leftarrow} \mathbb{Z}_{p}^{m \times n ; r_{1}}$ and $\mathbf{M}_{2} \stackrel{\$}{\leftarrow} \mathbb{Z}_{p}^{m \times n ; r_{2}}$.

In a recent result by Villar [Vil12] it was shown that the matrix rank problem can be reduced almost tightly to the DDH problem.

Theorem 4.5 ( [Vil12] Theorem 1, simplified). Assume there exists a PPT distinguisher $\mathcal{D}$ that solves MatrixRank( $\left.\mathbb{G}, m, n, r_{1}, r_{2}\right)$ problem with advantage $\epsilon$. Then, there exists a PPT distinguisher $\mathcal{D}^{\prime}$ (running in almost time as $\mathcal{D}$ ) that solves $\operatorname{DDH}$ problem over $\mathbb{G}$ with advantage at least $\frac{\epsilon}{\left|\log _{2}\left(r_{2} / r_{1}\right)\right\rangle}$.

We next give the construction of an HPS-friendly SSB hash function.
Construction. Our construction builds on the scheme of Okamoto et al. [OPWW15]. We will not delve into the details of their scheme and directly jump into our construction.

Let $n$ be an integer such that $n=2 \lambda$, and let $(\mathbb{G}, \cdot)$ be a cyclic group of order $p$ and with generator $g$. Let $\mathbf{T}_{i} \in \mathbb{Z}_{p}^{2 \times n}$ be a matrix which is zero everywhere except the $i$-th column, and the $i$-th column is equal to $\mathbf{t}=(0,1)^{\top}$. The three algorithms of the SSB hash function are defined as follows.

- $\operatorname{crsGen}\left(1^{\lambda}\right)$ : Pick a uniformly random matrix $\mathbf{H} \stackrel{\$}{\leftarrow} \mathbb{Z}_{p}^{2 \times n}$ and output $\hat{\mathbf{H}}=g^{\mathbf{H}}$.
- bindingCrsGen $\left(1^{\lambda}, i\right)$ : Pick a uniformly random vector $\left(w_{1}, w_{2}\right)^{\top}=\mathbf{w} \stackrel{\$}{\stackrel{\mathbb{Z}}{\mathbb{Z}}}{ }_{p}^{2}$ with the restriction that $w_{1}=1$, pick a uniformly random vector $\mathbf{a} \stackrel{\$}{\leftarrow} \mathbb{Z}_{p}^{n}$ and set $\mathbf{A} \leftarrow \mathbf{w} \cdot \mathbf{a}^{\top}$. Set $\mathbf{H} \leftarrow \mathbf{T}_{i}+\mathbf{A}$ and output $\hat{\mathbf{H}}=g^{\mathbf{H}}$.
- Hash (crs, $\mathbf{x})$ : Parse $\mathbf{x}$ as a vector in $\mathfrak{D}^{n}\left(\mathfrak{D}=\mathbb{Z}_{p}\right)$ and parse crs $=\hat{\mathbf{H}}$. Compute $\mathbf{y} \in \mathbb{G}^{2}$ as $\mathbf{y}=\hat{\mathbf{H}}^{\mathbf{x}}$. Parse $\mathbf{y}$ as a binary string and output the result.

Compression. Notice that we can get factor two compression for an input space $\{0,1\}^{2 \lambda}$ by restricting the domain to $\mathfrak{D}^{\prime}=\{0,1\} \subset \mathfrak{D}$. The input length $n=2 \lambda$, where $\lambda$ is set to be twice the number of bits in the bit representation of a group element in $\mathbb{G}$. In the following we will assume that $n=2 \lambda$ and that the bit-representation size of a group element in $\mathbb{G}$ is $\frac{\lambda}{2}$.

We will first show that the distributions crsGen $\left(1^{\lambda}\right)$ and bindingCrsGen $\left(1^{\lambda}, i\right)$ are computationally indistinguishable for every index $i \in[n]$, given that the DDH problem is computationally hard in the group $\mathbb{G}$.

Lemma 4.6 (Index Hiding). Assume that the MatrixRank $(\mathbb{G}, 2, n, 1,2)$ problem is hard. Then the distributions crsGen $\left(1^{\lambda}\right)$ and bindingCrsGen $\left(1^{\lambda}, i\right)$ are computationally indistinguishable, for every $i \in[n]$.

Proof. Assume there exists a PPT distinguisher $\mathcal{D}$ that distinguishes the distributions $\operatorname{crsGen}\left(1^{\lambda}\right)$ and bindingCrsGen $\left(1^{\lambda}, i\right)$ with non-negligible advantage $\epsilon$. We will construct a PPT distinguisher $\mathcal{D}^{\prime}$ that distinguishes MatrixRank $(\mathbb{G}, 2, n, 1,2)$ with non-negligible advantage.

The distinguisher $\mathcal{D}^{\prime}$ does the following on input $\hat{\mathbf{M}} \in \mathbb{G}^{2 \times n}$. It computes $\hat{\mathbf{H}} \in \mathbb{G}^{2 \times n}$ as elementwise multiplication of $\hat{\mathbf{M}}$ and $g^{\mathbf{T}_{i}}$ and runs $\mathcal{D}$ on $\hat{\mathbf{H}}$. If $\mathcal{D}$ outputs crsGen, then $\mathcal{D}^{\prime}$ outputs rank 2 , otherwise $\mathcal{D}^{\prime}$ outputs rank 1.

We will now show that $\mathcal{D}^{\prime}$ also has non-negligible advantage. Write $\mathcal{D}^{\prime}$ 's input as $\hat{\mathbf{M}}=g^{\mathbf{M}}$. If $\mathbf{M}$ is chosen uniformly random with rank 2 , then $\mathbf{M}$ is uniform in $\mathbb{Z}_{p}^{2 \times n}$ with overwhelming probability. Hence with overwhelming probability, $\mathbf{M}+\mathbf{T}_{i}$ is also distributed uniformly random and it follows that $\hat{\mathbf{H}}=g^{\mathbf{M}+\mathbf{T}_{i}}$ is uniformly random in $\mathbb{G}^{2 \times n}$ which is identical to the distribution generated by $\operatorname{crs} \operatorname{Gen}\left(1^{\lambda}\right)$. On the other hand, if $\mathbf{M}$ is chosen uniformly random with rank 1 , then there exists a vector $\mathbf{w} \in \mathbb{Z}_{p}^{2}$ such that each column of $\mathbf{M}$ can be written as $a_{i} \cdot \mathbf{w}$. We can assume that the first element $w_{1}$ of $\mathbf{w}$ is 1 , since the case $w_{1}=0$ happens only with probability $1 / p=\operatorname{negl}(\lambda)$ and if $w_{1} \neq 0$ we can replace all $a_{i}$ by $a_{i}^{\prime}=a_{i} \cdot w_{1}$ and replace $w_{i}$ by $w_{i}^{\prime}=\frac{w_{i}}{w_{1}}$. Thus,
 distributed, hence $\hat{\mathbf{H}}$ is identical to the distribution generated by bindingCrsGen $\left(1^{\lambda}, i\right)$. Since $\mathcal{D}$ can distinguish the distributions crsGen $\left(1^{\lambda}\right)$ and binding $\operatorname{CrsGen}\left(1^{\lambda}, i\right)$ with non-negligible advantage $\epsilon, \mathcal{D}^{\prime}$ can distinguish $\operatorname{MatrixRank}(\mathbb{G}, 2, n, 1,2)$ with advantage $\epsilon-$ negl $(\lambda)$, which contradicts the hardness of MatrixRank ( $\mathbb{G}, 2, n, 1,2)$.

A corollary of Lemma 4.6 is that for all $i, j \in[n]$ the distributions bindingCrsGen $\left(1^{\lambda}, i\right)$ and bindingCrsGen $\left(1^{\lambda}, j\right)$ are indistinguishable, stated as follows.

Corollary 4.7. Assume the MatrixRank $(\mathbb{G}, 2, n, 1,2)$ problem is computationally hard. Then it holds for all $i, j \in[n]$ that bindingCrsGen $\left(1^{\lambda}, i\right)$ and bindingCrsGen $\left(1^{\lambda}, j\right)$ are computationally indistinguishable.

We next show that if the common reference string crs $=\hat{\mathbf{H}}$ is generated by bindingCrsGen $\left(1^{\lambda}, i\right)$, then the hash value $\mathrm{Hash}(\mathrm{crs}, \mathbf{x})$ is statistically binded to $x_{i}$.

Lemma 4.8 (Statistically Binding at Position $i$ ). For every $i \in[n]$, every $\mathbf{x} \in \mathbb{Z}_{p}^{n}$, and all choices of crs in the support of bindingCrsGen $\left(1^{\lambda}, i\right)$ we have that for every $\mathbf{x}^{\prime} \in \mathbb{Z}_{p}^{n}$ such that $x_{i}^{\prime} \neq x_{i}$, Hash (crs, $\mathbf{x}) \neq$ Hash (crs, $\mathbf{x}^{\prime}$ ).
Proof. We first write crs as $\hat{\mathbf{H}}=g^{\mathbf{H}}=g^{\mathbf{w} \cdot \mathbf{a}^{\top}+\mathbf{T}_{i}}$ and $\operatorname{Hash}(\mathrm{crs}, \mathbf{x})$ as $\operatorname{Hash}(\hat{\mathbf{H}}, \mathbf{x})=g^{\mathbf{y}}=g^{\mathbf{H} \cdot \mathbf{x}}$. Thus, by taking the discrete logarithm with basis $g$ our task is to demonstrate that there exists a unique $x_{i}$ from $\mathbf{H}=\mathbf{w} \cdot \mathbf{a}^{\top}+\mathbf{T}_{i}$ and $\mathbf{y}=\mathbf{H} \cdot \mathbf{x}$. Observe that

$$
\begin{aligned}
\mathbf{y} & =\mathbf{H} \cdot \mathbf{x}=\left(\mathbf{w} \cdot \mathbf{a}^{\top}+\mathbf{T}_{i}\right) \cdot \mathbf{x}=\mathbf{w} \cdot\langle\mathbf{a}, \mathbf{x}\rangle+\mathbf{T}_{i} \cdot \mathbf{x} \\
& =\binom{1}{w_{2}} \cdot\langle\mathbf{a}, \mathbf{x}\rangle+\binom{0}{1} \cdot x_{i},
\end{aligned}
$$

where $\langle\mathbf{a}, \mathbf{x}\rangle$ is the inner product of $\mathbf{a}$ and $\mathbf{x}$. If $\mathbf{a} \neq \mathbf{0}$, then we can use any non-zero element of $\mathbf{a}$ to compute $w_{2}$ from $\mathbf{H}$, and recover $x_{i}$ by computing $x_{i}=y_{2}-w_{2} \cdot y_{1}$; otherwise $\mathbf{a}=\mathbf{0}$, so $x_{i}=y_{2}$.

### 4.3 A Hash Proof System for Knowledge of Preimage Bits

In this section, we give our desired hash proof systems. In particular, we need a hash proof system for membership in a subspace of a vector space. In our proof we need the following technical lemma.
Lemma 4.9. Let $\mathbf{M} \in \mathbb{Z}_{p}^{m \times n}$ be a matrix. Let $\operatorname{colsp}(\mathbf{M})=\left\{\mathbf{M} \cdot \mathbf{x} \mid \mathbf{x} \in \mathbb{Z}_{p}^{n}\right\}$ be its column space, and $\operatorname{rowsp}(\mathbf{M})=\left\{\mathbf{x}^{\top} \cdot \mathbf{M} \mid \mathbf{x} \in \mathbb{Z}_{p}^{m}\right\}$ be its row space. Assume that $\mathbf{y} \in \mathbb{Z}_{p}^{m}$ and $\mathbf{y} \notin \operatorname{colsp}(\mathbf{M})$. Let $\mathbf{r} \stackrel{\$}{\leftarrow} \mathbb{Z}_{p}^{m}$ be chosen uniformly at random. Then it holds that

$$
\left(\mathbf{M}, \mathbf{y}, \mathbf{r}^{\top} \mathbf{M}, \mathbf{r}^{\top} \mathbf{y}\right) \equiv\left(\mathbf{M}, \mathbf{y}, \mathbf{r}^{\top} \mathbf{M}, u\right)
$$

where $u \stackrel{\$}{\leftarrow} \mathbb{Z}_{p}$ is distributed uniformly and independently of $\mathbf{r}$. Here, $\equiv$ denotes distributional equivalence.
Proof. For any $\mathbf{t} \in \operatorname{rowsp}(\mathbf{M})$ and $s \in \mathbb{Z}_{p}$, consider following linear equation system

$$
\left\{\begin{array}{l}
\mathbf{r}^{\top} \mathbf{M}=\mathbf{t} \\
\mathbf{r}^{\top} \mathbf{y}=s
\end{array}\right.
$$

Let $\mathcal{N}$ be the left null space of $\mathbf{M}$. We know that $\mathbf{y} \notin \operatorname{colsp}(\mathbf{M})$, hence $\mathbf{M}$ has rank $\leq m-1$, therefore $\mathcal{N}$ has dimension $\geq 1$. Let $\mathbf{r}_{0}$ be an arbitrary solution for $\mathbf{r}^{\top} \mathbf{M}=\mathbf{t}$, and let $\mathbf{n}$ be a vector in $\mathcal{N}$ such that $\mathbf{n}^{\top} \mathbf{y} \neq \mathbf{0}$ (there must be such a vector since $\mathbf{y} \notin \operatorname{colsp}(\mathbf{M})$ ). Then there exists a solution $\mathbf{r}$ for the above linear equation system, that is,

$$
\mathbf{r}=\mathbf{r}_{0}+\left(\mathbf{n}^{\top} \mathbf{y}\right)^{-1} \cdot\left(s-\mathbf{r}_{0}^{\top} \mathbf{y}\right) \cdot \mathbf{n}
$$

where $\left(\mathbf{n}^{\top} \mathbf{y}\right)^{-1}$ is the multiplicative inverse of $\mathbf{n}^{\top} \mathbf{y}$ in $\mathbb{Z}_{p}$. Then two cases arise: (i) column vectors of ( $\mathbf{M} \mathbf{y}$ ) are full-rank, or (ii) not. In this first case, there is a unique solution for $\mathbf{r}$. In the second case the solution space has the same size as the left null space of ( $\mathbf{M} \mathbf{y}$ ). Therefore, in both cases, the number of solutions for $\mathbf{r}$ is the same for every $(\mathbf{t}, s)$ pair.

As $\mathbf{r}$ is chosen uniformly at random, all pairs $(\mathbf{t}, s) \in \operatorname{rowsp}(\mathbf{M}) \times \mathbb{Z}_{p}$ have the same probability of occurrence and the claim follows.

Construction. Fix a matrix $\hat{\mathbf{H}} \in \mathbb{G}^{2 \times n}$ and an index $i \in[n]$. We will construct a hash proof system HPS $=\left(\right.$ KeyGen, $\left.H_{\text {public }}, H_{\text {secret }}\right)$ for the following language $\mathcal{L}_{\hat{\mathbf{H}}, i}$ :

$$
\mathcal{L}_{\hat{\mathbf{H}}, i}=\left\{(\hat{\mathbf{y}}, b) \in \mathbb{G}^{2} \times\{0,1\} \mid \exists \mathbf{x} \in \mathbb{Z}_{p}^{n} \text { s.t. } \hat{\mathbf{y}}=\hat{\mathbf{H}}^{\mathbf{x}} \text { and } x_{i}=b\right\} .
$$

Note that in our hash proof system we only enforce that a single specified bit is $b$, where $b \in\{0,1\}$. However, our hash proof system does not place any requirement on the value used at any of the other locations. In fact the values used at the other locations may actually be from the full domain $\mathfrak{D}$ (i.e., $\mathbb{Z}_{p}$ ). Observe that the formal definition of the language $\mathcal{L}_{\hat{\mathbf{H}}, i}$ above incorporates this difference in how the honest computation of the hash function is performed and what the hash proof system is supposed to prove.

For ease of exposition, it will be convenient to work with a matrix $\hat{\mathbf{H}}^{\prime} \in \mathbb{G}_{p}^{3 \times n}$ :

$$
\hat{\mathbf{H}}^{\prime}=\binom{\hat{\mathbf{H}}}{g^{\mathbf{e}_{i}^{\top}}},
$$

where $\mathbf{e}_{i} \in \mathbb{Z}_{p}^{n}$ is the $i$-th unit vector, with all elements equal to zero except the $i^{\text {th }}$ one which is equal to one.

- KeyGen $\left(1^{\lambda},(\hat{\mathbf{H}}, i)\right)$ : Choose $\mathbf{r} \stackrel{\$}{\leftarrow} \mathbb{Z}_{p}^{3}$ uniformly at random. Compute $\hat{\mathbf{h}}=\left(\left(\hat{\mathbf{H}}^{\prime}\right)^{\top}\right)^{\mathbf{r}}$. Set $\mathrm{pk}=\hat{\mathbf{h}}$ and $\mathrm{sk}=\mathbf{r}$. Output ( $\mathrm{pk}, \mathrm{sk}$ ).
- $\mathrm{H}_{\text {public }}(\mathrm{pk},(\hat{\mathbf{y}}, b), \mathbf{x}):$ Parse pk as $\hat{\mathbf{h}}$. Compute $\hat{z}=\left(\hat{\mathbf{h}}^{\top}\right)^{\mathbf{x}}$ and output $\hat{z}$.
- $\mathrm{H}_{\text {secret }}(\mathrm{sk},(\hat{\mathbf{y}}, b))$ : Parse sk as $\mathbf{r}$ and $\operatorname{set} \hat{\mathbf{y}}^{\prime}=\binom{\hat{\mathbf{y}}}{g^{b}}$. Compute $\hat{z}=\left(\left(\hat{\mathbf{y}}^{\prime}\right)^{\top}\right)^{\mathbf{r}}$ and output $\hat{z}$.

Lemma 4.10. For every matrix $\hat{\mathbf{H}} \in \mathbb{G}^{2 \times n}$ and every $i \in[n]$, HPS is a hash proof system for the language $\mathcal{L}_{\hat{\mathbf{H}}, i}$.
Proof. Let $\hat{\mathbf{H}}=g^{\mathbf{H}}, \mathbf{r}=\left(\mathbf{r}^{*}, r_{3}\right)$ where $\mathbf{r}^{*} \in \mathbb{Z}_{p}^{2}$. Let $\mathbf{y}^{\prime}:=\log _{g} \hat{\mathbf{y}}^{\prime}, \mathbf{y}:=\log _{g} \hat{\mathbf{y}}, \mathbf{H}^{\prime}:=\log _{g} \hat{\mathbf{H}}^{\prime}$, $\mathbf{h}:=\log _{g} \hat{\mathbf{h}}$.

For perfect correctness, we need to show that for every $i \in[n]$, every $\hat{\mathbf{H}} \in \mathbb{G}^{2 \times n}$, and every (pk, sk) in the support of $\operatorname{Key} \operatorname{Gen}\left(1^{\lambda},\left(\hat{\mathbf{H}}^{\prime}, i\right)\right)$, if $(\hat{\mathbf{y}}, b) \in \mathcal{L}_{\hat{\mathbf{H}}, i}$ and $\mathbf{x}$ is a witness for membership (i.e., $\hat{\mathbf{y}}=\hat{\mathbf{H}}^{\mathbf{x}}$ and $\left.x_{i}=b\right)$, then it holds that $\mathrm{H}_{\text {public }}(\mathrm{pk},(\hat{\mathbf{y}}, b), \mathbf{x})=\mathrm{H}_{\text {secret }}(\mathrm{sk},(\hat{\mathbf{y}}, b))$.

To simplify the argument, we again consider the statement under the discrete logarithm with basis $g$. Then it holds that

$$
\begin{aligned}
& \log _{g}\left(\mathrm{H}_{\text {secret }}(\mathrm{sk},(\hat{\mathbf{y}}, b))\right) \\
= & \log _{g}\left(\left(\left(\hat{\mathbf{y}}^{\prime}\right)^{\top}\right)^{\mathbf{r}}\right)=\left\langle\mathbf{y}^{\prime}, \mathbf{r}\right\rangle=\left\langle\mathbf{y}, \mathbf{r}^{*}\right\rangle+b \cdot r_{3} \\
= & \left\langle\mathbf{H} \cdot \mathbf{x}, \mathbf{r}^{*}\right\rangle+x_{i} \cdot r_{3}=\left\langle\mathbf{H}^{\prime} \mathbf{x}, \mathbf{r}\right\rangle=\left\langle\left(\mathbf{H}^{\prime}\right)^{\top} \mathbf{r}, \mathbf{x}\right\rangle \\
= & \langle\mathbf{h}, \mathbf{x}\rangle=\log _{g}\left(\left(\hat{\mathbf{h}}^{\top}\right)^{\mathbf{x}}\right) \\
= & \log _{g}\left(\mathrm{H}_{\text {public }}(\mathrm{pk},(\hat{\mathbf{y}}, b), \mathbf{x})\right) .
\end{aligned}
$$

For perfect soundness, let $(\mathrm{pk}, \mathrm{sk}) \leftarrow \operatorname{KeyGen}\left(1^{\lambda},\left(\hat{\mathbf{H}}^{\prime}, i\right)\right)$. We will show that if $(\hat{\mathbf{y}}, b) \notin \mathcal{L}_{\hat{\mathbf{H}}, i}$, then $\mathrm{H}_{\text {secret }}(\mathrm{sk},(\hat{\mathbf{y}}, b))$ is distributed uniformly random in the range of $\mathrm{H}_{\text {secret }}$, even given $\hat{\mathbf{H}}$, $i$, and pk. Again under the discrete logarithm, this is equivalent to showing that $\left\langle\mathbf{y}^{\prime}, \mathbf{r}\right\rangle$ is distributed uniformly random given $\mathbf{H}^{\prime}$ and $\mathbf{h}=\left(\mathbf{H}^{\prime}\right)^{\top} \mathbf{r}$.

Note that we can re-write the language $\mathcal{L}_{\hat{\mathbf{H}}, i}=\left\{(\hat{\mathbf{y}}, b) \in \mathbb{G}^{2} \times \mathbb{Z}_{p} \mid \exists \mathbf{x} \in \mathbb{Z}_{p}^{n}\right.$ s.t. $\left.\mathbf{H}^{\prime} \mathbf{x}=\mathbf{y}^{\prime}\right\}$. It follows that if $(\hat{\mathbf{y}}, b) \notin \mathcal{L}_{\hat{\mathbf{H}}, i}$, then $\mathbf{y}^{\prime} \notin \operatorname{span}\left(\mathbf{H}^{\prime}\right)$. Now it follows directly from Lemma 4.9 that

$$
\mathbf{r}^{\top} \mathbf{y}^{\prime} \equiv u
$$

given $\mathbf{H}^{\prime}$ and $\mathbf{r}^{\top} \mathbf{H}^{\prime}$, where $u$ is distributed uniformly random. This concludes the proof.

Remark 4.11. While proving the security of our applications based on the above hash-proof system, we would generate $\hat{\mathbf{H}}$ to be the output of bindingCrsGen $\left(1^{\lambda}, i\right)$ and use the property that if $(\hat{\mathbf{y}}, b) \in$ $\mathcal{L}_{\hat{\mathbf{H}}, i}$, then $(\hat{\mathbf{y}},(1-b)) \notin \mathcal{L}_{\hat{\mathbf{H}}, i}$. This follows directly from Lemma 4.8 (that is, $\hat{\mathbf{H}}$ and $\hat{\mathbf{y}}$ uniquely fixes $x_{i}$ ).

### 4.4 The Laconic OT Scheme

We are now ready to put the pieces together and provide our $\ell O T$ scheme with factor- 2 compression.

Construction. Let SSBH = (SSBH.crsGen, SSBH.bindingCrsGen, SSBH.Hash) be the HPS-friendly SSB hash function constructed in Section 4.2 with domain $\mathfrak{D}=\mathbb{Z}_{p}$. Notice that we achieve factor-2 compression (namely, compressing $2 \lambda$ bits into $\lambda$ bits) by restricting the domain from $\mathfrak{D}^{n}$ to $\{0,1\}^{n}$ in our laconic OT scheme. Also, abstractly let the associated hash proof system be HPS $=($ HPS.KeyGen, HPS. Heblic , HPS. Hecret $)$ for the language

$$
\mathcal{L}_{\mathrm{crs}, i}=\left\{(\text { digest }, b) \in\{0,1\}^{\lambda} \times\{0,1\} \mid \exists D \in \mathfrak{D}^{2 \lambda}: \operatorname{SSBH} . H a s h(\operatorname{crs}, D)=\text { digest and } D[i]=b\right\} .
$$

Recall that the bit-representation size of a group element of $\mathbb{G}$ is $\frac{\lambda}{2}$, hence the language defined above is the same as the one defined in Section 4.3.

Now we construct the laconic OT scheme $\ell O T=$ (crsGen, Hash, Send, Receive) as follows.

- $\operatorname{crsGen}\left(1^{\lambda}\right)$ : Compute crs $\leftarrow \operatorname{SSBH} . \operatorname{crsGen}\left(1^{\lambda}\right)$ and output crs.
- Hash (crs, $\left.D \in\{0,1\}^{2 \lambda}\right)$ :
digest $\leftarrow$ SSBH.Hash (crs, $D$ )
$\hat{\mathrm{D}} \leftarrow(D$, digest $)$
Output (digest, $\hat{\text { D }}$ )
- Send(crs, digest, $\left.L, m_{0}, m_{1}\right)$ :

Let HPS be the hash-proof system for the language $\mathcal{L}_{\text {crs }, L}$
$(\mathrm{pk}, \mathrm{sk}) \leftarrow \operatorname{HPS} . \operatorname{KeyGen}\left(1^{\lambda},(\mathrm{crs}, L)\right)$
$c_{0} \leftarrow m_{0} \oplus$ HPS. $\mathrm{H}_{\text {secret }}($ sk, $($ digest, 0$))$
$c_{1} \leftarrow m_{1} \oplus \operatorname{HPS} . \mathrm{H}_{\text {secret }}($ sk, $($ digest, 1$))$
Outpute $=\left(\mathrm{pk}, c_{0}, c_{1}\right)$

- Receive ${ }^{\hat{\mathrm{D}}}(\mathrm{crs}, \mathrm{e}, L)$ :

Parse e $=\left(\mathrm{pk}, c_{0}, c_{1}\right)$
Parse $\hat{\mathrm{D}}=(D$, digest), and set $b \leftarrow D[L]$.
$m \leftarrow c_{b} \oplus$ HPS. Heblic $_{\text {pk, }}($ digest,$\left.b), D\right)$
Output $m$
We will now show that $\ell O T$ is a laconic OT protocol with factor-2 compression, i.e., it has compression factor 2, and satisfies the correctness and sender privacy requirements. First notice that SSBH.Hash is factor-2 compressing, so Hash also has compression factor 2. We next argue correctness and sender privacy in Lemmas 4.12 and 4.13, respectively.

Lemma 4.12. Given that HPS satisfies the correctness property, the lOT scheme also satisfies the correctness property.

Proof. Fix a common reference string crs in the support of $\operatorname{crsGen}\left(1^{\lambda}\right)$, a database string $D \in$ $\{0,1\}^{2 \lambda}$ and an index $L \in[2 \lambda]$. For any crs, $D, L$ such that $D[L]=b$, let digest $=$ Hash (crs, $\left.D\right)$. Then it clearly holds that (digest, $b) \in \mathcal{L}_{\text {crs }, L}$. Thus, by the correctness property of the hash proof system HPS it holds that

$$
\text { HPS. } \mathrm{H}_{\text {secret }}(\text { sk, }(\text { digest }, b))=\operatorname{HPS} . \mathrm{H}_{\text {public }}(\text { pk, }(\text { digest }, b), D) .
$$

By the construction of Send(crs, digest, $\left.L, m_{0}, m_{1}\right), c_{b}=m_{b} \oplus$ HPS. $H_{\text {secret }}($ sk, (digest, $b)$ ). Hence the output $m$ of Receive ${ }^{\hat{D}}$ (crs, e, $L$ ) is

$$
\begin{aligned}
m & =c_{b} \oplus \text { HPS. } \mathrm{H}_{\text {public }}(\text { pk },(\text { digest }, b), D) \\
& =m_{b} \oplus \text { HPS. } \mathrm{H}_{\text {secret }}(\text { sk, }(\text { digest }, b)) \oplus \mathrm{HPS} . \mathrm{H}_{\text {public }}(\text { pk, }(\text { digest }, b), D) \\
& =m_{b} .
\end{aligned}
$$

Lemma 4.13. Given that SSBH is index-hiding and has the statistically binding property and that HPS is sound, then the lOT scheme satisfies sender privacy against semi-honest receiver.

Proof. We first construct the simulator $\ell$ OTSim.

```
\ellOTSim(crs, D,L, mol[L]):
    digest }\leftarrow\mathrm{ SSBH.Hash(crs, D)
    Let HPS be the hash-proof system for the language }\mp@subsup{\mathcal{L}}{\mathrm{ crs,L}}{
    (pk, sk)}\leftarrowHPS.KeyGen(1^, (crs,L)
    co\leftarrow m m[L]}\oplus\mathrm{ HPS.H.H secret (sk, (digest, 0))
    c}\leftarrow\leftarrow\mp@subsup{m}{D[L]}{}\oplusHPS.\mp@subsup{H}{\mathrm{ secret }}{}(\mathrm{ sk, (digest, 1))
    Output (pk, co, c. )
```

For any database $D$ of size at most $M=\operatorname{poly}(\lambda)$ for any polynomial function poly $(\cdot)$, any memory location $L \in[M]$, and any pair of messages $\left(m_{0}, m_{1}\right) \in\{0,1\}^{\lambda} \times\{0,1\}^{\lambda}$, let $\operatorname{crs} \leftarrow \operatorname{crsGen}\left(1^{\lambda}\right)$ and digest $\leftarrow \operatorname{Hash}(\mathrm{crs}, D)$. Then we will prove that the two distributions (crs, Send(crs, digest, $\left.L, m_{0}, m_{1}\right)$ ) and (crs, $\ell \mathrm{OTSim}\left(\mathrm{crs}, D, L, m_{D[L]}\right)$ ) are computationally indistinguishable. Consider the following hybrids.

- Hybrid 0: This is the real experiment, namely (crs, Send(crs, digest, $\left.L, m_{0}, m_{1}\right)$ ).
- Hybrid 1: Same as hybrid 0 , except that crs is computed by crs $\leftarrow \operatorname{SSBH}$.bindingCrsGen $\left(1^{\lambda}, L\right)$.
- Hybrid 2: Same as hybrid 1, except that $c_{1-D[L]}$ is computed by $c_{1-D[L]} \leftarrow m_{D[L]} \oplus \mathrm{HPS} . \mathrm{H}_{\text {secret }}$ (sk, (digest, $1-D[L])$ ). That is, both $c_{0}$ and $c_{1}$ encrypt the same message $m_{D[L]}$.
- Hybrid 3: Same as hybrid 2, except that crs is computed by crs $\leftarrow \operatorname{SSBH} . c r s G e n\left(1^{\lambda}\right)$. This is the simulated experiment, namely (crs, $\ell \mathrm{OTSim}\left(\mathrm{crs}, D, L, m_{D[L]}\right)$ ).

Indistinguishability of hybrid 0 and hybrid 1 follows directly from Lemma 4.6, as we replace the distribution of crs from SSBH.crsGen $\left(1^{\lambda}\right)$ to $\operatorname{SSBH}$.bindingCrsGen $\left(1^{\lambda}, L\right)$. Indistinguishability of hybrids 2 and 3 also follows from Lemma 4.6, as we replace the distribution of crs from SSBH.bindingCrsGen $\left(1^{\lambda}, L\right)$ back to SSBH.crsGen $\left(1^{\lambda}\right)$.

We will now show that hybrids 1 and 2 are identically distributed. Since crs is in the support of SSBH.bindingCrsGen $\left(1^{\lambda}, i\right)$ and digest $=$ SSBH.Hash $($ crs, $D)$, by Lemma 4.8 it holds that (digest, $1-$ $D[L]) \notin \mathcal{L}_{\mathrm{crs}, L}$. By the soundness property of the hash-proof system HPS, it holds that

$$
\left(\mathrm{crs}, L, \mathrm{pk}, \mathrm{HPS} . \mathrm{H}_{\text {secret }}(\mathrm{sk},(\operatorname{digest}, 1-D[L]))\right) \equiv(\mathrm{crs}, L, \mathrm{pk}, u),
$$

for a uniformly random $u$. Furthermore, $c_{D[L]}$ can be computed by $m_{D[L]} \oplus H P S . H_{\text {public }}$ (pk, (digest, $D[L]), D)$. Hence

$$
\begin{aligned}
& \left(\text { crs, } L, \mathrm{pk}, m_{D[L]} \oplus \mathrm{HPS} . \mathrm{H}_{\text {secret }}(\mathrm{sk},(\text { digest, } 1-D[L])), c_{D[L]}\right) \\
\equiv & \left(\mathrm{crs}, L, \mathrm{pk}, u, c_{D[L]}\right) \\
\equiv & \left(\mathrm{crs}, L, \mathrm{pk}, m_{1-D[L]} \oplus \mathrm{HPS} . \mathrm{H}_{\text {secret }}(\text { sk, }(\text { digest }, 1-D[L])), c_{D[L]}\right) .
\end{aligned}
$$

This concludes the proof.

## 5 Construction of Updatable Laconic OT

In this section, we will construct an updatable laconic OT that supports a hash function that allows for compression from an input (database) of size an arbitrary polynomial in $\lambda$ to $\lambda$ bits. As every updatable laconic OT protocol is also a (standard) laconic OT protocol, we will only construct the former. Our main technique in this construction, is the use of garbled circuits to bootstrap a laconic OT with factor-2 compression into one with an arbitrary compression factor.

Below in Section 5.1 we describe some background on the primitives needed for realizing our laconic OT construction. Then we will give the construction of laconic OT along with its correctness and security proofs in Sections 5.2 and 5.3 , respectively.

### 5.1 Background

In this section we recall the needed background of garbled circuits and Merkle trees.

### 5.1.1 Garbled Circuits

Garbled circuits were first introduced by Yao [Yao82] (see Lindell and Pinkas [LP09] and Bellare et al. [BHR12] for a detailed proof and further discussion). A circuit garbling scheme GC is a tuple of PPT algorithms (GCircuit, Eval). Very roughly GCircuit is the circuit garbling procedure and Eval the corresponding evaluation procedure. Looking ahead, each individual wire $w$ of the circuit being garbled will be associated with two labels, namely key ${ }_{w, 0}$, key $_{w, 1}$.

- $\widetilde{\mathrm{C}} \leftarrow$ GCircuit $\left(1^{\lambda}, \mathrm{C},\left\{\operatorname{key}_{w, b}\right\}_{w \in \operatorname{inp}(\mathrm{C}), b \in\{0,1\}}\right):$ GCircuit takes as input a security parameter $\lambda$, a circuit C , and a set of labels key ${ }_{w, b}$ for all the input wires $w \in \operatorname{inp}(\mathrm{C})$ and $b \in\{0,1\}$. This procedure outputs a garbled circuit $\widetilde{\mathrm{C}}$.
- $y \leftarrow \operatorname{Eval}\left(\widetilde{\mathrm{C}},\left\{\operatorname{key}_{w, x_{w}}\right\}_{w \in \operatorname{inp}(\mathrm{C})}\right)$ : Given a garbled circuit $\widetilde{\mathrm{C}}$ and a garbled input represented as a sequence of input labels $\left\{\text { key }_{w, x_{w}}\right\}_{w \in \operatorname{inp}(\mathrm{C})}$, Eval outputs $y$.

Terminology of Keys and Labels. We note that, in the rest of the paper, we use the notation Keys to refer to both the secret values sampled for wires and the notation Labels to refer to exactly one of them. In other words, generation of garbled circuit involves Keys while computation itself depends just on Labels. Let Keys $=\left(\left(\right.\right.$ key $_{1,0}$, key $\left._{1,1}\right), \ldots,\left(\right.$ key $_{n, 0}$, key $\left.\left.{ }_{n, 1}\right)\right)$ be a list of $n$ key-pairs, we denote $\mathrm{Keys}_{x}$ for a string $x \in\{0,1\}^{n}$ to be a list of labels (key ${ }_{1, x_{1}}, \ldots$, key $_{n, x_{n}}$ ).

Correctness. For correctness, we require that for any circuit C and input $x \in\{0,1\}^{m}$ (here $m$ is the input length to C ) we have that:

$$
\operatorname{Pr}\left[\mathrm{C}(x)=\operatorname{Eval}\left(\widetilde{\mathrm{C}},\left\{\operatorname{key}_{w, x_{w}}\right\}_{w \in \operatorname{inp}(\mathrm{C})}\right)\right]=1
$$

where $\widetilde{\mathrm{C}} \leftarrow \operatorname{GCircuit}\left(1^{\lambda}, \mathrm{C},\left\{\text { key }_{w, b}\right\}_{w \in \operatorname{inp}(\mathrm{C}), b \in\{0,1\}}\right)$.
Security. For security, we require that there is a PPT simulator CircSim such that for any $\mathrm{C}, x$, and uniformly random keys $\left\{\operatorname{key}_{w, b}\right\}_{w \in \operatorname{inp}(C), b \in\{0,1\}}$, we have that

$$
\left(\widetilde{\mathrm{C}},\left\{\operatorname{key}_{w, x_{w}}\right\}_{w \in \operatorname{inp}(\mathrm{C})}\right) \stackrel{\mathrm{c}}{\approx} \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}, y\right)
$$

where $\widetilde{\mathrm{C}} \leftarrow \operatorname{GCircuit}\left(1^{\lambda}, \mathrm{C},\left\{\text { key }_{w, b}\right\}_{w \in \operatorname{inp}(\mathrm{C}), b \in\{0,1\}}\right)$ and $y=\mathrm{C}(x)$.

### 5.1.2 Merkle Tree

In this section we briefly review Merkle trees. A Merkle tree is a hash based data structure that generically extend the domain of a hash function. The following description will be tailored to the hash function of the laconic OT scheme that we will present in Section 5.2. Given a two-to-one hash function Hash : $\{0,1\}^{2 \lambda} \rightarrow\{0,1\}^{\lambda}$, we can use a Merkle tree to construct a hash function that compresses a database of an arbitrary (a priori unbounded polynomial in $\lambda$ ) size to a $\lambda$-bit string. Now we briefly illustrate how to compress a database $D \in\{0,1\}^{M}$ (assume for ease of exposition that $M=2^{d} \cdot \lambda$ ). First, we partition $D$ into strings of length $2 \lambda$; we call each string a leaf. Then we use Hash to compress each leaf into a new string of length $\lambda$; we call each string a node. Next, we bundle the new nodes in pairs of two and call these pairs siblings, i.e., each pair of siblings is a string of length $2 \lambda$. We then use Hash again to compress each pair of siblings into a new node of size $\lambda$. We continue the process till a single node of size $\lambda$ is obtained. This process forms a binary tree structure, which we refer to as a Merkle tree. Looking ahead, the hash function of the laconic OT scheme has output ( $\hat{\mathrm{D}}$, digest), where $\hat{\mathrm{D}}$ is the entire Merkle tree, and digest is the root of the tree.

A Merkle tree has the following property. In order to verify that a database $D$ with hash root digest has a certain value $b$ at a location $L$ (namely, $D[L]=b$ ), there is no need to provide the entire Merkle tree. Instead, it is sufficient to provide a path of siblings from the Merkle tree root to the leaf that contains location $L$. It can then be easily verified if the hash values from the leaf to the root are correct.

Moreover, a Merkle tree can be updated in the same fashion when the value at a certain location of the database is updated. Instead of recomputing the entire tree, we only need to recompute the nodes on the path from the updated leaf to the root. This can be done given the path of siblings from the root to the leaf.

### 5.2 Construction

We will now provide our construction to bootstrap an $\ell O T$ scheme with factor-2 compression into an updatable $\ell O T$ scheme with an arbitrary compression factor, which can compress a database of an arbitrary (a priori unbounded polynomial in $\lambda$ ) size.

Overview. We first give an overview of the construction. Consider a database $D \in\{0,1\}^{M}$ such that $M=2^{d} \cdot \lambda$. Given a laconic OT scheme with factor- 2 compression (denoted as $\ell O T_{\text {const }}$ ), we will first use a Merkle tree to obtain a hash function with arbitrary (polynomial) compression factor. As described in Section 5.1.2, the Hash function of the updatable $\ell O T$ scheme will have an output ( $\hat{\mathrm{D}}$, digest), where $\hat{\mathrm{D}}$ is the entire Merkle tree, and digest is the root of the tree.

In the Send algorithm, suppose we want to send a message depending on a bit $D[L]$, we will follow the natural approach of traversing the Merkle tree layer by layer until reaching the leaf containing $L$. In particular, $L$ can be represented as $L=\left(b_{1}, \ldots, b_{d-1}, t\right)$, where $b_{1}, \ldots, b_{d-1}$ are bits representing the path from the root to the leaf containing location $L$, and $t \in[2 \lambda]$ is the position within the leaf. The Send algorithm first takes as input the root digest of the Merkle tree, and it will generate a chain of garbled circuits, which would enable the receiver to traverse the Merkle tree from the root to the leaf. And upon reaching the leaf, the receiver will be able to evaluate the last garbled circuit and retrieve the message corresponding to the $t$-th bit of the leaf.

We briefly explain the chain of garbled circuits as follows. The chain consists of $d-1$ traversing circuits along with a reading circuit. Every traversing circuit takes as input a pair of siblings $\mathrm{sbl}=\left(\mathrm{sbl}_{0}, \mathrm{sbl}_{1}\right)$ at a certain layer of the Merkle tree, chooses $\mathrm{sbl}_{b}$ which is the node in the path from root to leaf, and generates a laconic OT ciphertext (using $\ell O T_{\text {const }}$.Send) which encrypts the input keys of the next traversing garbled circuit and uses $s b l_{b}$ as the hash value. Looking ahead, when the receiver evaluates the traversing circuit and obtains the laconic OT ciphertext, he can then use the siblings at the next layer to decrypt the ciphertext (by $\ell O T_{\text {const }}$. Receive) and obtain the corresponding input labels for the next traversing garbled circuit. Using the chain of traversing garbled circuits the receiver can therefore traverse from the first layer to the leaf of the Merkle tree. Furthermore, the correct keys for the first traversing circuit are sent via the $\ell O T_{\text {const }}$ with digest (i.e., root of the tree) as the hash value.

Finally, the last traversing circuit will transfer keys for the last reading circuit to the receiver in a similar fashion as above. The reading circuit takes the leaf as input and outputs $m_{\text {leaf }[t]}$, i.e., the message corresponding to the $t$-th bit of the leaf. Hence, when evaluating the reading circuit, the receiver can obtain the message $m_{\text {leaf }[t]}$.

SendWrite and ReceiveWrite are similar as Send and Receive, except that (a) ReceiveWrite updates the Merkle tree from the leaf to the root, and (b) the last writing circuit recomputes the root of the Merkle tree and outputs messages corresponding to the new root. To enable (b), the writing circuit will take as input the whole path of siblings from the root to the leaf. The input keys for the writing circuit corresponding to the siblings at the $(i+1)$-th layer are transferred via the $i$-th traversing circuit. That is, the $i$-th traversing circuit transfers the keys for the $(i+1)$-th transferring circuit as well as partial keys for the writing circuit. In the actual construction, both the reading circuit and writing circuit take as input the entire path of siblings (for the purpose of symmetry).

The Construction. Let $\ell O T_{\text {const }}=\left(\ell O T_{\text {const }}\right.$.crsGen, $\ell O T_{\text {const }}$. Hash, $\ell O T_{\text {const }}$. Send, $\ell O T_{\text {const }}$. Receive $)$ be a laconic OT protocol with factor-2 compression. Let $\mathrm{GC}=(\mathrm{GCircuit}$, Eval) be a circuit garbling scheme. Without loss of generality, let $D \in\{0,1\}^{M}$ be a database such that $|M|=2^{d}$. $\lambda$. A location $L \in[M]$ can be represented as $\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right) \in\{0,1\}^{d-1} \times[2 \lambda]$, where the bits $b_{i}$ 's define the path from the root to a leaf in the Merkle tree, and $t \in[2 \lambda]$ defines a position in that leaf.

Before delving into the construction, we first describe three gadget circuits: the traversing circuit $C^{\text {trav }}$, the reading circuit $C^{\text {read }}$, and the writing circuit $C^{\text {write }}$. These circuits are defined
formally in Figures 4, 5, and 6, respectively.
The traversing circuit has hardwired inside it a common reference string crs, a bit $b$ and two vectors of input keys Keys, Keys, each containing $2 \lambda$ key-pairs (a key-pair is a pair of $\lambda$-bit strings). It takes as input a pair of siblings $\mathrm{sbl}=\left(\mathrm{sbl}_{0}, \mathrm{sbl}_{1}\right)$, each of length $\lambda$, and generates two laconic OT Send messages with $\mathrm{sbl}_{\mathrm{b}}$ as the digest and Keys, $\widetilde{\text { Keys }}$ as message vectors respectively. Further, it also has the randomness needed for $\ell O T_{\text {const }}$.Send hardwired inside it.

The reading circuit $\mathrm{C}^{\text {read }}$ has a location $t \in[2 \lambda]$, and messages $m_{0}, m_{1} \in\{0,1\}^{\lambda}$ hardwired inside it. It takes as input a path of siblings from the root to a leaf, reads the $t$-th bit of the leaf, and outputs either $m_{0}$ or $m_{1}$ depending on that bit.

The writing circuit $C^{\text {write }}$ has hardwired inside it a common reference string crs, a location $L \in[M]$, a bit $b$ and a vector of messages Keys consisting of $\lambda$ key-pairs. It takes as input a path of siblings from the root to a leaf, changes the $t$-th bit of the leaf to $b$ (where $L$ corresponds to $t$-th location in the leaf), recomputes the Merkle tree root along the path, and outputs the corresponding labels for the new root/digest.

```
Circuit C \({ }^{\text {trav }}\)
    Hardwired Values: crs, \(b\), Keys, \(\widetilde{\text { Keys, } r, ~} \widetilde{r}\)
    Input: sbl
    Parse sbl as \(\left(\mathrm{sbl}_{0}\right.\), sbl \(\left._{1}\right)\)
    \(\mathrm{e} \leftarrow \ell O T_{\text {const }}\).Send (crs, sbl \({ }_{b}\), Keys; \(r\) )
    \(\widetilde{\mathrm{e}} \leftarrow \ell O T_{\text {const }}\).Send \(\left(\mathrm{crs}, \mathrm{sbl}_{b}, \widehat{\text { Keys; } ; \widetilde{r})}\right.\)
    Output (e, \(\widetilde{\mathrm{e}}\) )
```

Figure 4: The Traversing Circuit $\mathrm{C}^{\text {trav }}[\mathrm{crs}, b$, Keys, Keys, $r, \widetilde{r}]$

```
Circuit Cread
    Hardwired Values: t, mo, m
    Input: path
    Parse path = (sbl}\mp@subsup{|}{}{1},\ldots,\mathrm{ sbl }\mp@subsup{|}{}{d-1},\mathrm{ leaf }
    Output m}\mp@subsup{m}{\mathrm{ leaf }[t]}{
```

Figure 5: The Reading Circuit $\mathrm{C}^{\text {read }}\left[t, m_{0}, m_{1}\right]$
Now we construct the updatable $\ell O T$, namely (crsGen, Hash, Send, Receive, SendWrite, ReceiveWrite) as follows.

- $\operatorname{crsGen}\left(1^{\lambda}\right):$ Sample crs $\leftarrow \ell O T_{\text {const. }} \cdot \operatorname{crsGen}\left(1^{\lambda}\right)$ and output crs.
- Hash (crs, $D \in\{0,1\}^{M}$ ):

Build a Merkle tree $\hat{D}$ of $D$ using $\ell O T_{\text {const }}$.Hash (crs, •), as in Section 5.1.2.
Let digest be the root of $\hat{D}$.
Output (digest, $\hat{\text { D }}$ ).

- Send(crs, digest, $\left.L, m_{0}, m_{1}\right)$ :

Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$.
Pick $\left(\widetilde{\text { Keys }}^{1}, \ldots, \widetilde{\text { Keys }}^{d}\right)$ as input keys for $C^{\text {read }}$,

## Circuit Cwrite

Hardwired Values: crs, $L, b$, Keys
Input: path
Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$
Parse path $=\left(\left.\mathbf{s b}\right|^{1}, \ldots,\left.s\right|^{d-1}\right.$, leaf $)$, and parse $\left.\mathbf{s b}\right|^{i}=\left(\operatorname{sbl}_{0}^{i}, \operatorname{sbl}_{1}^{i}\right)$ for $i \in[d-1]$
$\operatorname{leaf}[t] \leftarrow b$
$\mathrm{sbl}^{d} \leftarrow$ leaf
For $i=d-1$ downto 1 :
$\mathrm{sbl}_{b_{i}}^{i} \leftarrow \ell O T_{\text {const }}$. Hash (crs, sbli+1 $)$
digest $^{*} \leftarrow \ell O T_{\text {const }} \cdot$ Hash $\left(\right.$ crs, sbl $\left.{ }^{1}\right)$.
Output Keys ${ }_{\text {digest }}{ }^{*}$
Figure 6: The Writing Circuit $\mathrm{C}^{\text {write }}$ [crs, $L, b$, Keys]
where $\widetilde{\text { Keys }}^{i}$ corresponds to the input keys of sbli for $i \in[d-1]$, and $\widetilde{\text { Keys }}^{d}$ corresponds to the input keys of leaf.

Let Keys ${ }^{d}$ be $0^{*}$
For $i=d-1$ downto 1 :
Pick $\mathrm{Keys}^{i}$ as input keys for $\mathrm{C}^{\text {trav }}$
Pick $r_{i}, \widetilde{r}_{i}$ as random coins for $\ell O T_{\text {const }}$.Send
$\widetilde{\mathrm{C}}_{i} \leftarrow \operatorname{GCircuit}\left(1^{\lambda}, \mathrm{C}^{\text {trav }}\left[\mathrm{crs}, b_{i}\right.\right.$, Keys $\left.^{i+1},{\widetilde{\text { Keys }^{2}}}^{i+1}, r_{i}, \widetilde{r}_{i}\right]$, Keys $\left.^{i}\right)$
$\mathrm{e}_{0} \leftarrow \ell O T_{\text {const }}$.Send(crs, digest, Keys ${ }^{1}$ )
$\widetilde{\mathrm{e}}_{0} \leftarrow \ell O T_{\text {const }}$.Send (crs, digest, $\widehat{\text { Keys }}{ }^{1}$ )
Outpute $=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {read }}\right)$

- Receive ${ }^{\hat{D}}(\mathrm{crs}, L, \mathrm{e})$ :

Parse e $=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {read }}\right)$
Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$
Parse D as a Merkle tree.
Denote the end node of path $b_{1} b_{2} \ldots b_{i}$ by $\hat{\mathrm{D}}_{b_{1} b_{2} \ldots b_{i}}$.
For $i=1$ to $d-1$ :

$$
\left.\mathrm{sb}\right|^{i} \leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 1}\right)
$$

Labels $^{i} \leftarrow \ell O T_{\text {const }}$.Receive $\left(\right.$ crs, $\left.\mathrm{e}_{i-1}, \mathrm{sbl}^{i}\right)$
$\widetilde{\text { Labels }}^{i} \leftarrow \ell O T_{\text {const }}$.Receive(crs, $\left.\widetilde{\mathrm{e}}_{i-1}, \mathrm{sbl}^{i}\right)$
$\left(\mathrm{e}_{i}, \widetilde{\mathrm{e}}_{i}\right) \leftarrow \operatorname{Eval}\left(\widetilde{\mathrm{C}}_{i}\right.$, Labels $\left.^{i}\right)$
leaf $\leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 1}\right)$
$\widetilde{\text { Labels }}^{d} \leftarrow \ell O T_{\text {const }}$.Receive (crs, $\widetilde{\mathrm{e}}_{d-1}$, leaf)
$m \leftarrow \operatorname{Eval}\left(\widetilde{\mathrm{C}}^{\text {read }},\left(\widetilde{\text { Labels }}^{1}, \ldots, \widetilde{\text { Labels }}^{d}\right)\right)$
Output $m$

- SendWrite(crs, digest, $\left.L, b,\left\{m_{j, 0}, m_{j, 1}\right\}_{j=1}^{\lambda}\right)$ :

Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$.
Pick $\left(\widetilde{\text { Keys }}^{1}, \ldots,{\widetilde{\text { Keys }^{\prime}}}^{d}\right)$ as input keys for Cwrite ,
where $\widetilde{\text { Keys }}^{i}$ corresponds to the input keys of sbl ${ }^{i}$ for $i \in[d-1]$,
and $\widetilde{\text { Keys }}^{d}$ corresponds to the input keys of leaf.
$\widetilde{\mathrm{C}}^{\text {write }} \leftarrow \operatorname{GCircuit}\left(1^{\lambda}, \mathrm{C}^{\text {write }}\left[\right.\right.$ crs $\left.\left., L, b,\left\{m_{j, 0}, m_{j, 1}\right\}_{j=1}^{\lambda}\right],\left(\widetilde{\text { Keys }}^{1}, \ldots, \widetilde{\text { Keys }}^{d}\right)\right)$
Let Keys ${ }^{d}$ be $0^{*}$
For $i=d-1$ downto 1 :
Pick $\mathrm{Keys}^{i}$ as input keys for $\mathrm{C}^{\text {trav }}$
Pick $r_{i}, \widetilde{r}_{i}$ as random coins for $\ell O T_{\text {const }}$. Send
$\widetilde{\mathrm{C}}_{i} \leftarrow$ GCircuit $\left(1^{\lambda}, \mathrm{C}^{\text {trav }}\left[\mathrm{crs}, b_{i}, \mathrm{Keys}^{i+1}, \widetilde{\text { Keys }}^{i+1}, r_{i}, \widetilde{r}_{i}\right]\right.$, Keys $\left.^{i}\right)$
$\mathrm{e}_{0} \leftarrow \ell O T_{\text {const }}$.Send(crs, digest, Keys ${ }^{1}$ )
$\widetilde{\mathrm{e}}_{0} \leftarrow \ell O T_{\text {const }}$.Send (crs, digest, $\widetilde{\text { Keys }}{ }^{1}$ )
Output $\mathrm{e}_{\mathrm{w}}=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {write }}\right)$

- ReceiveWrite ${ }^{\hat{\mathrm{D}}}\left(\mathrm{crs}, L, b, \mathrm{e}_{\mathrm{w}}\right)$ :

Parse $\mathrm{e}_{w}=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {write }}\right)$
Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$
Parse D as a Merkle tree.
Denote the end node of path $b_{1} b_{2} \ldots b_{i}$ by $\hat{\mathrm{D}}_{b_{1} b_{2} \ldots b_{i}}$.
Computing messages corresponding to the new digest:
For $i=1$ to $d-1$ :
$\left.\mathbf{s b}\right|^{i} \leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 1}\right)$
Labels $^{i} \leftarrow \ell O T_{\text {const }}$.Receive $\left(\right.$ crs, $\left.\mathrm{e}_{i-1}, \mathrm{sbl}^{i}\right)$
$\widetilde{\text { Labels }}^{i} \leftarrow \ell O T_{\text {const }}$.Receive(crs, $\widetilde{\mathrm{e}}_{i-1}$, sbl $\left.{ }^{i}\right)$
$\left(\mathrm{e}_{i}, \widetilde{\mathrm{e}}_{i}\right) \leftarrow \operatorname{Eval}\left(\widetilde{\mathrm{C}}_{i}\right.$, Labels $\left.^{i}\right)$
leaf $\leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 1}\right)$
${\left.\widetilde{\text { Labels }}^{d} \leftarrow \ell O T_{\text {const }} \text {.Receive(crs, } \widetilde{\mathrm{e}}_{d-1}, \text { leaf) }\right) ~}_{d}$
$\left\{m_{j}\right\}_{j=1}^{\lambda} \leftarrow \operatorname{Eval}\left(\widetilde{\mathrm{C}}^{\text {write }},\left(\widetilde{\text { Labels }}^{1}, \ldots, \widetilde{\text { Labels }}^{d}\right)\right)$

## Updating the Merkle tree:

$\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 0} \| \hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 1}\right)[t] \leftarrow b$
For $i=d-1$ downto 0 :

$$
\hat{\mathrm{D}}_{b_{1} \ldots b_{i}} \leftarrow \ell O T_{\text {const }} \cdot \operatorname{Hash}\left(\text { crs, } \hat{\mathrm{D}}_{b_{1} \ldots b_{\imath} 0} \| \hat{\mathrm{D}}_{b_{1} \ldots b_{i} 1}\right)
$$

Update digest with the new root of $\hat{D}$
Output $\left\{m_{j}\right\}_{j=1}^{\lambda}$
Efficiency. It can be seen from the scheme that the length of digest is $\lambda$. The algorithm Hash runs in time $|D| \cdot \operatorname{poly}(\log |D|, \lambda)$. Furthermore, Send, Receive, SendWrite, SendWrite all run in time poly $(\log |D|, \lambda)$.

Correctness. We briefly argue (perfect) correctness of the updatable laconic OT scheme. Given a ciphertext $\mathrm{e}=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {read }}\right)$ computed by Send, correctness of $\ell O T_{\text {const }}$ ensures that Labels ${ }^{1} \leftarrow \ell O T_{\text {const }}$.Receive (crs, $\mathrm{e}_{0}$, sbl ${ }^{1}$ ) outputs the correct labels for $\widetilde{\mathrm{C}}_{1}$ and that $\widetilde{\text { Labels }}^{1} \leftarrow$ $\ell O T_{\text {const }}$.Receive (crs, $\widetilde{\mathrm{e}}_{0}, \mathrm{sbl}^{1}$ ) outputs the correct labels for $\widetilde{\mathrm{C}}^{\text {read }}$, namely Labels ${ }^{1}=$ Keys $_{\text {sbl }}{ }^{1}$ and $\widetilde{\text { Labels }}^{1}=\widetilde{\text { Keys }_{\text {sbl }}{ }^{1}}{ }^{1}$. In turn, correctness of the garbling scheme guarantees that $\widetilde{\mathrm{C}}_{1}$ outputs the correct $\left(\mathrm{e}_{1}, \widetilde{\mathrm{e}}_{1}\right)$, namely $\mathrm{e}_{1}=\ell O T_{\text {const }}$.Send (crs, sbl $\left.\right|_{b_{1}} ^{1}$, Keys $\left.^{2} ; r_{1}\right)$ and $\widetilde{\mathrm{e}}_{1}=\ell O T_{\text {const }}$.Send (crs, sbl $\left.\left.\right|_{b_{1}} ^{1}, \widetilde{\text { Keys }}^{2} ; \widetilde{r}_{1}\right)$. It follows inductively that for every $i=1,2, \ldots, d-1$, Labels $^{i}=$ Keys $_{\text {sbl }}^{i}{ }^{i}, \widetilde{\text { Labels }}^{i}={\widetilde{\mathrm{Keys}_{\mathrm{sbl}}{ }^{i}}}_{i}, \mathrm{e}_{i}=$ $\ell O T_{\text {const }}$ Send (crs, sbl $b_{i}^{i}$, Keys $^{i+1} ; r_{i}$ ), $\widetilde{\mathrm{e}}_{i}=\ell O T_{\text {const }}$.Send (crs, sbl $\left.{ }_{b_{i}}^{i},{\widetilde{\text { Keys }^{2}}}^{i+1} ; \widetilde{r}_{i}\right)$. Again by using the correctness of $\ell O T_{\text {const }},{\widetilde{\text { Labels }^{2}}}^{d} \leftarrow \ell O T_{\text {const }}$.Receive (crs, $\widetilde{\mathrm{e}}_{d-1}$, leaf) gives $\widetilde{\text { Labels }}^{d}=\widetilde{\text { KeyS }}_{\text {leaf }}^{d}$. Then by using correctness of the garbling scheme it follows that evaluating $\widetilde{\mathrm{C}}^{\text {read }}$ gives the correct output $m_{D[L]}$. Correctness with regard to writes can be argued analogously.

### 5.3 Security

In this section, we will prove the security of the above updatable laconic OT scheme.
Theorem 5.1 (Sender Privacy against Semi-honest Receivers). Given that $\ell O T_{\text {const }}$ has sender privacy and that the garbled circuit scheme GCircuit is secure, the updatable laconic OT scheme $\ell O T$ has sender privacy.

Proof. Let $\ell$ OTSim const be the simulator for $\ell O T_{\text {const }}$ and CircSim be the simulator for the garbling scheme GCircuit. Below, we provide the two simulators $\ell$ OTSim for a read and $\ell$ OTSimWrite for the write.

- $\ell$ OTSim(crs, $D, L, m)$ :
(digest, $\hat{\mathrm{D}}) \leftarrow$ Hash $($ crs,$D)$
Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$
$\left(\widetilde{\mathrm{C}}^{\text {read }},\left(\widetilde{\text { Labels }}^{1}, \ldots, \widetilde{\text { Labels }}^{d}\right)\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {read }}, m\right)$
leaf $\leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 1}\right)$
$\mathrm{e}_{d-1} \leftarrow \ell$ OTSim $_{\text {const }}\left(\right.$ crs, leaf, $\left.0^{*}\right)$
$\widetilde{\mathrm{e}}_{d-1} \leftarrow \ell$ OTSim ${ }_{\text {const }}\left(\mathrm{crs}\right.$, leaf, $\left.\widetilde{\text { Labels }}^{d}\right)$
For $i=d-1$ downto 1 :

$$
\left(\widetilde{\mathrm{C}}_{i}, \text { Labels }{ }^{i}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {trav }},\left(\mathrm{e}_{i}, \widetilde{\mathrm{e}}_{i}\right)\right)
$$

$$
\mathrm{sbl}^{i} \leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 1}\right)
$$

$$
\mathrm{e}_{i-1} \leftarrow \ell \mathrm{OTSim}_{\text {const }}\left(\mathrm{crs}^{2} \mathrm{sbl}^{i}, \text { Labels }^{i}\right)
$$

$$
\widetilde{\mathrm{e}}_{i-1} \leftarrow \ell \mathrm{OTSim}_{\text {const }}\left(\mathrm{crs}, \widetilde{\sim}_{\mathrm{sbl}}{ }^{i}, \widetilde{\sim}_{\text {Labels }}{ }^{i}\right)
$$

Output e $=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {read }}\right)$

- $\ell$ OTSimWrite(crs, $\left.D, L, b,\left\{m_{j}\right\}_{j=1}^{\lambda}\right)$ :
$($ digest, $\hat{\mathrm{D}}) \leftarrow$ Hash $($ crs,$D)$
Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$
$\left(\widetilde{\mathrm{C}}^{\text {write }},\left(\widetilde{\text { Labels }}^{1}, \ldots, \widetilde{\text { Labels }}^{d}\right)\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {write }},\left\{m_{j}\right\}_{j=1}^{\lambda}\right)$

$$
\begin{aligned}
& \text { leaf } \leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{d-1} 1}\right) \\
& \mathrm{e}_{d-1} \leftarrow \ell \mathrm{OTSim} \text { const }\left(\text { crs, leaf, } 0^{*}\right) \\
& \widetilde{\mathrm{e}}_{d-1} \leftarrow \ell \widetilde{\mathrm{OTSim}}_{\text {const }}\left(\mathrm{crs}, \text { leaf },{\text { Labels } \left.^{d}\right)}^{\text {and }}\right. \text {. }
\end{aligned}
$$

For $i=d-1$ downto 1 :
$\left(\widetilde{\mathrm{C}}_{i}\right.$, Labels $\left.{ }^{i}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {trav }},\left(\mathrm{e}_{i}, \widetilde{\mathrm{e}}_{i}\right)\right)$
$\mathrm{sb}^{i} \leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 1}\right)$
$\mathrm{e}_{i-1} \leftarrow \ell$ OTSim $_{\text {const }}\left(\right.$ crs, sbl ${ }^{i}$, Labels $^{i}$ )
$\widetilde{\mathrm{e}}_{i-1} \leftarrow \ell$ OTSim const $\left(\mathrm{crs}, \mathrm{sbl}^{i}, \widetilde{\text { Labels }}^{i}\right)$
Output $\mathrm{e}_{\mathrm{w}}=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {write }}\right)$
In the following we will only prove sender security with regard to reads. Since (Send, $\ell$ OTSim) and (SendWrite, $\ell$ OTSimWrite) are very similar, sender security with regard to writes can be argued analogously.

We prove security via a hybrid argument. In the first hybrid, we replace the ciphertexts $\mathrm{e}_{0}$ and $\widetilde{\mathrm{e}}_{0}$ computed by $\ell O T_{\text {const }}$. Send with ciphertexts computed by $\ell O T \operatorname{Sim}_{\text {const }}$.

Afterwards, we can use security of the garbling scheme to replace the honestly generated $\widetilde{\mathrm{C}}_{1}$ ${\underset{\mathrm{C}}{2}}^{\text {with }}$ a simulated one, and run $\ell$ OTSim const using the simulated input labels of $\widetilde{\mathrm{C}}_{1}$. As the output of $\widetilde{C}_{1}$ is again a pair of ciphertexts ( $e_{1}, \widetilde{e}_{1}$ ), we will simulate it using $\ell O T S \operatorname{Sim}_{\text {const }}$ in the next hybrid. We continue alternating between simulating the garbled circuits and simulating the ciphertexts, until reaching the reading circuit. Once we reach the reading circuit, it holds that all $\widetilde{\text { Labels }}^{i}$ are information theoretically fixed to the path from the root to the leaf containing $L$. We will then invoke the garbled circuit security of the reading circuit, and conclude the hybrid argument.

The formal proof is as follows. For every PPT machine $\mathcal{A}$, let $\operatorname{crs} \leftarrow \operatorname{crsGen}\left(1^{\lambda}\right)$, and let $\left(D, L, m_{0}, m_{1}\right) \leftarrow \mathcal{A}$ (crs). Further let digest $\leftarrow \operatorname{Hash}(\operatorname{crs}, D)$. Then we will prove that the two distributions (crs, Send(crs, digest, $\left.L, m_{0}, m_{1}\right)$ ) and (crs, $\ell$ OTSim (crs, $\left.\left.D, L, m_{D[L]}\right)\right)$ are computationally indistinguishable. Consider the following hybrids.

- Hybrid 0: This is the real experiment, i.e., (crs, Send(crs, digest, $\left.L, m_{0}, m_{1}\right)$ ).
- Hybrid 1: Same has hybrid 0, except that $\mathrm{e}_{0}$ and $\widetilde{\mathrm{e}}_{0}$ are computed as follows.

$$
(\text { digest }, \hat{\mathrm{D}}) \leftarrow \operatorname{Hash}(\operatorname{crs}, D)
$$

Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$.
Pick $\left(\widetilde{\text { Keys }}{ }^{1}, \ldots, \widetilde{\text { Keys }}^{d}\right)$ as input keys for Cread
$\widetilde{\mathrm{C}}^{\text {read }} \leftarrow \operatorname{GCircuit}\left(1^{\lambda}, \mathrm{C}^{\text {read }}\left[t, m_{0}, m_{1}\right],\left(\widetilde{\text { Keys }}^{1}, \ldots, \widetilde{\text { Keys }}^{d}\right)\right)$
Let Keys ${ }^{d}$ be $0^{*}$
For $i=d-1$ downto 1 :
Pick Keys ${ }^{i}$ as input keys for $\mathrm{C}^{\text {trav }}$
Pick $r_{i}, \widetilde{r}_{i}$ as random coins for $\ell O T_{\text {const }}$.Send

| $\quad \widetilde{\mathrm{C}}_{i} \leftarrow$ GCircuit $\left(1^{\lambda}, \mathrm{C}^{\text {trav }}\left[\mathrm{crs}, b_{i}\right.\right.$, Keys $\left.^{i+1}, \widetilde{\text { Keys }}^{i+1}, r_{i}, \widetilde{r}_{i}\right]$, Keys $\left.^{i}\right)$ |
| :--- |
| sbl $^{1} \leftarrow\left(\hat{\mathrm{D}}_{0}, \hat{\mathrm{D}}_{1}\right)$ |
| Labels $^{1} \leftarrow$ Keys $_{\text {sbl }^{1}}$ |


| $\widetilde{\text { Labels }}^{1} \leftarrow{\widetilde{\text { Keys }_{\text {sbl }}}}^{1}$ |
| :--- |
| $\mathrm{e}_{0} \leftarrow \ell \mathrm{OTSim}_{\text {const }}\left(\mathrm{crs}\right.$, sbl $^{1}$, Labels $\left.^{1}\right)$ |
| $\widetilde{\mathrm{e}}_{0} \leftarrow \ell \mathrm{OTSim}$ |
| const |
| Output $\left.\mathrm{ers}, \widetilde{\mathrm{sbl}}^{1}, \widetilde{\text { Labels }}^{1}\right)$ |
| $\left.\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {read }}\right)$ |

The differences between hybrid 0 and hybrid 1 have been marked with boxes. Indistinguishability between hybrid 0 and hybrid 1 can be argued from the multi-execution sender security of $\ell O T_{\text {const }}$ via the following reduction. Given crs by the experiment and the adversarial input $D$, compute hybrid 1 until $\mathrm{e}_{0}$ and $\widetilde{\mathrm{e}}_{0}$ are computed. In particular, compute $\hat{\mathrm{D}}$, Keys $^{1}, \widetilde{\text { Keys }}^{1}$, sbl ${ }^{1}$, Labels $^{1}=$ Keys $_{\text {sbl }^{1}}{ }^{1}, \widetilde{\text { Labels }}^{1}=\widetilde{\text { Keys }_{\text {sbl }}} 1$. Then choose sbl ${ }^{1}$ as the database and (Keys ${ }^{1}, \widetilde{\text { Keys }}^{1}$ ) as the messages for $\ell O T_{\text {const }}$, and obtain the challenge $\left(e_{0}^{*}, \widetilde{e}_{0}^{*}\right)$, which is from one of the following two distributions:

$$
\begin{aligned}
& \left.\left(\ell O T_{\text {const }} \text {.Send (crs, sbl }{ }^{1} \text {, Keys }{ }^{1} \text { ), } \ell O T_{\text {const }} \text {.Send(crs, sbl }{ }^{1}, \widetilde{\text { Keys }}^{1}\right)\right) ; \\
& \left(\ell \text { OTSim } \text { const }\left(\text { crs, sbl }{ }^{1}, \text { Labels }^{1}\right), \ell \text { OTSim }{ }_{\text {const }}\left(\text { crs, sbl }{ }^{1}, \widetilde{\text { Labels }}^{1}\right)\right) \text {. }
\end{aligned}
$$

If $\left(e_{0}^{*}, \widetilde{e}_{0}^{*}\right)$ is from the first distribution, then it results in hybrid 0 ; otherwise it results in hybrid 1. Hence the indistinguishability of the two distributions implies indistinguishability of the two hybrids.

- Hybrid $2 k(k=1,2, \ldots, d-1)$ : Same has hybrid $2 k-1$, except that $\widetilde{\mathrm{C}}_{k}$ is computed as follows.
(digest, $\hat{\mathrm{D}}) \leftarrow \operatorname{Hash}($ crs,$D)$
Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$.
Pick $\left(\widetilde{\text { Keys }^{1}}, \ldots, \widetilde{\text { Keys }}^{d}\right)$ as input keys for $C^{\text {read }}$
$\widetilde{\mathrm{C}}^{\text {read }} \leftarrow \operatorname{GCircuit}\left(1^{\lambda}, \mathrm{C}^{\text {read }}\left[t, m_{0}, m_{1}\right],\left(\widetilde{\text { Keys }}^{1}, \ldots, \widetilde{\text { Keys }^{d}}{ }^{d}\right)\right)$
Let Keys ${ }^{d}$ be $0^{*}$
For $i=d-1$ downto $k+1$ :
Pick Keys ${ }^{i}$ as input keys for $\mathrm{C}^{\text {trav }}$
Pick $r_{i}, \widetilde{r}_{i}$ as random coins for $\ell O T_{\text {const }}$.Send
$\widetilde{\mathrm{C}}_{i} \leftarrow$ GCircuit $\left(1^{\lambda}, \mathrm{C}^{\text {trav }}\left[\right.\right.$ crs,$b_{i}$, Keys $\left.^{i+1}, \widetilde{\text { Keys }}^{i+1}, r_{i}, \widetilde{r}_{i}\right]$, Keys $\left.^{i}\right)$

| $\mathrm{sbl}^{k} \leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{k-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{k-1} 1}\right)$ |
| :--- |
| $\mathrm{e}_{k} \leftarrow \ell O T_{\text {const }}$ Send $\left(\mathrm{crs}, \mathrm{sbl}_{b_{k}}^{k}\right.$, Keys $\left.^{k+1}\right)$ |
| $\widetilde{\mathrm{e}}_{k} \leftarrow \ell O T_{\text {const }}$.Send $\left(\mathrm{crs}, \mathrm{sbl}_{b_{k}}^{k}, \widetilde{\text { Keys }}^{k+1}\right)$ |

For $i=k$ downto 1 :
$\left(\widetilde{\mathrm{C}}_{i}\right.$, Labels $\left.{ }^{i}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {trav }},\left(\mathrm{e}_{i}, \widetilde{\mathrm{e}}_{i}\right)\right)$
$\mathrm{sbl}^{i} \leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 1}\right)$

$$
\begin{aligned}
& \widehat{\text { Labels }}^{i} \leftarrow{\widetilde{\text { Keys }_{\text {sbl }}}}^{i} \\
& \mathrm{e}_{i-1} \leftarrow \ell \text { OTSim } \text { const }\left(\mathrm{crs}, \text { sbl }^{i}, \text { Labels }^{i}\right) \\
& \widetilde{\mathrm{e}}_{i-1} \leftarrow \ell \text { OTSim } \text { const }\left(\mathrm{crs}, \mathrm{sbl}^{i}, \widetilde{\text { Labels }}^{i}\right) \\
& \text { Output } \mathrm{e}=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {read }}\right)
\end{aligned}
$$

- Hybrid $2 k+1(k=1,2, \ldots, d-1)$ : Same has hybrid $2 k$, except that $\mathrm{e}_{k}$ and $\widetilde{\mathrm{e}}_{k}$ are computed as follows.
(digest, $\hat{\mathrm{D}}) \leftarrow$ Hash $($ crs,$D)$
Parse $L=\left(b_{1}, b_{2}, \ldots, b_{d-1}, t\right)$.
Pick $\left(\widetilde{\text { Keys }}^{1}, \ldots, \widetilde{\text { Keys }}^{d}\right)$ as input keys for Cead

Let Keys ${ }^{d}$ be $0^{*}$
For $i=d-1$ downto $k+1$ :
Pick Keys ${ }^{i}$ as input keys for $\mathrm{C}^{\text {trav }}$
Pick $r_{i}, \widetilde{r}_{i}$ as random coins for $\ell O T_{\text {const }}$. Send


$$
\begin{aligned}
& \left(\widetilde{\mathrm{C}}_{i}, \text { Labels }^{i}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {trav }},\left(\mathrm{e}_{i}, \widetilde{\mathrm{e}}_{i}\right)\right) \\
& \mathrm{sbl} i^{i} \leftarrow\left(\hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 0}, \hat{\mathrm{D}}_{b_{1} \ldots b_{i-1} 1}\right) \\
& \widetilde{\text { Labels }}^{i} \leftarrow \widetilde{\text { Keys }}_{\text {sbl }}{ }^{i} \\
& \mathrm{e}_{i-1} \leftarrow \ell \text { OTSim }_{\text {const }}\left(\text { crs, sbl }{ }^{i}, \text { Labels }^{i}\right. \text { ) } \\
& \widetilde{\mathrm{e}}_{i-1} \leftarrow \ell \text { OTSim }_{\text {const }}\left(\mathrm{crs}, \mathrm{sbl}{ }^{i}, \widetilde{\text { Labels }}^{i}\right) \\
& \text { Output } \mathrm{e}=\left(\mathrm{e}_{0}, \widetilde{\mathrm{e}}_{0}, \widetilde{\mathrm{C}}_{1}, \ldots, \widetilde{\mathrm{C}}_{d-1}, \widetilde{\mathrm{C}}^{\text {read }}\right)
\end{aligned}
$$

We will first show that hybrids $2 k-1$ and $2 k$ are indistinguishable via a reduction to the security of the garbling scheme GCircuit. Notice that the only difference between hybrids $2 k-1$ and $2 k$ is $\left(\mathrm{C}_{k}, \mathrm{e}_{k-1}\right)$. Consider the following two distributions:

$$
\left.\begin{array}{l}
\left(\widetilde{\mathrm{C}}_{k}, \text { Labels }^{k}\right) \leftarrow\left(\operatorname{GCircuit}\left(1^{\lambda}, \mathrm{C}^{\text {trav }}\left[\mathrm{crs}, b_{k}, \text { Keys }^{k+1},{\widetilde{\text { Keys }^{k+1}}}^{k+1}, r_{i}, \widetilde{r}_{k}\right], \text { Keys }^{k}\right), \text { Keys }_{\text {sbl }^{k}}^{k}\right) ; \\
\left(\widetilde{\mathrm{C}}_{k},\right. \text { Labels }
\end{array}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {trav }},\left(\mathrm{e}_{k}, \widetilde{\mathrm{e}}_{k}\right)\right), ~ \$
$$

where $\mathrm{e}_{k} \leftarrow \ell O T_{\text {const }}$.Send(crs, sbl $\left.\right|_{b_{k}} ^{k}$, Keys $^{k+1}$ ) and $\widetilde{\mathrm{e}}_{k} \leftarrow \ell O T_{\text {const }}$.Send(crs, sbl $\left.\right|_{b_{k}} ^{k}, \widetilde{\mathrm{Keys}}^{k+1}$ ). Notice that $\left(\mathrm{e}_{k}, \widetilde{\mathrm{e}}_{k}\right)$ is the output of ( $\widetilde{\mathrm{C}}_{k}$, Labels $\left.{ }^{k}\right)$ from the first distribution. By security of
the garbled circuit scheme, the above two distributions are computationally indistinguishable. Furthermore, if $\widetilde{C}_{k}$ is generated using the first distribution and $\mathrm{e}_{k-1}$ is computed using Labels ${ }^{k}$ from the first distribution, then it results in hybrid $2 k-1$; otherwise it results in hybrid $2 k$. Hence the two hybrids are computationally indistinguishable.
Indistinguishability of hybrids $2 k$ and $2 k+1$ follows again from sender security of $\ell O T_{\text {const }}$, in the same fashion as the indistinguishability between hybrids 0 and 1.

- Hybrid 2d: This is the simulated experiment, namely (crs, $\ell$ OTSim (crs, $\left.D, L, m_{D[L]}\right)$ ).

The difference between hybrids $2 d-1$ and $2 d$ is $\left(\widetilde{\mathrm{C}}^{\text {read }}, \widetilde{\mathrm{e}}_{0}, \ldots, \widetilde{\mathrm{e}}_{d-1}\right)$. The indistinguishability would follow from the security of garbled circuit scheme, similarly as when we argue indistinguishability hybrids $2 k-1$ and $2 k$.

## 6 Warm-Up Application: Non-Interactive Secure Computation (NISC) on Large Inputs in RAM Setting

In this section, we consider the application of non-interactive secure computation in the RAM (random access machine) setting.

### 6.1 Background

We recall the needed background of RAM computation model and two-message oblivious transfer in this section. We will also use garbled circuits (see Section 5.1.1) as building blocks.

### 6.1.1 Random Access Machine (RAM) Model of Computation

Now we define the RAM model of computation. Parts of this subsection have been taken verbatim from [GLO15].

Notation for the RAM Model of Computation. The RAM model consists of a CPU and a memory storage of size $M$. The CPU executes a program that can access the memory by using read/write operations. In particular, for a program $P$ with memory of size $M$ we denote the initial contents of the memory data by $D \in\{0,1\}^{M}$. Additionally, the program gets a "short" input $x \in\{0,1\}^{m}$, which we alternatively think of as the initial state of the program. We use the notation $P^{D}(x)$ to denote the execution of program $P$ with initial memory contents $D$ and input $x$. The program $P$ can read from and write to various locations in memory $D$ throughout its execution. ${ }^{10}$

We will also consider the case where several different programs are executed sequentially and the memory persists between executions. We denote this process as $\left(y_{1}, \ldots, y_{\ell}\right)=\left(P_{1}\left(x_{1}\right), \ldots, P_{\ell}\left(x_{\ell}\right)\right)^{D}$ to indicate that first $P_{1}^{D}\left(x_{1}\right)$ is executed, resulting in some memory contents $D_{1}$ and output $y_{1}$, then $P_{2}^{D_{1}}\left(x_{2}\right)$ is executed resulting in some memory contents $D_{2}$ and output $y_{2}$ etc. As an example,

[^6]imagine that $D$ is a huge database and the programs $P_{i}$ are database queries that can read and possibly write to the database and are parameterized by some values $x_{i}$.

CPU-Step Circuit. Consider an execution of a RAM program which involves at most $t$ CPU steps. We represent a RAM program $P$ via $t$ small CPU-Step Circuits each of which executes one CPU step. In this work we will denote one CPU step by:

$$
C_{\mathrm{CPU}}^{P}(\text { state }, \mathrm{rData})=\left(\text { state }^{\prime}, \mathrm{R} / \mathrm{W}, L, \mathrm{wData}\right)
$$

This circuit takes as input the current CPU state state and a bit rData. Looking ahead the bit rData will be read from the memory location that was requested by the previous CPU step. The circuit outputs an updated state state', a read or write bit R/W, the next location to read/write from $L \in[M]$, and a bit wData to write into that location (wData $=\perp$ when reading). The sequence of locations and read/write values collectively form what is known as the access pattern, namely MemAccess $=\left\{\left(\mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}, \mathrm{wData}^{\tau}\right): \tau=1, \ldots, t\right\}$.

Note that in the description above without loss of generality we have made some simplifying assumptions. We assume that each CPU-step circuit always reads from or write some location in memory. This is easy to implement via a dummy read and write step. Moreover, we assume that the instructions of the program itself are hardwired into the CPU-step circuits.

Representing RAM computation by CPU-Step Circuits. The computation $P^{D}(x)$ starts with the initial state set as state $_{1}=x$. In each step $\tau \in\{1, \ldots t\}$, the computation proceeds as follows: If $\tau=1$ or $\mathrm{R} / \mathrm{W}^{\tau-1}=$ write, then $\mathrm{rData}{ }^{\tau}:=0$; otherwise $\mathrm{rData}{ }^{\tau}:=D\left[L^{\tau-1}\right]$. Next it executes the CPU-Step Circuit $C_{\mathrm{CPU}}^{P}\left(\right.$ state $\left.^{\tau}, \mathrm{rData}^{\tau}\right)=\left(\right.$ state $\left.^{\tau+1}, \mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}, \mathrm{wData}^{\tau}\right)$. If R/W ${ }^{\tau}=$ write, then set $D\left[L^{\tau}\right]=\mathrm{wData}^{\tau}$. Finally, when $\tau=t$, then state ${ }^{\tau+1}$ is the output of the program.

### 6.1.2 Oblivious Transfer

[AIR01, NP01, ?] gave two-message oblivious transfer (OT) protocols. We describe the definition below and refer the reader to [AIR01,NP01, ?] for details.

Definition 6.1 (Two-Message Oblivious Transfer). A two-message oblivious transfer protocol $\mathrm{OT}=$ $\left(\mathrm{OT}_{1}, \mathrm{OT}_{2}, \mathrm{OT}_{3}\right)$ is a protocol between a sender $S$ and a receiver $R$ where $S$ gets as input two strings $s_{1}, s_{2}$ of equal length and $R$ gets as input a choice bit $x \in\{0,1\}$. The algorithms have the following syntax:

- $\left(m_{1}\right.$, secret $) \leftarrow \mathrm{OT}_{1}\left(1^{\lambda}, x\right):$ It takes as input the security parameter $1^{\lambda}$ and receiver's choice bit $x \in\{0,1\}$ and outputs the first OT message $m_{1}$ (sent by the receiver) and receiver's secret state secret.
- $m_{2} \leftarrow \mathrm{OT}_{2}\left(m_{1}, s_{0}, s_{1}\right)$ : It takes as input the first OT message and the sender's input $\left(s_{0}, s_{1}\right)$, and outputs the second OT message $m_{2}$ (sent back to the receiver).
- $s \leftarrow \mathrm{OT}_{3}\left(m_{2}\right.$, secret $):$ It takes $m_{2}$ and secret as input, and outputs a string $s$.

The following conditions are satisfied:

- Perfect Correctness: For all security parameter $\lambda$, sender input strings $\left(s_{1}, s_{2}\right)$ of equal length, and receiver's choice bit $x$, let $\left(m_{1}\right.$, secret $) \leftarrow \mathrm{OT}_{1}\left(1^{\lambda}, x\right)$, $m_{2} \leftarrow \mathrm{OT}_{2}\left(m_{1}, s_{0}, s_{1}\right)$, and $s \leftarrow \mathrm{OT}_{3}\left(m_{2}\right.$, secret), then it holds that

$$
\operatorname{Pr}\left[s=s_{x}\right]=1 .
$$

- Receiver Security: The following two distributions are computationally indistinguishable:

$$
\mathrm{OT}_{1}\left(1^{\lambda}, 0\right) \stackrel{\mathrm{c}}{\approx} \mathrm{OT}_{1}\left(1^{\lambda}, 1\right) .
$$

- Sender Security: There exists a PPT simulator OTSim such that for all sender input strings $\left(s_{1}, s_{2}\right)$ of equal length and receiver's choice bit $x$, and any first message $m_{1}$ in the support of $\mathrm{OT}_{1}\left(1^{\lambda}, x\right)$, the following two distributions are statistically close:

$$
\mathrm{OT}_{2}\left(m_{1}, s_{0}, s_{1}\right) \stackrel{\mathrm{s}}{\approx} \mathrm{OTSim}\left(1^{\lambda}, x, s_{x}, m_{1}\right) .
$$

We described the above definition with respect to one OT, but the same formalism naturally extends to support multiple parallel executions of OT. We will use the following shorthand notations (generalizing the above notions) to run multiple parallel executions. Let Keys $=$ $\left(\left(\mathrm{Key}_{1,0}, \mathrm{Key}_{1,1}\right), \ldots,\left(\mathrm{Key}_{n, 0}, \mathrm{Key}_{n, 1}\right)\right)$ be a list of $n$ string-pairs, and $x \in\{0,1\}^{n}$ be an $n$-bit choice string. Then we define

- $\left(m_{1}\right.$, secret $) \leftarrow \mathrm{OT}_{1}\left(1^{\lambda}, x\right)=\left(\mathrm{OT}_{1}\left(1^{\lambda}, x_{1}\right), \ldots, \mathrm{OT}_{1}\left(1^{\lambda}, x_{n}\right)\right)$.
- $m_{2} \leftarrow \mathrm{OT}_{2}\left(m_{1}, \mathrm{Keys}\right)=\left(\mathrm{OT}_{2}\left(m_{1,1}, \mathrm{Key}_{1,0}, \mathrm{Key}_{1,1}\right), \cdots, \mathrm{OT}_{2}\left(m_{1, n}, \mathrm{Key}_{n, 0}, \mathrm{Key}_{n, 1}\right)\right)$.
- Labels $\leftarrow \mathrm{OT}_{3}\left(m_{2}\right.$, secret $)=\left(\mathrm{OT}_{3}\left(m_{2,1}\right.\right.$, secret $\left._{1}\right), \ldots, \mathrm{OT}_{3}\left(m_{2, n}\right.$, secret $\left.\left._{n}\right)\right)$.

In the above $m_{1}=\left(m_{1,1}, \ldots, m_{1, n}\right), m_{2}=\left(m_{2,1}, \ldots, m_{2, n}\right)$, secret $=\left(\right.$ secret $_{1}, \ldots$, secret $\left._{n}\right)$. Correctness guarantees that Labels $=\mathrm{Keys}_{x}=\left(\mathrm{Key}_{1, x_{1}}, \ldots, \mathrm{Key}_{n, x_{n}}\right)$.

Moreover, we will use two important properties of the oblivious transfer [NP01] for our applications: (1) Security holds for multiple second OT messages with regard to the same first OT message. This will be crucial for extending NISC for RAM to support multiple senders with the same receiver. (2) The second OT message is re-randomizable. This will be crucial for the application of multi-hop homomorphic encryption for RAM.

### 6.2 Formal Model for NISC in RAM Setting

Suppose the receiver owns a large confidential database $D \in\{0,1\}^{M}$. It first publishes a short message, denoted by $m_{1}$, which hides $D$. Afterwards, if a sender wants to run a RAM program $P$ (with input $x$ ) on $D$, it can send a single message $m_{2}$ to the receiver. For security we require that $m_{2}$ only reveals the output $P^{D}(x)$ and the memory access pattern MemAccess of the execution to the receiver. We require that once $m_{1}$ is published, the computational cost of both the sender (in computing $m_{2}$ ) and the receiver (in evaluation), as well as the size of $m_{2}$, should grow only with the running time of the RAM computation and the size of $m_{1}$, and is independent of the size of $D$.

Moreover, the sender can run a sequence of programs on a persistent database by sending one message per program to the receiver. Finally, the receiver can run the protocol in parallel with multiple senders, where the same $m_{1}$ is used. For ease of exposition, below we will describe the setting of one single sender executing one program with the receiver. We provide details on above extensions in Section 6.6.

The Model. A non-interactive secure RAM computation scheme NISC-RAM = (Setup, EncData, EncProg, Dec) has the following syntax. It is a two-party protocol between a receiver holding a large secret database $D$ and a sender holding secret program $P$ of running time $t$ and a short input $x$.

- Setup: crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}\right)$.

On input the security parameter $1^{\lambda}$, it outputs a common reference string.

- Database Encryption: $\left(m_{1}, \widetilde{D}\right) \leftarrow$ EncData(crs, $\left.D\right)$.

On input the common reference string crs and a database $D \in\{0,1\}^{M}$, it outputs a message $m_{1}$ and a secret state $\widetilde{D}$. The receiver publishes $m_{1}$ as the short message corresponding to $D$.

- Program Encryption: $m_{2} \leftarrow \operatorname{EncProg}\left(\mathrm{crs}, m_{1},(P, x, t)\right)$.

It takes as input the crs, a message $m_{1}$, a RAM program $P$ with input $x$ and maximum run-time $t$. It then outputs another message $m_{2}$. The sender sends the message $m_{2}$.

- Decryption: $y \leftarrow \operatorname{Dec}^{\widetilde{D}}\left(\mathrm{crs}, m_{2}\right)$.

The procedure Dec is modeled as a RAM program that can read and write to arbitrary locations of its database initially containing $\widetilde{D}$. This procedure is run by the receiver. On input the crs and $m_{2}$, it outputs $y$.

The following conditions are satisfied:

- Correctness: For every database $D \in\{0,1\}^{M}$ where $M=\operatorname{poly}(\lambda)$ for any polynomial function poly $(\cdot)$, for every RAM program ( $P, x, t$ ), it holds that

$$
\operatorname{Pr}\left[\operatorname{Dec}^{\widetilde{D}}\left(\operatorname{crs}, m_{2}\right)=P^{D}(x)\right]=1,
$$

where $\operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}\right),\left(m_{1}, \widetilde{D}\right) \leftarrow \operatorname{EncData}(\operatorname{crs}, D), m_{2} \leftarrow \operatorname{EncProg}\left(\operatorname{crs}, m_{1},(P, x, t)\right)$.

- Receiver Privacy: For every pair of databases $D_{0} \in\{0,1\}^{M}, D_{1} \in\{0,1\}^{M}$ where $M$ is polynomial in $\lambda$, for every crs in the support of $\operatorname{Setup}\left(1^{\lambda}\right)$, let $\left(m_{0}, \widetilde{D}_{0}\right) \leftarrow \operatorname{EncData}\left(\operatorname{crs}, D_{0}\right)$, $\left(m_{1}, \widetilde{D}_{1}\right) \leftarrow \operatorname{EncData}\left(\mathrm{crs}, D_{1}\right)$. Then it holds that

$$
\left(\mathrm{crs}, m_{0}\right) \stackrel{\mathrm{c}}{\approx}\left(\mathrm{crs}, m_{1}\right) .
$$

- Sender Privacy: There exists a PPT simulator niscSim such that for every database $D \in$ $\{0,1\}^{M}$ where $M=\operatorname{poly}(\lambda)$ for any polynomial function poly $(\cdot)$, and for every RAM program $(P, x, t)$, let $y=P^{D}(x)$ be the output of the program, and MemAccess be the memory access pattern, then it holds that

$$
\left(\operatorname{crs}, D,\left(m_{1}, \widetilde{D}\right), m_{2}\right) \stackrel{\mathrm{c}}{\approx} \operatorname{niscSim}\left(1^{\lambda}, D,(y, \text { MemAccess })\right)
$$

where $\operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}\right),\left(m_{1}, \widetilde{D}\right) \leftarrow \operatorname{EncData}(\operatorname{crs}, D)$ and $m_{2} \leftarrow \operatorname{EncProg}\left(\operatorname{crs}, m_{1},(P, x, t)\right)$.

- Efficiency: The length of $m_{1}$ is a fixed polynomial in $\lambda$ independent of the size of the database. Moreover, the algorithm EncData runs in time $M \cdot \operatorname{poly}(\lambda, \log M)$, EncProg and Dec run in time $t \cdot \operatorname{poly}(\lambda, \log M)$.


### 6.3 Construction

Overview. We first give an overview of the construction. For ease of exposition, consider a readonly program where each CPU step outputs the next location to be read based on the value read from last location.

We first describe the EncProg procedure. As already mentioned in technical overview (see Section 2.2), our construction is based on high level ideas of garbled RAM (introduced by Lu and Ostrovsky [LO13]) to make sender and receiver complexity grow only with the running time of the program. In particular, the sender would generate a garbled RAM program consisting of a sequence of $t$ garbled step circuits. Similar to the RAM computation model described in Section 6.1.1, every step circuit takes as input the current CPU state and the last read bit and outputs the updated state and the next read location, say $L$. Note that the next step circuit would take the new value read from database as input.

The main challenge in program garbling is revealing the correct labels for the next circuit based on the value of $D[L]$. Moreover, it is crucial for garble circuit security that the receiver does not learn the label corresponding to $1-D[L]$. Prior works [LO13, GHL ${ }^{+} 14$, GLOS15, GLO15] proposed several different solutions to the above problem. Here we present a new and arguably simpler solution for achieving this using laconic oblivious transfer.

Let digest be the hash value of $D$ that would be fed into the first step circuit and passed along the sequence of circuits. That is, each circuit would take this digest as input and also output the correct input labels corresponding to the digest for the next circuit. Now, to transfer the correct label corresponding to the value in the database, a step circuit would output a laconic OT ciphertext (using algorithm Send) that encrypts the input keys of the next step circuit and uses digest as the hash value. Looking ahead, when the receiver evaluates the step circuit which outputs the laconic OT ciphertext, he can use $D$ to decrypt it to obtain the correct labels (using the procedure Receive of laconic OT).

We would show that the sender privacy follows from the sender privacy of laconic OT and security of circuit garbling. In order to achieve receiver privacy, the receiver does not publish digest in the clear, but instead, the labels for digest of the first step circuit are transferred from the sender to the receiver via a two-message OT. In particular, the EncData procedure outputs the first OT message of digest, and EncProg will output the garbled step circuits along with the second OT message for digest's labels.

Finally, note that a general program can also write to the database, in which case we need to update the database as well as the step circuits need to know the updated digest for the correctness of laconic OT and future reads/writes. This is achieved via the updatability property of the laconic OT which allows a sender to generate a ciphertext that allows the receiver to learn messages corresponding to the updated digest. In our case, the messages encrypted would be the input digest keys of the next step circuit.

Next, we give a more formal construction of our scheme.

The Construction. Let $\ell O T=$ (crsGen, Hash, Send, Receive, SendWrite, ReceiveWrite) be an updatable laconic OT protocol as per Definition 3.2. Let OT $=\left(\mathrm{OT}_{1}, \mathrm{OT}_{2}, \mathrm{OT}_{3}\right)$ be a two-message secure oblivious transfer, and let $\mathrm{GC}=$ (GCircuit, Eval) be a circuit garbling scheme. The noninteractive secure RAM computation scheme NISC-RAM $=$ (Setup, EncData, EncProg, Dec) is constructed as follows.

Setup: $\quad \operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}\right)$.
The set up algorithm is described in Figure 7. It generates the common reference string for the updatable laconic OT scheme.

Set up. $\quad \operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}\right)$.

1. $\operatorname{crs} \leftarrow \operatorname{crsGen}\left(1^{\lambda}\right)$.
2. Output crs.

Figure 7: Set up procedure of NISC-RAM

Database Encryption: $\quad\left(m_{1}, \widetilde{D}\right) \leftarrow$ EncData $($ crs,$D)$.
The algorithm is formally described in Figure 8. It hashes the database $D$ using laconic OT Hash function and obtains digest. Then digest is encrypted using the $\mathrm{OT}_{1}$ procedure of two-message OT protocol.

Database Encryption. $\quad\left(m_{1}, \widetilde{D}\right) \leftarrow$ EncData $($ crs,$D)$.

1. (digest, $\hat{\mathrm{D}}) \leftarrow \operatorname{Hash}($ crs,$D)$.
2. $\left(m_{1}\right.$, secret $) \leftarrow \mathrm{OT}_{1}\left(1^{\lambda}\right.$, digest $)$.
3. Output $\left(m_{1}, \widetilde{D}=(\right.$ digest,$\hat{\mathrm{D}}$, secret $\left.)\right)$.

Figure 8: Database encryption procedure of NISC-RAM

Program Encryption: $\quad m_{2} \leftarrow$ EncProg $\left(\operatorname{crs}, m_{1},(P, x, t)\right)$.
The program encryption procedure is formally described in Figure 9. As mentioned above, it generates $t$ garbled step circuits $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=1}^{t}$, where every step circuit implements the functionality of a CPU-step circuit. We describe the structure of a step circuit $C^{\text {step }}$ below. The program encryption also consists of the second OT message corresponding to the short message $m_{1}$ of the receiver (for digest) where the sender's messages consist of the input keys for the first garbled circuit. Finally, it also outputs the keys for decrypting the output of the last step circuit.

Now we elaborate on the logic of a step circuit. The pseudocode of a step circuit $C^{\text {step }}$ is formally described in Figure 10, and the structure is illustrated in Figure 11. The input of a step circuit can be partitioned into (state, rData, digest), where state is the current CPU state, rData is the bit read from the database, and digest is the up-to-date digest of the database. If the previous step is a write, then $\mathrm{rData}=0$. The program encryption outputs garbled circuits for these step circuits, hence, the first step of EncProg is to pick the input keys for all the circuits. The $\tau$-th step circuit $C_{\tau}^{\text {step }}$ has hardwired in it the input keys nextKeys $=$ (stateKeys, dataKeys, digestKeys) for the next step circuit $\mathrm{C}_{\tau+1}^{\text {step }}$.

The logic of the step circuit is as follows: It first computes the new (state ${ }^{\prime}, \mathrm{R} / \mathrm{W}, L, \mathrm{wData}$ ). Then, in the case of a "read" it outputs stateKeys corresponding to state', labels for rData via laconic OT procedure Send $(\cdot)$, and digestKeys corresponding to digest. The case of a write is similar, but now the labels of new updated digest are transferred via laconic OT procedure SendWrite(•).

Program Encryption. $\quad m_{2} \leftarrow$ EncProg (crs, $\left.m_{1},(P, x, t)\right)$.

1. Generate the garbled program for $P$ : Generate garbled circuits $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=1}^{t}$.
(a) Sample stateKeys ${ }^{\tau}$, dataKeys $^{\tau}$, digestKeys $^{\tau}$ for each $\tau \in\{1, \ldots, t+1\}$.
(b) For each $\tau \in\{1, \ldots, t\}$

$$
\widetilde{\mathrm{C}}_{\tau}^{\text {step }} \leftarrow \operatorname{GCircuit}\left(1^{\lambda}, \mathrm{C}^{\text {step }}\left[\mathrm{crs}, P, \text { Keys }^{\tau+1}\right], \text { Keys }^{\tau}\right),
$$

where Keys $^{\tau}=\left(\right.$ stateKeys $^{\tau}$, dataKeys $^{\tau}$, digestKeys $\left.^{\tau}\right)$.
(c) For $\tau=1$, embed labels dataKeys ${ }_{0}^{1}$ and stateKeys ${ }_{x}^{1}$ in $\widetilde{\mathrm{C}}_{1}^{\text {step }}$.
2. Compute $\mathrm{L} \leftarrow \mathrm{OT}_{2}\left(m_{1}\right.$, digestKeys $\left.{ }^{1}\right)$.
3. Output $m_{2}=\left(\mathrm{L},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=1}^{t}\right.$, stateKeys $\left.^{t+1}\right)$.

Figure 9: Program encryption procedure of NISC-RAM
Hardwired Parameters: [crs, $P$, nextKeys $=$ (stateKeys, dataKeys, digestKeys)].
Input: (state, rData, digest).

```
    (state \(\left.{ }^{\prime}, \mathrm{R} / \mathrm{W}, L, \mathrm{wData}\right):=C_{\mathrm{CPU}}^{P}(\) state, rData\()\).
    if \(R / W=\) read then
    \(\mathrm{e}_{\text {data }} \leftarrow\) Send(crs, digest, \(L\), dataKeys).
    return \(\left(\left(\right.\right.\) stateKeys \(_{\text {state }^{\prime}}, \quad e_{\text {data }}, \quad\) digestKeys \(\left.\left._{\text {digest }}\right), \mathrm{R} / \mathrm{W}, L\right)\).
    else
        \(\mathrm{e}_{\text {digest }} \leftarrow\) SendWrite (crs, digest, \(L\), wData, digestKeys).
        return ((stateKeys state,\(\quad\) dataKeys \(\left.\left._{0}, \quad e_{\text {digest }}, w D a t a\right), R / W, L\right)\).
```

Figure 10: Pseudocode of a step circuit $\mathrm{C}^{\text {step }}$ [crs, $P$, nextKeys].
Decryption: $\quad y \leftarrow \operatorname{Dec}^{\widetilde{D}}\left(\mathrm{crs}, m_{2}\right)$.
The decryption procedure is described in Figure 13. At a high level the receiver evaluates the garbled step circuits one by one from $\widetilde{\mathrm{C}}_{1}^{\text {step }}$ to $\widetilde{\mathrm{C}}_{t}^{\text {step }}$, and uses the database to decrypt $\ell O T$ ciphertexts between two consecutive circuits. The output of the last step circuit can be decrypted using stateKeys ${ }^{t+1}$ and hence $y$ is obtained.

More precisely, the receiver first obtains the digestLabels for the first step circuit by running $\mathrm{OT}_{3}$. Note that the first garbled step circuit already has labels for the rData and state embedded. Hence the receiver can obtain all the labels for the first step circuit and evaluate it. Then the receiver executes the circuits $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=1}^{t}$ one by one, and learns the labels for the next circuit by running the receiver algorithms of laconic OT on its database.

### 6.4 Correctness

For correctness, we require that for every database $D \in\{0,1\}^{M}$, for every RAM program $(P, x, t)$, it holds that

$$
\operatorname{Pr}\left[\operatorname{Dec}{ }^{\widetilde{D}}\left(\mathrm{crs}, m_{2}\right)=P^{D}(x)\right]=1
$$



Figure 11: A step circuit $\mathrm{C}^{\text {step }}$ [crs, $P$, nextKeys]
where crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}\right),\left(m_{1}, \widetilde{D}\right) \leftarrow \operatorname{EncData}(\operatorname{crs}, D), m_{2} \leftarrow \operatorname{EncProg}\left(\operatorname{crs}, m_{1},(P, x, t)\right)$. Correctness follows from Lemma 6.3 that we will prove below.
Claim 6.2. The first garbled step circuit $\widetilde{\mathrm{C}}_{1}^{\text {step }}$ gets evaluated on ( $x, 0$, digest), where (digest, $\left.\hat{\mathrm{D}}\right)=$ Hash (crs, D).

Proof. Since $\left(m_{1}\right.$, secret $) \leftarrow \mathrm{OT}_{1}\left(1^{\lambda}\right.$, digest $), \mathrm{L} \leftarrow \mathrm{OT}_{2}\left(m_{1}\right.$, digestKeys $\left.{ }^{1}\right)$, and digestLabels ${ }^{1} \leftarrow$ $\mathrm{OT}_{3}(\mathrm{~L}$, secret $)$, by correctness of OT, digestLabels ${ }^{1}=$ digestKeys ${ }_{\text {digest }}^{1}$. Moreover, $\widetilde{\mathrm{C}}_{1}^{\text {step }}$ already has labels stateKeys ${ }_{x}^{1}$ and dataKeys ${ }_{0}^{1}$ embedded in it, by correctness of the circuit garbling scheme, $\widetilde{\mathrm{C}}_{1}^{\text {step }}$ gets evaluated on ( $x, 0$, digest).

Lemma 6.3. Consider the execution of $P^{D}(x)$. Let (state ${ }^{\tau}, \mathrm{rData}^{\tau}$ ) be the input to the $\tau$-th $C P U$ step. Let $D^{\tau}$ be the database at the beginning of step $\tau$, and let $\left(\operatorname{digest}^{\tau}, \hat{\mathrm{D}}^{\tau}\right)=\operatorname{Hash}\left(\mathrm{crs}, D^{\tau}\right)$. During the Dec procedure, for every $\tau \in[t], \widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ is evaluated on inputs (state ${ }^{\tau}$, $\mathrm{rData}^{\tau}$, digest $^{\tau}$ ). Moreover, the state of the database held by the receiver at the beginning of evaluating $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ is $\hat{\mathrm{D}}^{\tau}$.

Proof. We will prove this lemma by induction on $\tau$. The base case follows from Claim 6.2. Assume that the lemma holds for $\tau=\rho$, then we prove that the lemma holds for $\rho+1$ in the following. We know that $\left(\hat{\mathrm{D}}^{\rho}\right.$, digest $\left.^{\rho}\right)=\operatorname{Hash}\left(\mathrm{crs}, D^{\rho}\right)$, and that $\widetilde{\mathrm{C}}_{\rho}^{\text {step }}$ is executed on (state ${ }^{\rho}, \mathrm{rData}^{\rho}$, digest $\left.^{\rho}\right)$. By correctness of GC, $\widetilde{\mathrm{C}}_{\rho}^{\text {step }}$ implements its code of a CPU step, namely (state $, \mathrm{R} / \mathrm{W}, L$, wData) $=$ $C_{\mathrm{CPU}}^{P}\left(\right.$ state $^{\rho}$, rData $\left.^{\rho}\right)$. Also notice that nextKeys $=($ stateKeys, dataKeys, digestKeys) hardwired in $\widetilde{\mathrm{C}}_{\rho}^{\text {step }}$ are the input keys for $\widetilde{\mathrm{C}}_{\rho+1}^{\text {step }}$. There are two cases:

- $\mathrm{R} / \mathrm{W}=$ read: In this case, it follows directly from the Dec procedure that stateLabels ${ }^{\rho+1}=$ stateKeys $_{\text {state }^{\prime}}$ and digestLabels ${ }^{\rho+1}=$ digestKeys $_{\text {digest }}$. Since $\mathrm{e}_{\text {data }} \leftarrow$ Send $\left(\right.$ crs, digest $^{\rho}, L$, dataKeys) and dataLabels ${ }^{\rho+1}=$ Receive $^{\hat{D}^{\rho}}\left(\mathrm{crs}, \mathrm{e}_{\text {data }}, L\right)$, by correctness of the $\ell O T$ scheme, dataLabels ${ }^{\rho+1}=$ dataKeys $_{D^{\rho}[L]}$. Hence $\widetilde{C}_{\rho+1}^{\text {step }}$ is evaluated on inputs (state ${ }^{\prime}, D^{\rho}[L]$, digest), which is exactly $\left(\right.$ state $^{\rho+1}$, rData $^{\rho+1}$, digest $^{\rho+1}$ ). And ( $\hat{\mathrm{D}}^{\rho}$, digest $\left.^{\rho}\right)$ remains unchanged.

Decryption. $y \leftarrow \operatorname{Dec}^{\widetilde{D}}\left(\mathrm{crs}, m_{2}\right)$.

1. Parse $\widetilde{D}=($ digest, $\hat{\mathrm{D}}$, secret $)$.
2. Parse $m_{2}=\left(\mathrm{L},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=1}^{t}\right.$, stateKeys $\left.^{t+1}\right)$.
3. Compute digestLabels ${ }^{1} \leftarrow \mathrm{OT}_{3}(\mathrm{~L}$, secret $)$.
4. Parse $\widetilde{\mathrm{C}}_{1}^{\text {step }}=\left(\widetilde{\mathrm{C}}_{1}^{\text {step }}\right.$, dataLabels ${ }^{1}$, stateLabels $\left.{ }^{1}\right)$.
5. For $\tau=1$ to $t$ do the following:

$$
(X, \mathrm{R} / \mathrm{W}, L):=\operatorname{Eval}\left(\widetilde{\mathrm{C}}_{\tau}^{\widetilde{s t e p}^{s}},\left(\text { stateLabels }^{\tau}, \text { dataLabels }^{\tau}, \operatorname{digestLabels}^{\tau}\right)\right)
$$

if $R / W=$ read then
Parse $X=\left(\right.$ stateLabels $^{\tau+1}, \mathrm{e}_{\text {data }}$, digestLabels $\left.{ }^{\tau+1}\right)$
dataLabels $^{\tau+1}=$ Receive $^{\hat{\mathrm{D}}}\left(\right.$ crs $\left., \mathrm{e}_{\text {data }}, L\right)$
else
Parse $X=\left(\right.$ stateLabels ${ }^{\tau+1}$, dataLabels $^{\tau+1}, \mathrm{e}_{\text {digest }}$, wData $)$
digestLabels ${ }^{\tau+1}=\operatorname{ReceiveWrite}^{\hat{\mathrm{E}}}$ (crs, $L$, wData, $\mathrm{e}_{\text {digest }}$ )
6. Use stateKeys ${ }^{t+1}$ to decode stateLabels ${ }^{t+1}$ and obtain $y$.

Figure 13: Decryption procedure of NISC-RAM

- $\mathrm{R} / \mathrm{W}=$ write: In this case, it follows from the Dec procedure that stateLabels ${ }^{\rho+1}=$ stateKeys $_{\text {state }}{ }^{\prime}$ and dataLabels ${ }^{\rho+1}=$ dataKeys $_{0}$. Since $\mathrm{e}_{\text {digest }} \leftarrow$ SendWrite $\left(\right.$ crs, digest ${ }^{\rho}, L$, wData, digestKeys) and digestLabels ${ }^{\rho+1}=$ ReceiveWrite $^{\hat{\mathrm{D}}^{\rho}}$ (crs, $L$, wData, $\mathrm{e}_{\text {digest }}$ ), by correctness of the $\ell O T$ scheme, digestLabels $^{\rho+1}=$ digestKeys $_{\text {digest }^{\prime}}$ where $\left(\hat{\mathrm{D}}^{\prime}\right.$, digest $\left.^{\prime}\right)=$ Hash $\left(\mathrm{crs}, D^{\prime}\right)$ for an updated database $D^{\prime}\left(D^{\prime}\right.$ is identical to $D^{\rho}$ except that $D^{\prime}[L]=w$ Data $)$. Hence $\widetilde{\mathrm{C}}_{\rho+1}^{\text {step }}$ is evaluated on inputs (state ${ }^{\prime}, 0$, digest ${ }^{\prime}$ ), which is exactly ( state $^{\rho+1}, \mathrm{rData}^{\rho+1}$, digest ${ }^{\rho+1}$ ). And ( $\hat{\mathrm{D}}^{\rho}$, digest $^{\rho}$ ) gets updated to ( $\hat{\mathrm{D}}^{\prime}$, digest ${ }^{\prime}$ ), which is exactly ( $\hat{\mathrm{D}}^{\rho+1}$, digest $^{\rho+1}$ ).


### 6.5 Security Proof

In this section we prove sender privacy and receiver privacy as defined in Section 6.2 under the decisional Diffie-Hellman (DDH) assumption. The receiver privacy follows directly from the receiver security of OT. Below we prove sender privacy by describing a PPT simulator niscSim such that for every database $D \in\{0,1\}^{M}$ where $M$ is polynomial in $\lambda$, and for every RAM program ( $P, x, t$ ), let $y=P^{D}(x)$ be the output of the program, and MemAccess be the memory access pattern, then it holds that

$$
\left(\operatorname{crs},\left(m_{1}, \widetilde{D}\right), \operatorname{EncProg}\left(\operatorname{crs}, m_{1},(P, x, t)\right)\right) \stackrel{\mathrm{c}}{\approx}\left(\operatorname{crs},\left(m_{1}, \widetilde{D}\right), \operatorname{niscSim}\left(\operatorname{crs}, m_{1}, D, y, \text { MemAccess }\right)\right)
$$

where crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}\right),\left(m_{1}, \widetilde{D}\right) \leftarrow \operatorname{EncData}($ crs, $D)$. Notice that this definition is slightly different from the definition in Section 6.2, but in the semi-honest case it implies a simulator as defined in Section 6.2

1. Sample input keys (stateKeys ${ }^{t+1}$, dataKeys ${ }^{t+1}$, digestKeys $^{t+1}$ ) for $\mathrm{C}^{\text {step }}$.
2. Parse MemAccess as $\left\{\left(\mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}, \mathrm{wData}^{\tau}\right): \tau \in[t]\right\}$, where $\left(\mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}, \mathrm{wData}^{\tau}\right)$ is partial output of the $\tau$-th CPU step circuit. Compute ( $\mathrm{rData}^{\tau}, D^{\tau}$, $\mathrm{digest}^{\tau}$ ) at the beginning of step $\tau$ for every $\tau \in[t+1]$.
3. Compute (stateLabels ${ }^{t+1}$, dataLabels ${ }^{t+1}$, digestLabels $^{t+1}$ ):

$$
\begin{aligned}
& \text { stateLabels }^{t+1} \leftarrow \text { stateKeys }_{y}^{t+1} . \\
& \text { digestLabels }^{t+1} \leftarrow \text { digestKeys }_{\text {digest }}^{t+1} \\
& \text { dataLabels }^{t+1} \leftarrow \text { dataKeys }_{\text {rData }^{t+1}} .
\end{aligned}
$$

4. For $\tau=t$ downto 1, proceed as follows:
if $\mathrm{R} / \mathrm{W}^{\tau}=$ read then
$\mathrm{e}_{\text {data }} \leftarrow \ell$ OTSim $\left(\mathrm{crs}, D^{\tau}, L^{\tau}\right.$, dataLabels $\left.{ }^{\tau+1}\right)$.
$X \leftarrow\left(\right.$ stateLabels ${ }^{\tau+1}, \mathrm{e}_{\text {data }}$, digestLabels
else
$\mathrm{e}_{\text {digest }} \leftarrow \ell \mathrm{OTSimWrite}\left(\mathrm{crs}, D^{\tau}, L^{\tau}\right.$, wData $^{\tau}$, digestLabels $\left.{ }^{\tau+1}\right)$.
$X \leftarrow\left(\right.$ stateLabels $^{\tau+1}$, dataLabels $\left.^{\tau+1}, \mathrm{e}_{\text {digest }}, \mathrm{wData}^{\tau}\right)$.
$\left(\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right.$, stateLabels ${ }^{\tau}$, dataLabels $^{\tau}$, digestLabels $\left.{ }^{\tau}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {step }},\left(X, \mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}\right)\right)$.
5. $\mathrm{L} \leftarrow \mathrm{OT}_{2}\left(m_{1},\left(\right.\right.$ digestLabels $^{1}$, digestLabels $\left.\left.^{1}\right)\right)$.
6. Output $\left(\mathrm{L},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=1}^{t}\right.$, stateKeys $\left.^{t+1}\right)$.

We show that the above simulation is indistinguishable from the real execution through a sequence of hybrids where the first hybrid outputs the real execution and the last hybrid outputs the simulated one.

- $\mathrm{H}_{2 i}$ for $i \in\{0,1, \ldots, t\}$ : Notice that in the output, there are $t$ garbled step circuits $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=1}^{t}$. In hybrid $\mathrm{H}_{2 i}$, the garbled step circuits from 1 to $i$ are simulated while the remaining step circuits $(i+1$ to $t)$ are generated honestly. Given the program, all the intermediate outputs of every step circuit can all be computed. Given the correct output of circuit $C_{i}^{\text {step }}$, the step circuits from 1 to $i$ can be simulated one by one from the $i$-th to the first similarly as niscSim. More formally, it proceeds as follows.

1. Execute $P^{D}(x)$ to obtain ( $\mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}$, $\mathrm{wData}^{\tau}$ ) for every $\tau \in[t]$ and state $^{t+1}=y$. Compute ( $\mathrm{rData}{ }^{\tau}, D^{\tau}$, digest $^{\tau}$ ) at the beginning of step $\tau$ for every $\tau \in[t+1]$.
2. Generate the garble circuits $\left\{\widetilde{\mathcal{C}}_{\tau}^{\text {step }}\right\}_{\tau=i+1}^{t}$ honestly (same as Step 1 in EncProg).
3. Let (stateKeys ${ }^{i+1}$, dataKeys $^{i+1}$, digestKeys ${ }^{i+1}$ ) be the input keys of $\widetilde{\mathrm{C}}_{i+1}^{\text {step }}$.
4. Compute (stateLabels ${ }^{i+1}$, dataLabels ${ }^{i+1}$, digestLabels ${ }^{i+1}$ ):

$$
\begin{aligned}
& \text { stateLabels }^{i+1} \leftarrow \text { stateKeys }_{\text {statate }_{i+1}^{i+1}}^{i+1} \\
& \text { digestLabels }^{i+1} \leftarrow \text { digestKeys }_{\text {digest }}^{+1+1} \\
& \text { dataLabels }
\end{aligned}
$$

5. For $\tau=i$ downto 1, proceed as in Step 4 of the simulator niscSim.
6. $\mathrm{L} \leftarrow \mathrm{OT}_{2}\left(m_{1},\left(\right.\right.$ digestLabels $^{1}$, digestLabels $\left.\left.{ }^{1}\right)\right)$.
7. Output (L, $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=1}^{t}$, stateKeys $\left.^{t+1}\right)$.

- ${\underset{\widetilde{\sim}}{2 i+1}}^{\mathrm{H}_{2}}$ for $i \in\{0, \ldots, t-1\}$ : Hybrid $\mathrm{H}_{2 i+1}$ is identical to $\mathrm{H}_{2 i}$ except that $\mathrm{H}_{2 i+1}$ simulates $\widetilde{\mathrm{C}}_{i+1}^{\text {step }}$ based on the real output of $\mathrm{C}_{i+1}^{\text {step }}$. In particular, $\mathrm{H}_{2 i+1}$ is the same as $\mathrm{H}_{2 i}$ except that Steps 2, 3, 4 proceed as follows:

2. Generate the garble circuits $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau=i+2}^{t}$ honestly (same as Step 1 in EncProg).
3. Let $\left(\right.$ stateKeys $^{i+2}$, dataKeys ${ }^{i+2}$, digestKeys ${ }^{i+2}$ ) be the input keys of $\widetilde{\mathrm{C}}_{i+2}^{\text {step }}$.
4. if $\mathrm{R} / \mathrm{W}^{i+1}=$ read then

$$
\begin{aligned}
& \mathrm{e}_{\text {data }} \leftarrow \text { Send }\left(\text { crs, } \text { digest }^{i+1}, L^{i+1} \text {, dataKeys }{ }^{i+2}\right) \text {. } \\
& X \leftarrow\left(\text { stateKeys }_{\text {state }^{i+2}}^{i+2}, \text { e data } \text {, digestKeys }{ }_{\text {digest }}^{i+2}\right) . \\
& \text { else } \\
& \mathrm{e}_{\text {digest }} \leftarrow \text { SendWrite }\left(\text { crs, } \text { digest }^{i+1}, L^{i+1}, \text { wData }^{i+1} \text {, digestKeys }{ }^{i+2}\right) . \\
& X \leftarrow\left(\text { stateKeys }_{\text {state }^{i+1}}^{i+1}, \text { dataKeys }_{0}^{i+1}, \text { edigest }, \text { wData }^{i+1}\right) . \\
& \left(\widetilde{\mathrm{C}}_{i+1}^{\text {step }}, \text { stateLabels }{ }^{i+1}, \text { dataLabels }^{i+1}, \text { digestLabels }^{i+1}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathrm{C}^{\text {step }},\left(X, \mathrm{R} / \mathrm{W}^{i+1}, L^{i+1}\right)\right) .
\end{aligned}
$$

It is easy to see that $H_{0}$ is the output of the real execution, and $\mathrm{H}_{2 t}$ is the simulated output. Now we prove that the consecutive hybrids are computationally indistinguishable. Below we prove that $\mathrm{H}_{2 i} \stackrel{\mathrm{c}}{\approx} \mathrm{H}_{2 i+1} \stackrel{c}{\approx} \mathrm{H}_{2(i+1)}$ for every $i \in\{0, \ldots, t-1\}$. Since hybrid $\mathrm{H}_{2 i+1}$ simulates $\widetilde{\mathrm{C}}_{i+1}^{\text {step }}$ based on the real output of $\mathrm{C}_{i+1}^{\text {step }}$, the output of $\widetilde{\mathrm{C}}_{i+1}^{\text {step }}$ is identical for hybrids $\mathrm{H}_{2 i}$ and $\mathrm{H}_{2 i+1}$. That said, indistinguishibility of hybrids $\mathrm{H}_{2 i}$ and $\mathrm{H}_{2 i+1}$ follows from the garbled circuit security. Next, indistinguishability between $\mathrm{H}_{2 i+1}$ and $\mathrm{H}_{2(i+1)}$ follows from the sender's privacy property of the updatable laconic OT since the laconic OT responses are simulated in $\mathrm{H}_{2(i+1)}$. This concludes the proof.

### 6.6 Extension

For simplicity of exposition, the protocol we described so far is for a single sender executing a single program with the receiver. It can be extended to the setting where a sender can execute a sequence of programs on a persistent database. Moreover, the message $m_{1}$ published by the receiver can be used by multiple senders, in which case the receiver maintains a different copy of the database for every sender.

Executing multiple programs on a persistent database. After receiving the first message $m_{1}$ from the receiver, a sender can run multiple programs on a persistent database (with initial content $D$ ) by sending one message per program to the receiver. For security we require only the output and the memory access pattern of every program execution are revealed to the receiver. We also require that once $m_{1}$ is published, the computational cost of both the sender and the receiver for every program should grow only with the running time of the RAM computation, and is independent of the size of $D$. The NISC-RAM scheme we constructed in Section 6.3 can be naturally extended to the multi-program setting. We explain the extension by describing the changes of EncProg and Dec procedures for the second program. Encryption and evaluation of more programs would follow analogously.

- EncProg: When encrypting the first program, the sender should store locally digestKeys* $=$ digestKeys ${ }^{t+1}$. Then, when encrypting the second program, there are two changes in EncProg compared to encrypting the first program: (1) digestKeys* is used as the digest keys of the first step circuit, (2) L is not generated.
- Dec: When evaluating the first program, the receiver should store locally digestLabels* $=$ digestLabels ${ }^{t+1}$. Then, when evaluating the second program, the sender should use digestLabels* as the digest labels for the first step circuit.

Multiple senders with a single receiver. The above protocol also works for multiple parallel senders. That is, after the receiver publishes the first message $m_{1}$, every sender $S$ can send a message $m_{S}$ to the receiver enabling the execution of $P_{S}^{D}\left(x_{S}\right)$, where $D$ is the initial database of the receiver, and $\left(P_{S}, x_{S}\right)$ is the program of $S$. Security follows from the security of OT which supports multiple second OT messages with the same first OT message. Moreover, every sender can execute a sequence of programs on a persistent database. In this case, the receiver keeps a different copy of her initial database for every sender.

## 7 Main Application: Multi-Hop Homomorphic Encryption for RAM Programs

### 7.1 Our Model

Consider a server $S$ and a collection of clients $Q_{1}, Q_{2}, \ldots$ with private databases $D_{1}, D_{2}, \ldots$, respectively. The clients ship their encrypted databases to $S$ to be computed on later in multiple executions in a persistent manner. At the beginning of any execution, the server $S$ encrypts his private input $x$ as $\mathrm{ct}_{0}$, chooses a subset of clients $Q_{i_{1}}, \ldots, Q_{i_{n}}$ and sends the $\mathrm{ct}_{0}$ to client $Q_{i_{1}}$. Next, for all $j \in[n]$, client $Q_{i}$ homomorphically evaluates an arbitrary program $P_{j}$ of his choice on $\mathrm{ct}_{j-1}$ to obtain $\mathrm{ct}_{j}$. Finally, client $Q_{i_{n}}$ sends $\mathrm{ct}_{n}$ to the server $S$. The server decrypts this ciphertext using his secret key of encryption as well as encrypted databases sent earlier to learn $P_{n}^{D_{i_{n}}}\left(\ldots P_{1}^{D_{i_{1}}}(x) \ldots\right)$. During this execution, the databases get updated and future execution of any client happens on respective updated databases.

We require that the size of the ciphertext only grows with the cumulative running time of all programs in an execution and is independent of the size of the databases. For security, we require program and data privacy for all honest clients against an adversary that corrupts the server and any subset of the clients. Next, we describe the model formally.

We say that an ordered sequence of RAM programs $P_{1}, \cdots, P_{n}$ are compatible if the output length of $P_{i}$ is the same as the input length of $P_{i+1}$ for every $i \in[n-1]$. A multi-hop RAM homomorphic encryption scheme mhop-RAM $=($ Setup, KeyGen, InpEnc, EncData, Eval, Dec) has the following syntax. We define the algorithms w.r.t. clients $Q_{1}, \ldots, Q_{n}$.

- Setup: crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}\right)$.

On input the security parameter $1^{\lambda}$, it outputs a common reference string.

- Key Generation: $(\mathrm{pk}, \mathrm{sk}) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}\right)$.

On input the security parameter $1^{\lambda}$, it outputs a public/secret key pair ( $\mathrm{pk}, \mathrm{sk}$ ).

- Database Encryption: $(\widetilde{D}=(\hat{D}$, digest $)) \leftarrow \operatorname{EncData}($ crs,$D)$.

On input the common reference string crs and database $D \in\{0,1\}^{M}$, it outputs an encrypted database $\widetilde{D}=(\hat{D}$, digest $)$, where digest is a short digest of the database.

- Input Encryption: (ct, x_secret) $\leftarrow \operatorname{InpEnc}(x)$.

On input $S$ 's input $x$, it outputs a ciphertext ct and secret state for $S$ denoted by x_secret.

- Homomorphic Evaluation: $\mathrm{ct}^{\prime} \leftarrow \operatorname{Eval}\left(\mathrm{crs}, i,\left\{\mathrm{pk}_{j}\right\}_{j=i+1}^{n}, \mathrm{ct}, \mathrm{sk},(P, t)\right.$, digest $)$.

It takes as input the crs, the client number $i$, the public keys of the clients later in the sequence, i.e., $Q_{i+1}, \cdots, Q_{n}$, a ciphertext from the previous client, $Q$ 's secret key sk, $Q$ 's RAM program $P$ with maximum run-time $t$ and the digest digest of the database $D$ of $Q$. It then outputs a new ciphertext ct'.

- Decryption: $y=\operatorname{Dec}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}$ (crs, x_secret, ct).

On input the crs, server's state x_secret, the final ciphertext ct from client $Q_{n}$, and RAM access to encrypted databases $\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}$, it outputs $y$. The procedure Dec is itself modeled as a RAM program that can read and write to arbitrary locations of its database initially containing $\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}$.

Next, we describe how these algorithms are used in a real execution.
Real Scenario. In our multi-hop scheme for RAM programs, after the initialization phase that generates the common-random string crs, each cleint runs key generation to generate the public key and the secret key, followed by the database encryption. The encrypted database is sent to the server, and the cleint stores the digest of the database locally. After this initialization phase, the server $S$ can initiate various executions of RAM computations with different subsets of the clients. After each execution, the database of a client gets updated by the server during the decryption phase. It is ensured that the server also learns the updated digest of the database that is communicated to the clients during the start of the next execution.

At the onset of any execution, the server $S$ encrypts his input and sends the ciphertext ct ${ }_{0}$ to the first client $Q_{1}$ and maintains x_secret to be used later. The client $Q_{1}$ generates the ciphertext $\mathrm{ct}_{1}$ using his program $P_{1}$ and digest digest ${ }_{1}$ and sends it to $Q_{2}$. Similarly, when a client $Q_{i}$ receives $\mathrm{ct}_{i-1}$ from $Q_{i-1}$, it uses program $P_{i}$ and digest ${ }_{i}$ to generate $\mathrm{ct}_{i}$ and sends it to $Q_{i+1}$. This continues and finally, $Q_{n}$ sends $\mathrm{ct}_{n}$ back to the server $S$. Then, the server runs the decryption procedure on $\mathrm{ct}_{n}$ using all the encrypted databases and secret state x_secret to obtain output $y$.

For the case of multiple executions, each of the above procedures take the session identifier sid as additional input. We denote by $\widetilde{D}_{1}^{\text {(sid) }}, \cdots, \widetilde{D}_{n}^{\text {(sid) }}$ the encrypted databases before the execution with session identifier sid. Initially $\widetilde{D}_{1}^{(1)}=\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}^{(1)}=\widetilde{D}_{n}$.

We require the algorithms above to satisfy the correctness, sender-privacy, client-privacy and efficiency properties described below.

Correctness. We require that in a sequence of executions, each output of homomorphic evaluation equals the output of the corresponding computation in the clear. We formalize this as follows: For every set of keys $\left\{\left(\mathrm{pk}_{i}, \mathrm{sk}_{i}\right)\right\}_{i=1}^{n}$ in the support of KeyGen, and any collection of initial databases $D_{1}, \cdots, D_{n}$, for any unbounded polynomial N number of executions the following holds:

For sid $\in \mathrm{N}$, let $P_{1}^{(\text {sid })}, \ldots, P_{n}^{(\text {sid })}$ be the sequence of programs, $x^{(\text {sid })}$ be the server's input and $D_{i}^{(\text {sid })}$ be the resulting database after executing the session sid-1 in the clear, then

$$
\operatorname{Pr}\left[\operatorname{Dec}^{\widetilde{D}_{1}^{(\text {sid })}, \ldots, \widetilde{D}_{n}^{(\text {sid })}}\left(\operatorname{crs}, \mathrm{x} \_\operatorname{secret}^{(\text {sid })}, \mathrm{ct}_{n}^{(\text {sid })}\right)=P_{n}^{(\text {sid })} D_{n}^{(\text {sid })}\left(\cdots\left(P_{1}^{(\text {sid })} D_{1}^{(\text {sid })}\left(x^{(\text {sid })}\right)\right) \cdots\right)\right]=1
$$

where $\widetilde{D}_{i}^{\text {(sid) }}$ is the resulting garbled database after executing sid- 1 homomorphic evaluations, $\left(\mathrm{ct}_{0}^{(\text {sid })}, \mathrm{x} \_\right.$secret $\left.^{(\text {sid })}\right) \leftarrow \operatorname{InpEnc}\left(x^{(\text {sid })}\right), \mathrm{ct}_{i}^{(\text {sid })} \leftarrow \operatorname{Eval}\left(\mathrm{crs}, i,\left\{\mathrm{pk}_{j}\right\}_{j=i+1}^{n}, \mathrm{ct}_{i-1}^{(\text {sid })}, \mathrm{sk}, P_{i}^{(\text {sid })}, t_{i}^{(\text {sid })}\right.$, digest $\left._{i}^{(\text {sid })}\right)$.

Server Privacy (Semantic Security). For server privacy, we require that for every pair of inputs $\left(x_{0}, x_{1}\right)$, let $\left(\mathrm{CT}_{b}\right.$, x_secret $\left.{ }^{b}\right) \leftarrow \operatorname{InpEnc}\left(x_{b}\right)$ for $b \in\{0,1\}$, then

$$
\mathrm{CT}_{0} \stackrel{\mathrm{c}}{\approx} \mathrm{CT}_{1} .
$$

Client Privacy (Program Privacy) with Unprotected Memory Access (UMA). We define client privacy against a semi-honest adversary that corrupts the server $S$ as well as an arbitrary subset of clients $\mathcal{I} \subset[n]$. Intuitively, we require program-privacy for the honest clients such that the adversary cannot learn anything beyond the output of the honest client's program on one input. We formalize this as follows:

There exists a PPT simulator ihopSim such that the following holds. Let crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}\right)$, for every set of keys $\left\{\left(\mathrm{pk}_{i}, \mathrm{sk}_{i}\right)\right\}_{i=1}^{n}$ in the support of KeyGen, and any collection of initial databases $D_{1}, \cdots, D_{n}$, for any unbounded polynomial N number of executions: For sid $\in \mathrm{N}$, let $P_{1}^{(\text {sid })}, \ldots, P_{n}^{\text {(sid) }}$ be the sequence of programs, $x^{(\text {sid })}$ be the server's input, then

$$
\begin{aligned}
& \left(\operatorname{crs},\left(\widetilde{D}_{1}, \ldots, \widetilde{D}_{n}\right),\left\{\mathrm{ct}_{0}^{(\text {sid })}, \mathrm{ct}_{1}^{(\text {sid })}, \ldots, \mathrm{ct}_{n}^{(\text {sid })}\right\}_{\text {sid } \in[\mathrm{N}]}\right) \stackrel{\mathrm{c}}{\approx} \\
& \left(\mathrm{crs}, \operatorname{ihopSim}\left(\mathrm{crs},\left\{\left\{\mathrm{pk}_{i}, \mathrm{sk}_{i}\right\}_{i \in[n]},\left(\left\{D_{j}, P_{j}^{(\text {sid })}\right\}_{j \in \mathcal{I}}, x^{(\text {sid })}\right),\left\{D_{j}, \text { MemAccess }_{j}^{(\text {sid })}, y_{j}^{(\text {sid })}\right\}_{j \in[n] \backslash \mathcal{I}}\right\}_{\text {sid } \in[\mathbb{N}]}\right)\right)
\end{aligned}
$$

where $\widetilde{D}_{i}, \mathrm{ct}_{i}^{(\text {sid })}$ corresponds to outputs in the real execution given all the programs and databases and $y_{j}^{(\text {sid })}=P_{j}^{(\text {sid })} D_{j}^{(\text {sid })}\left(\cdots\left(P_{1}^{(\text {sid }) D_{1}^{(\text {sid })}}\left(x^{(\text {sid })}\right)\right) \cdots\right)$.

Remark 7.1. We note that the above definition also captures the security against a semi-malicious adversary who may choose his randomness for KeyGen maliciously but behaves honestly in the protocol.

Client Privacy (Program Privacy) with Full Security. In the case of full client privacy, the simulator ihopSim does not get the database and the access pattern for the honest clients. That is, the simulator ihopSim takes as input $\left\{\left\{\mathrm{pk}_{i}, \mathrm{sk}_{i}\right\}_{i \in[n]},\left(\left\{D_{j}, P_{j}^{(\text {sid })}\right\}_{j \in \mathcal{I}}, x^{(\text {sid })}\right),\left\{1^{M_{j}}, 1^{1_{j}^{\text {(sid) }}}, y_{j}^{(\text {sid })}\right\}_{j \in[n] \backslash \mathcal{I}}\right\}_{\text {sid } \in[\mathrm{N}]}$, where $M_{j}$ is the size of $D_{j}$ and $t_{j}^{(\text {sid })}$ is the running time of $P_{j}^{(\text {sid })}$.

Efficiency. We require the following efficiency guarantees from mhop-RAM. Let $M_{i}=\left|D_{i}\right|$.

- $\left|\widetilde{D}_{i}\right|=M_{i} \cdot \operatorname{poly}\left(\lambda, \log M_{i}\right)$ for all $i \in[n]$.
- $\left|\mathrm{ct}_{0}\right|=|x| \cdot \operatorname{poly}(\lambda)$, where $\mathrm{ct}_{0}$ is the output of $\operatorname{InpEnc}(x)$.
- $\left|\mathrm{ct}_{i}\right|=\sum_{j=1}^{i} n \cdot t_{j} \cdot \operatorname{poly}\left(\lambda, \log M_{j}, \log t_{j}\right)$ for all $i \in[n]$.

Extension. This definition (and our construction) can be extended to the setting where in each execution all the clients do not necessarily join the homomorphic evaluation. We allow for different set of clients to participate in different executions. In particular, before the first execution, the initial database of every client is encrypted. Later before each execution, a sequence of distinct clients $\left\langle i_{1}, \cdots, i_{m}\right\rangle$ can be specified.

The input encryption is the same as above, while the homomorphic evaluation is executed in the specified order (as specified by the server) as $\mathrm{ct}_{j} \leftarrow \operatorname{Eval}\left(\mathrm{crs}, j,\left\{\mathrm{pk}_{i_{u}}\right\}_{u=j+1}^{n}, \mathrm{sk}_{i_{j}}, \mathrm{ct}_{j-1}, P_{i_{j}}, t_{i_{j}}\right.$, digest $_{i_{j}}$ )
 $\widetilde{D}_{i_{1}}, \cdots, \widetilde{D}_{i_{m}}$ are the up-to-date garbled databases of clients $Q_{i_{1}}, \ldots, Q_{i_{m}}$. The correctness and privacy properties can be naturally extended to this setting.

### 7.2 Building Blocks Needed

In this section we introduce building blocks needed in our construction. In addition to these building blocks, we will also need RAM computation model (see Section 6.1.1) and two-message oblivious transfer (see Section 6.1.2). We use $[n]$ to denote the set $\{1, \ldots, n\}$.

### 7.2.1 2-message Secure Function Evaluation based on Garbled Circuits

A two-message secure function evaluation (SFE) based on garbled circuits is as follows: Let $\mathcal{U}(\cdot, \cdot)$ be a particular "universal circuit evaluator" that takes as input the description of a circuit $C$ and an argument $x$, and returns $\mathcal{U}(C, x)$. We write $C(x)$ as a shorthand for $\mathcal{U}(C, x)$. Let Alice be the client with private input $x$ and Bob have private input a circuit $C$. The protocol is as follows:

1. $\left(m_{1}\right.$, x_secret $) \leftarrow \operatorname{SFE}_{1}(x)$ : Alice computes $\left(m_{1}, \mathrm{x} \_\right.$secret $) \leftarrow \mathrm{OT}_{1}(x)$ and sends $m_{1}$ to Bob.
2. $m_{2} \leftarrow \operatorname{SFE}_{2}\left(C, m_{1}\right)$ : Bob computes $\widetilde{\mathrm{C}} \leftarrow \operatorname{GCircuit}\left(\mathrm{C},\left\{\operatorname{key}_{b}^{w}\right\}_{w \in \operatorname{inp}(\mathrm{C}), b \in\{0,1\}}\right)$ and $\mathrm{L} \leftarrow \mathrm{OT}_{2}\left(m_{1},\left\{\operatorname{key}_{b}^{w}\right\}_{w \in \operatorname{inp}(\mathrm{C}), b \in\{0,1\}}\right)$. Sends $m_{2}:=(\widetilde{\mathrm{C}}, \mathrm{L})$.
3. $y=\operatorname{SFE}_{\text {out }}\left(\mathrm{x} \_\right.$secret, $\left.m_{2}\right)$ : Alice locally computes the output: $\left\{\operatorname{key}_{x_{w}}^{w}\right\}_{w \in \operatorname{inp}(\mathrm{C})}=\mathrm{OT}_{3}(\mathrm{~L}$, s_secret $)$, and $y=\operatorname{Eval}\left(\widetilde{\mathrm{C}},\left\{\operatorname{key}_{x_{w}}^{w}\right\}_{w \in \operatorname{inp}(\mathrm{C})}\right)$.
The correctness of the above protocol follows from the correctness of Yao garbled circuits. It can be shown that the above protocol is a secure function evaluation protocol satisfying both semi-honest client privacy and semi-honest server privacy.

### 7.2.2 Re-Randomizable Secure Function Evaluation based on Garbled Circuits

[GHV10] defined the tool of "re-randomizable secure function evaluation" that was used to realize multi-hop homomorphic computation for circuits. This tool was constructed under the DDH assumption by instantiating Yao's garbled circuits with a special encryption scheme (BHHO [BHHO08]) and using re-randomizable two-message oblivious transfer [NP01].

Definition 7.2. A secure function evaluation protocol is said to be re-randomizable if there exists an efficient procedure Re-rand such that for every input $x$ and function $f$ and every $\left(m_{1}, x_{\text {_secret }}\right) \in$ $\operatorname{SFE}_{1}(x)$ and $m_{2} \in \operatorname{SFE}_{2}\left(C, m_{1}\right)$, the distributions $\left\{x, C, m_{1}, x \_\right.$secret, $\left.m_{2}, \operatorname{Re}-r a n d\left(m_{1}, m_{2}\right)\right\}$ and $\left\{x, C, m_{1}, x_{-}\right.$secret, $\left.m_{2}, \mathrm{SFE}_{2}\left(C, m_{1}\right)\right\}$ are computationally indistinguishable.
[GHV10] proved the following:
Theorem 7.3 ( [GHV10]). Under the DDH assumption, there exists a re-randomizable secure function evaluation protocol satisfying Definition 7.2.

Below, we abstract out the scheme of [GHV10] by stating some of the procedures implicitly provided by [GHV10] that will be needed for this paper.

Definition 7.4 (Re-randomizable Yao garbled circuits.). The scheme in [GHV10] provides the following algorithms (implicitly) for their re-randomizable Yao scheme.

1. Keys $=$ SampleKeys $\left(1^{\lambda}, \mathrm{W} ; r\right)$ : Takes as input a set of input wires W as well randomness $r$ and outputs the input-keys for set of wires W for re-randomizable Yao. Note that it is a deterministic function given the randomness $r$. When clear from context, we will skip mentioning the inputs in the calls to this function.
2. $\widetilde{\mathrm{C}} \leftarrow \operatorname{ReGCircuit}(C$, InpKeys): Takes as input a circuit $C$ and InpKeys for the input wires of $C$ and outputs a re-randomizable garbled circuit $\widetilde{\mathrm{C}}$ where input wires have keys as InpKeys.
3. $\widetilde{\mathrm{C}}^{\prime} \leftarrow \operatorname{ReGCircuit}{ }^{\prime}(C$, InpKeys, OutKeys): Takes as input a circuit $C$, InpKeys for input wires of $C$ and OutKeys for output wires of $C$, and outputs a re-randomizable garbled circuit $\widetilde{\mathrm{C}}$ where input wires have keys as InpKeys and output wires have keys as OutKeys.
4. Keys $^{\dagger}=$ Transform $($ Keys, $r)$ : Takes as input Keys and randomness $r$ and outputs randomized keys Keys ${ }^{\dagger}$. Also, we use Transform(Keys, $\left\{r_{1}, \ldots, r_{k}\right\}$ ) to denote

$$
\operatorname{Transform}\left(\operatorname{Transform}\left(\ldots\left(\operatorname{Transform}\left(\text { Keys }, r_{1}\right), \ldots\right), r_{k-1}\right), r_{k}\right)
$$

5. $\left(\widetilde{C}^{\prime}, \mathrm{L}^{\prime}\right) \leftarrow \operatorname{Re}-\mathrm{rand}\left((\widetilde{\mathrm{C}}, \mathrm{L}),\left\{r_{w}\right\}_{w \in \mathrm{Wires}(C)}\right)$ : Takes as input a re-randomizable garbled circuit $\widetilde{\mathrm{C}}$ and OT second messages L for the keys of input wires of $C$ and randomness to re-randomize each wire of $C$ and outputs a new functionally equivalent re-randomizable garbled circuit $\widetilde{\mathrm{C}}^{\prime}$ and consistent OT second messages $\mathrm{L}^{\prime}$. This procedure satisfies the property of re-randomizable SFE. Moreover, the guarantee is that after randomization, for any wire $w$, the new keys for $w$ in $\widetilde{\mathrm{C}}^{\prime}$ are Transform $\left(\mathrm{key}^{w}, r_{w}\right)$. Finally, re-randomization of OT messages only requires ${ }^{11}$ $\left\{r_{w}\right\}_{w \in \operatorname{inp}(\mathrm{C})}$.

For our multi-hop homomorphic scheme for RAM it will be useful to define SampleKeys(•) as follows: SampleKeys $\left(1^{\lambda}, \mathrm{W}, r\right)=\operatorname{Transform}\left(\operatorname{SampleKeys}\left(1^{\lambda}, \mathrm{W}, 0^{\star}\right), r\right)$.

[^7]
### 7.3 Our Construction of Multi-hop RAM Scheme

Below, in Section 7.3 .1 we provide our construction for multi-hop scheme for RAM satisfying UMAsecurity for clients (see Section 7.1). We give a general transformation for UMA to full security in Section 7.3.5.

### 7.3.1 UMA Secure Construction

In this section, we first describe the UMA-secure scheme for one execution and then explain how this scheme can be extended naturally for multiple executions in Section 7.3.3. Also, as we shall see our scheme can easily be extended to the setting where different subset of parties participate in each session.

UMA-secure multi-hop RAM scheme for a single execution involving all parties. Let $Q_{1}, \ldots, Q_{n}$ be the clients holding databases $D_{1}, \ldots, D_{n}$, respectively, and $S$ be the server. Let the server's private input be $x$ and secret programs of clients be $P_{1}, \ldots, P_{n}$, respectively. Let $\ell O T=$ (crsGen, Commit, Send, Receive, SendWrite, ReceiveWrite) be an updatable laconic OT scheme with sender privacy as defined in Definition 3.2. Let Re-GC $=$ (SampleKeys, ReGCircuit, ReGCircuit', Transform, Re-rand) be a re-randomizable scheme for Yao's garbled circuits given by [GHV10] (see Definition 7.4). Let OT $=\left(\mathrm{OT}_{1}, \mathrm{OT}_{2}, \mathrm{OT}_{3}\right)$ be a two-message oblivious transfer protocol as defined in Section 6.1.2.

The multi-hop RAM scheme mhop-RAM = (Setup, KeyGen, EncData, InpEnc, Eval, Dec) is as follows: The algorithms Setup, KeyGen, EncData are formally described in Figure 14.

Setup: $\quad \operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}\right)$
Setup algorithm generates the common reference string for laconic OT.
Key Generation: $(\mathrm{pk}, \mathrm{sk}) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}\right)$
Each client runs this algorithm once to generate the secret-key sk and public-key pk. A client $Q$ picks a PRF seed $s$ as the secret key. Next, it generates the public key as the first message of OT for $s$ and secret-key as the secret state for OT as well as PRF key.

Looking ahead, the client $Q$ will use the PRF key $s$ to garble his own $P$ and to randomize the garbled program generated by all previous clients in any execution.

Database Encryption: $\quad \widetilde{D} \leftarrow$ EncData (crs, $D$ )
Each client runs this algorithm at the beginning to garble the database and sends the garbled database to the server $S$. The garbled database is generated by executing the Hash procedure of laconic OT. This outputs an encoded database $\hat{D}$ and a digest digest, both of which are given to the server $S$.

Input Encryption: $(\mathrm{ct}, \mathrm{x}$ _secret $) \leftarrow \operatorname{InpEnc}(x)$
In each execution, the server $S$ encrypts its input $x$ as follows: It computes the first message of OT as the ciphertext and stores the secret state of OT to be used for decryption of computation later. The ciphertext is sent the first client, w.l.o.g. $Q_{1}$. The algorithm InpEnc is described formally in Figure 15.

Set up. $\quad$ crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}\right)$.

1. Sample crs $\leftarrow \operatorname{crsGen}\left(1^{\lambda}\right)$.
2. Output crs.

Key Generation. $\quad(\mathrm{pk}, \mathrm{sk}) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}\right)$.

1. Sample a PRF key $s \leftarrow\{0,1\}^{\lambda}$ and generate (pk, s_secret) $\leftarrow \mathrm{OT}_{1}(s)$.
2. Output ( $\mathrm{pk}, \mathrm{sk}:=(s$, s_secret $)$ ).

Database Encryption. $\quad \widetilde{D} \leftarrow$ EncData(crs, $D$ ).

1. (digest, $\hat{\mathrm{D}}) \leftarrow \mathrm{Hash}(\mathrm{crs}, D)$.
2. Output $\widetilde{D}=($ digest,$\hat{\mathrm{D}})$.

Figure 14: Formal description of the set up, key generation and database encryption algorithms.
Input Encryption. $\quad\left(\mathrm{ct}, \mathrm{x} \_\right.$secret $) \leftarrow \operatorname{InpEnc}(\mathrm{crs}, x)$.

1. $(\mathrm{ct}, \mathrm{x}$ _secret $) \leftarrow \mathrm{OT}_{1}(x)$.

Figure 15: Input encryption algorithm.

Homomorphic Evaluation: $\quad \mathrm{ct}^{\prime} \leftarrow \operatorname{Eval}\left(\mathrm{crs}, i,\left\{\mathrm{pk}_{j}\right\}_{j=i+1}^{n}, \mathrm{ct}, \mathrm{sk}=(s\right.$, s_secret $),(P, t)$, digest $)$.
This algorithm is executed by client $Q_{i}$ to generate the next ciphertext ct' given ct from client $Q_{i-1}$, and is described formally in Figure 16. This is the most involved procedure in our construction, and hence, we first provide an informal description. At a very high level, as illustrated in Figure 17, the client $Q_{i}$ generate the garbled program for $P$ consisting of $t$ garbled circuits and also re-randomize all the circuits in ct. As mentioned before, this re-randomization step is crucial to get program privacy for this client. Moreover, the re-randomization has to be done carefully so that the previous ct is consistent with the new garbled program. ${ }^{12}$

This procedure consists of four main steps: Let $T$ be the number of step-circuits in ct.

1. Garble the new program $P$ : For each $\tau \in[T+1, T+t]$, client does the following: It generates a "super-circuit" that is illustrated in Figure 18 consisting of a CPU step circuit $C_{\tau}^{\text {step }}$ (see Figure 19) and PRF circuits $C_{\tau, i+1}^{P R F}, \ldots, C_{\tau, n}^{P R F}$ (see Figure 20). A step circuit, encodes the logic of a CPU step of a program $P$ and PRF circuits provide a part of the randomness used in re-randomization. We will elaborate on the functionality of PRF circuits later.

The garbled program will consists of garbled circuits corresponding to all the step circuits and PRF circuits. The first step is to pick the keys for the input wires of all of these circuits. Next, we begin by describing the step circuits.

Step Circuits $C_{\tau}^{\text {step }}$ (Figure 19) : The inputs of a step circuit (see Figure 18) can be partitioned into ((state, rData, digest), Rd), where state is the current CPU state, rData is the bit-read from database, and digest is the up-to-date digest of the database. Rd corresponds to the randomness given as input to the step-circuit computed from the PRF circuits. A step

[^8]
## Homomorphic Evaluation.

$\mathrm{ct}^{\prime} \leftarrow \operatorname{Eval}\left(\mathrm{crs}, i,\left\{\mathrm{pk}_{j}\right\}_{j=i+1}^{n}, \mathrm{ct}=\left(\mathrm{L}_{0},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\mathrm{PRF}}, \mathrm{L}_{\tau, j}\right\}_{j=i}^{n}\right\}_{\tau \in[T]}\right), \mathrm{sk}=(s\right.$, s_secret $),(P, t)$, digest $)$.

1. Generate the "new" garbled program for $P$ : Generate garbled circuits $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\text {PRF }}\right\}_{j=i+1}^{n}\right\}_{\tau=T+1}^{T+t}$.
(a) Set stateKeys ${ }^{\tau}$, dataKeys $^{\tau}$, digestKeys $^{\tau}$, RdKeys $^{\tau, j}$, PKeys $^{\tau, j}$ for each $\tau \in[T+1, T+t]$ and $j \in[i+1, n]$ as $\operatorname{Sample} \operatorname{Keys}\left(F_{s}\left(\mathrm{GC}_{-} \star \| \tau\right)\right)$ where $F_{s}(\mathrm{GC}-\star \| \tau)$ is the randomness used and $\star \in\{$ STATE, DATA, DIGEST, RD, P\}, respectively.
Set stateKeys ${ }^{\tau}$, dataKeys $^{\tau}$, digestKeys $^{\tau}$ to SampleKeys $\left(0^{*}\right)$ for $\tau=T+t+1$.
(b) Garble $\mathrm{C}^{\text {step }}$ circuits: For each $\tau \in[T+1, T+t]$

$$
\widetilde{\mathrm{C}}_{\tau}^{\text {step }} \leftarrow \operatorname{ReGCircuit}\left(\mathrm{C}^{\text {step }}\left[i, \operatorname{crs}, P, \text { Keys }^{\tau+1}, F_{s}(\mathrm{PSI} \| \tau)\right],\left(\text { Keys }^{\tau},\left\{\text { RdKeys }^{\tau, j}\right\}_{j=i+1}^{n}\right)\right),
$$

where Keys $^{\tau}=\left(\right.$ stateKeys $^{\tau}$, dataKeys $^{\tau}$, digestKeys $\left.^{\tau}\right)$.
Embed labels dataKeys ${ }_{0}^{T+1}$ and digestKeys digest $_{T+1}$ in $\widetilde{\mathrm{C}}_{T+1}^{\text {step }}$.
(c) Garble $C^{\text {PRF }}$ circuits: For each $[T+1, T+t]$ and $j \in[i+1, n]$, compute

$$
\widetilde{\mathrm{C}}_{\tau, j}^{\mathrm{PRF}} \leftarrow \operatorname{ReGCircuit}^{\prime}\left(\mathrm{C}^{\text {PRF }}[\tau], \mathrm{PKeys}^{\tau, j}, \operatorname{RdKeys}^{\tau, j}\right) .
$$

2. Generate the OT second messages for newly generated circuits: For all $\tau \in[T+$ $1, T+t]$ and $j \in[i+1, n]$ compute $\mathrm{L}_{\tau, j} \leftarrow \mathrm{OT}_{2}\left(\mathrm{pk}_{j}, \mathrm{PKeys}^{\tau, j}\right)$.
3. Obtain partial labels for previous circuits:
(a) For every $\tau \in[T]$, compute $\mathrm{M}_{\tau, i}=\mathrm{OT}_{3}\left(\mathrm{~L}_{\tau, i}\right.$, s_secret) and $\widetilde{\mathrm{C}}_{\tau, i}^{\text {PRF }}$ using input labels $\mathrm{M}_{\tau, i}$ and embed the labels in $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$.
(b) If $i=1$, then $\mathrm{L}_{0} \leftarrow \mathrm{OT}_{2}\left(\mathrm{~L}_{0}\right.$, stateKeys $\left.{ }^{1}\right)$.
4. Re-randomize previous garbled circuits: If $i>1$, do the following:
(a) For each $\tau \in[T]$, re-randomize the circuit $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ using Re-rand $(\cdot)$ (see Definition 7.4) such that the input wire keys are randomized using $F_{s}\left(\mathrm{GC}_{-} \star \| \tau\right)$, where $\star \in$ \{STATE, DATA, DIGEST, RD\} for different input wires appropriately.
(b) For each $\tau \in[T]$, re-randomize the circuits $\left\{\mathrm{C}_{\tau, j}^{\mathrm{PRF}}\right\}_{j \in[i+1, n]}$ and $\left\{\mathrm{L}_{\tau, j}\right\}_{j \in[i+1, n]}$ using Re-rand $(\cdot)$ such that the input wires are randomized using $F_{s}(\mathrm{GC}-\mathrm{P} \| \tau)$ and output wires are randomized using $F_{s}\left(\mathrm{GC} \_\mathrm{RD} \| \tau\right)$.
(c) Re-randomize $\mathrm{L}_{0}$ using $F_{s}($ GC_STATE $| | 1)$.
5. Output $\mathrm{ct}^{\prime}=\left(\mathrm{L}_{0},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\text {PRF }}, \mathrm{L}_{\tau, j}\right\}_{j=i+1}^{n}\right\}_{\tau=1}^{T+t}\right)$.

Figure 16: Homomorphic evaluation algorithm


Figure 17: Homomorphic Evaluation by $Q_{i}: Q_{i}$ contributes new circuits (denoted in white in the lower layer) and processes the input circuits as follows: (i) computes the yellow circuits, and (ii) re-randomizes all input circuits. The re-randomized circuits are shown in gray color.
circuit executes one CPU step and passes on the updated state, new bit read, and new digest to the next step circuit. Note that we do not achieve this by passing the output wires of $\tau$ into input wires of $\tau+1$. That is, the output wire of $\tau$ will not have same keys as input wires of $\tau+1$ (Note that the two consecutive step circuits are not connected by solid lines in Figure 17.). Hence, the step circuit $\tau$ will have the keys of the next circuit hard-coded inside it.

Next, we explain the logic of a step-circuit. First, it computes the new (state', R/W, $L$, wData). Next, it computes the transformed keys nextKeys ${ }^{\dagger}$ of the next step-circuit using the hard-coded keys and the input randomness (this uses the transform functionality of re-randomizable Yao from Section 7.2.2). Then, in the case of a "read" it outputs stateKeys ${ }^{\dagger}$ corresponding to new state', labels for data via laconic OT procedure Send (•) for location $L$ where the sender's inputs are dataKeys ${ }_{0}^{\dagger}$, dataKeys ${ }_{1}^{\dagger}$ and digestKeys ${ }^{\dagger}$ corresponding to digest. The case of a write is similar, but now the labels of new updated digest are transferred via laconic OT procedure SendWrite(•). Note that it follows via correctness of reads and writes of the laconic OT that the evaluator would be able to recover the correct labels for the read-data and the new digest.

The down-bend in output and input wires of step-circuits for data and digest in Figure 17 represents that these keys are not output in the clear, but are output using laconic OT.


Figure 18: One step circuit along with the attached PRF circuits.

Correct labels will be learnt during execution using the encoded database $\widetilde{D}_{i}$ and laconic OT procedures.

PRF circuits $C_{\tau, j}^{\text {PRF }}$ (Figure 20) : This circuit takes as input a PRF key $s_{j}$ of client $Q_{j}$ and outputs the PRF value corresponding to time-step $\tau$. The use of these circuits will be clear when we describe the re-randomization step below.
All these circuits are garbled such that the keys for output wires of PRF circuits are same as keys for Rd input keys of step circuits. In Figure 17, this is depicted by joining the output wires of PRF circuits with Rd input wires of step circuit with a solid line. The garbled program consists of garbled step circuits and garbled PRF circuits $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\text {PRF }}\right\}_{j \in\{i+1, \ldots, n\}}\right\}_{\tau \in[T+1, T+t]}$. The client also embeds labels for $\mathrm{rData}=0$ and digest $_{i}$ in the first step circuit.
2. Generate OT messages for $\mathrm{C}_{\tau, j}^{\mathrm{PRF}}$ : Recall that this circuit takes as input a PRF key $s_{j}$ of client $Q_{j}$ whose OT first message is present in $\mathrm{pk}_{j}$. Client $Q_{i}$ generates the OT second message $\mathrm{L}_{\tau, j}$ for the input keys of $\widetilde{\mathrm{C}}_{\tau, j}^{\mathrm{PRF}}$.
3. Evaluating the PRF circuits for itself: Note that the ciphertext ct consists of a sequence of step circuits and PRF circuits for each step circuit corresponding to $j \in\{i, \ldots, n\}$. See Figure 17 where the PRF circuits for client $Q_{i}$ are depicted in yellow. $Q_{i}$ computes the input labels for $\widetilde{\mathrm{C}}_{\tau, i}^{\text {PRF }}$ using the OT message $\mathrm{L}_{\tau, i}$ and embeds the output labels in to $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ for

```
Hard-coded parameters: \([i\), crs, \(P\), nextKeys \(=(\) stateKeys, dataKeys, digestKeys),\(\psi]\).
Input: ((state, rData, digest), \(\left.\left(\left\{\omega_{j}, \phi_{j}\right\}_{j>i}\right)\right)\).
    (state \(\left.{ }^{\prime}, \mathrm{R} / \mathrm{W}, L, \mathrm{wData}\right):=C_{\mathrm{CPU}}^{P}(\) state, rData) .
    nextKeys \({ }^{\dagger}:=\) Transform(nextKeys, \(\left\{\omega_{j}\right\}_{j>i}\) ).
    Parse nextKeys \({ }^{\dagger}\) as (stateKeys \({ }^{\dagger}\), dataKeys \({ }^{\dagger}\), digestKeys \({ }^{\dagger}\) ).
    if \(R / W=\) read then
    \(\mathrm{e}_{\text {data }} \leftarrow\) Send (crs, digest, \(L\), dataKeys \(\left.{ }^{\dagger} ; \psi \oplus \bigoplus_{j>i} \phi_{j}\right)\).
    return \(\left(\left(\right.\right.\) stateKeys \(_{\text {state }^{\prime}}^{\dagger}, \quad e_{\text {data }}, \quad\) digestKeys \(\left.\left.{ }_{\text {digest }}^{\dagger}\right), \mathrm{R} / \mathrm{W}, L\right)\).
else
    \(\mathrm{e}_{\text {digest }} \leftarrow\) SendWrite \(\left(\right.\) crs, digest, \(L\), wData, digestKeys \(\left.{ }^{\dagger} ; \psi \oplus \bigoplus_{j>i} \phi_{j}\right)\).
    return \(\left(\left(\right.\right.\) stateKeys \(_{\text {state }}^{\dagger}, \quad\) dataKeys \(\left.\left.{ }_{0}^{\dagger}, \quad e_{\text {digest }}, w D a t a \quad\right), R / W, L\right)\).
```

Figure 19: Pseudocode of the step circuit $\mathrm{C}^{\text {step }}[i$, crs, $P$, nextKeys, $\psi]$.

```
Hard-coded parameters: [\tau].
Input: s.
Output: }({\mp@subsup{F}{s}{}(\mp@subsup{\textrm{GC}}{-}{*}||\tau+1)\mp@subsup{}}{*\in{\mathrm{ STATE,DATA,DIGEST }}}{},\mp@subsup{F}{s}{(LACONIC_OT|}|\tau))
```

Figure 20: PRF circuit $\mathrm{C}^{\mathrm{PRF}}[\tau]$.
all $\tau \in[T]$. In other words, $Q_{i}$ consumes the first PRF circuits from each step of previous ciphertext ct.
4. Re-randomize the previous circuits: After consuming the first PRF circuit from each step, $Q_{i}$ randomizes all the remaining circuits using appropriate randomness. Note that the input keys of $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}$ are randomized using the exact randomness that was fed into $\mathrm{C}_{\tau}^{\text {step }}$ via the PRF circuit for $Q_{i}$. This makes sure that the hard-coded input keys of step $\tau+1$ are randomized consistently in the same way as how $Q_{i}$ will randomize the circuit $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}$.
Hence, to conclude, the PRF circuits are present to provide the randomness needed to randomize the hard-coded keys inside the step circuits ${ }^{13}$.

Homomorphic Decryption: $y=\operatorname{Dec}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}\left(\right.$ crs, x_secret, $\left.\mathrm{ct}=\left(\mathrm{L}_{0},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau \in[T]}\right)\right)$.
The algorithm is described formally in Figure 21. It takes as input the secret state of the server x _secret and the final ciphertext $\mathrm{ct}_{n}$ consisting of OT message for $x$ and sequence of $T$ step-circuits, where $T=\sum_{i \in[n]} t_{i}$. Note that all the PRF circuits have been evaluated already by correct parties and correct labels for RdKeys have been embedded into the step-circuits. The server does the following for decryption:

1. It obtains the stateLabels for the first circuit by running $\mathrm{OT}_{3}$. Note that the first circuit of program of any client has labels for data and digest already embedded. Hence, now the server knows all the labels for the first circuit.
[^9]
## Homomorphic Decryption.

$y=\operatorname{Dec}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}\left(\mathrm{crs}, \mathrm{x} \_\right.$secret, $\left.\mathrm{ct}=\left(\mathrm{L}_{0},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right\}_{\tau \in[T]}\right)\right)$, where $T=\sum_{i \in[n]} t_{i}$.

1. For all $i \in[n]$, parse $\widetilde{D}_{i}=\left(\operatorname{digest}_{i}, \hat{\mathrm{D}}_{i}\right)$.
2. Compute $\mathrm{M}_{0}=\mathrm{OT}_{3}$ ( $\mathrm{L}_{0}$, x_secret).
3. Parse $\widetilde{\mathrm{C}}_{1}^{\text {step }}=\left(\widetilde{\mathrm{C}}_{1}^{\text {step }}\right.$, dataLabels, digestLabels).
4. Define Labels ${ }^{1}=\left(\right.$ stateLabels $^{1}=\mathrm{M}_{0}$, dataLabels $^{1}=$ dataLabels, digestLabels $^{1}=$ digestLabels $^{\text {}}$.
5. For $\tau=1$ to $T$ do the following:

> Define $i$ s.t. $\tau \in\left[\sum_{j \in[i-1]} t_{j}+1, \sum_{j \in[i]} t_{j}\right]$.
> $(X, \mathrm{R} / \mathrm{W}, L):=\operatorname{ReEval}\left(\widetilde{\mathrm{C}}_{\tau}^{\text {step }}\right.$, Labels $\left.^{\tau}\right)$.
if $\mathrm{R} / \mathrm{W}=$ read then
Parse $X=\left(\right.$ stateLabels $^{\tau+1}, \mathrm{e}_{\text {data }}$, digestLabels $\left.{ }^{\tau+1}\right)$.
dataLabels ${ }^{\tau+1}:=\operatorname{Receive}^{\hat{\mathrm{D}}_{i}}\left(\mathrm{crs}, \mathrm{e}_{\text {data }}, L\right)$.
else
Parse $X=\left(\right.$ stateLabels $^{\tau+1}$, dataLabels $^{\tau+1}, \mathrm{e}_{\text {digest }}$, wData $)$.
digestLabels $^{\tau+1}:=\operatorname{ReceiveWrite}^{\hat{\mathrm{D}}_{i}}\left(\right.$ crs,$L$, wData, $\left.\mathrm{e}_{\text {digest }}\right)$.
if $\tau=\sum_{j \in[i]} t_{j}$ and $\tau<T$ then
Parse $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}=\left(\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}\right.$, dataLabels, digestLabels $)$
Set dataLabels ${ }^{\tau+1}=$ dataLabels and digestLabels ${ }^{\tau+1}=$ digestLabels.

$$
\text { Labels }^{\tau+1}:=\left(\text { stateLabels }^{\tau+1}, \text { dataLabels }^{\tau+1}, \text { digestLabels }^{\tau+1}\right)
$$

6. Decode output $y$ using (stateLabels ${ }^{T+1}$, SampleKeys $\left(0^{*}\right)$ ).

Figure 21: Decryption algorithm for multi-hop scheme for RAM.
2. For $\tau \in[T]$, the server executes the circuit $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$, and learns the labels for the next circuit via running the receiver algorithms of laconic OT correctly.

### 7.3.2 Correctness

Here we first prove correctness (as defined in Section 7.1) for a single execution. In fact, we would prove something stronger that would help us extend the scheme to multiple executions in a straightforward manner. We prove the following two properties:

Property 1. For the above scheme, $y=P_{n}^{D_{n}}\left(\ldots P_{1}^{D_{1}}(x) \ldots\right)$, where programs, databases and input $x$ are as defined above.

Property 2. Let $\hat{D}_{i}^{\prime}$, digest $_{i}^{\prime}$ denote the updated encoded database and digest with the server after
the execution. We show that these are equal to Hash (crs, $D_{i}^{\prime}$ ), where $D^{\prime}$ results after executing $P_{i}^{D_{i}}\left(P_{i-1}^{D_{i-1}}\left(\ldots P_{1}^{D_{1}}(x) \ldots\right)\right)$.

Below we prove correctness via a sequence of facts and claims.
Fact 7.5. At any point in homomorphic evaluation, the circuit $\widetilde{\mathcal{C}}_{\tau, j}^{\text {PRF }}$ and the second OT message for its input keys $\mathrm{L}_{\tau, j}$ are consistent.

This follows from correctness of $\mathrm{OT}_{2}(\cdot, \cdot)$ procedure when it is generated and the fact that re-randomization happens consistently in Re-rand procedure of re-randomizable garbled circuits.

Fact 7.6. During the homomorphic evaluation of client $Q_{i}$, in Step 3a, Figure 16 while obtaining partial labels, $\mathrm{M}_{\tau, i}=\mathrm{PKeys}_{s}^{\tau, i}$, where $s$ is the PRF key of $Q_{i}$.

This follows from the correctness of OT protocol as well as Fact 7.5.
Fact 7.7. During the homomorphic evaluation of client $Q_{i}$, in Step 3a, Figure 16 the labels embedded in circuit $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ correspond to $\operatorname{RdKeys}{ }_{\omega_{i}, \phi_{i}}^{\tau, i}$ where $\phi_{i}=F_{s}\left(\operatorname{LACONIC\_ OT} \| \tau\right)$ and $\omega_{i}=$ $\left\{F_{s}\left(\mathrm{GC}_{-} \star \| \tau+1\right)\right\}_{\star \in\{\text { STATE,DATA,DIGEST }\}}$.

This is because the functionality of the $C^{\text {PRF }}$ is preserved in randomization so far, Fact 7.6 and because the output keys of $\widetilde{\mathrm{C}}_{\tau, i}^{\text {PRF }}$ and RdKeys ${ }^{\tau, i}$ are same when they are generated and are re-randomized using same randomness in Step 4b of Figure 16.

Recall that $\mathrm{ct}_{n}$ consists of garbled step-circuits of client $Q_{1}$ followed by $Q_{2}$ and so on. We prove the following fact about garbled step circuits belonging to some client $Q_{i}$ in final ciphertext $\mathrm{ct}_{n}$.
Claim 7.8. Consider circuits $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ and $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}$ such that both belong to program $P_{i}$ for some $i$. Since all the PRF circuits $\widetilde{\mathrm{C}}^{\text {PRF }}$ have been evaluated, the value nextKeys ${ }^{\dagger}$ in $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ is well defined. Then, nextKeys ${ }^{\dagger}=$ Keys $^{\tau+1}$ where Keys $^{\tau+1}$ corresponds to the input keys for $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}$ in ct $_{n}$.
Proof. Initially, $Q_{i}$ picks $\operatorname{Keys}_{\sim}^{\tau+1}$ as SampleKeys $\left(F_{s}\left(\mathrm{GC}_{-} \star \| \tau+1\right)\right)$, where $\star \in\{$ STATE, DATA, DIGEST $\}$ and uses them in garbling of $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}$ as well as are hardcoding inside $\mathrm{C}_{\tau}^{\text {step }}$.

Then, $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}$ is randomized by clients $Q_{i+1}, \ldots, Q_{n}$ such that the stateKeys, dataKeys, digestKeys are randomized sequentially using $\omega_{j}=\left(F_{s_{j}}\left(\mathrm{GC} \_\right.\right.$STATE $\left.\| \tau+1\right), F_{s_{j}}\left(\mathrm{GC} \_\right.$DATA $\left.\| \tau+1\right), F_{s_{j}}\left(\mathrm{GC} \_\right.$DIGEST $\| \tau+$ 1)). This is same as $\operatorname{Transform}\left(\operatorname{Keys}^{\tau+1},\left\{\omega_{j}\right\}_{j>i}\right)$ inside $\widetilde{\mathbb{C}}_{\tau}^{\text {step }}$. By Fact $7.7, \omega_{j}$ is the value used for Transform in $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$.
Claim 7.9. The above claim also holds for $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ and $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}$ when $\tau$ is the last circuit of a program for $Q_{i}$ and $\tau+1$ is the first circuit for $Q_{i+1}$.

Proof. When the client $Q_{i}$ generates $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$, the keys hard-coded are SampleKeys $\left(0^{\star}\right)$. Then, this circuit is re-randomized by $Q_{i+1}$ resulting in keys SampleKeys $\left(F_{s_{i+1}}\left(\mathrm{GC}_{-} \star \| \tau+1\right)\right)$ which same as the value used by $Q_{i+1}$ to generate the step-circuit $\widetilde{\mathrm{C}}_{\tau+1}^{\text {step }}$.
Fact 7.10. The first garbled step circuit $\widetilde{\mathrm{C}}_{1}^{\text {step }}$ gets evaluated on ( $x, 0$, digest $_{1}$ ).
This follows from correctness of OT and consistency of re-rerandomization of OT and garbled circuits similar to Fact 7.5.

Now, we will prove a lemma about the execution of circuits generated by client $Q_{1}$. Then, we will prove a claim about the inputs on which circuit of $Q_{2}$ is executed. Finally, the correctness of execution programs of all clients would follow in a similar manner.

Lemma 7.11. Consider the program $P_{1}$ and the database $D_{1}$ of the first client and the input $x$ of the server. Consider the execution $P_{1}^{D_{1}}(x)$ execution in the clear as $\left(\operatorname{state}_{\tau}, \mathrm{rData}_{\tau}\right)$ as the values on which $\mathrm{C}^{\text {step }}$ is executed. Also, let $\left(\hat{\mathrm{D}}_{1}^{\tau}\right.$, digest ${ }_{\tau}$ ) denote the Hash $\left(\mathrm{crs}, D_{1}^{\tau}\right)$, where $D_{1}^{\tau}$ is the database at beginning of step $\tau$. Then, while decryption, $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ is executed on inputs ( state $_{\tau}, \mathrm{rData}_{\tau}$, digest $_{\tau}$ ). Moreover, the encoded database held by the server before step $\tau$ is $\hat{\mathrm{D}}_{1}^{\tau}$.

Proof. We will prove this lemma by induction on $\tau$. The base case follows from Fact 7.10. Assume that the lemma holds for $\tau=\rho$, then we prove that the lemma holds for $\rho+1$ as follows: So it holds that $\widetilde{\mathrm{C}}_{\rho}^{\text {step }}$ is executed on ( state $_{\rho}$, rData $_{\rho}$, digest $_{\rho}$ ). Moreover, ( $\hat{\mathrm{D}}_{1}^{\rho}$, digest $_{\rho}$ ) denote the Hash (crs, $\left.D_{1}^{\rho}\right)$. Note that $\widetilde{\mathrm{C}}_{\rho}^{\text {step }}$ correctly implements its code that includes one CPU step of $P_{1}$. Hence, $\left(\right.$ state $\left.^{\prime}, \mathrm{R} / \mathrm{W}, L, \mathrm{wData}\right)=C_{\mathrm{CPU}}^{P}\left(\right.$ state $_{\rho}$, rData $\left._{\rho}\right)$. Also, by Claim 7.8, nextKeys ${ }^{\dagger}$ in $\widetilde{\mathrm{C}}_{\rho}^{\text {step }}$ are correct input keys for $\widetilde{\mathrm{C}}_{\rho+1}^{\text {step }}$. There are following two cases:

- $R / W=$ read: In this case, database and the digest are unchanged. New CPU state and digest are output correctly. Moreover, the labels for bit read from the memory will be learnt via Receive of updatable laconic OT. Correctness of these labels follows from correctness of read of laconic OT.
- R/W = write: Similar to above, in this case new state and data keys are correctly output. Moreover, the digest keys w.r.t. the new updated digest are output via laconic OT. The correctness of these labels follows from correctness of laconic OT write function ReceiveWrite. Finally, in this function, the encoded database is updated correctly.

Lemma 7.12. Let $\widetilde{\mathrm{C}}_{t_{1}}^{\text {step }}$ be the last circuit of client $Q_{1}$ or program $P_{1}$. Then, during decryption, $\widetilde{\mathrm{C}}_{t_{1}+1}^{\text {step }}$ is executed on $\left(y_{1}, 0\right.$, digest $\left.{ }_{2}\right)$, where $y_{1}=P_{1}^{D_{1}}(x)$ and digest ${ }_{2}$ is the digest for $D_{2}$.

Proof. At the time of homomorphic evaluation, in Step 1b, Figure 16, labels dataKeys ${ }_{0}^{t_{1}+1}$ and digestKeys ${ }_{\text {digest }}^{t_{1}+1}$ are embedded in $\widetilde{\mathrm{C}}_{t_{1}+1}^{\text {step }}$. Also, by Claim 7.9, in the final ciphertext ct, nextKeys ${ }^{\dagger}$ inside $\widetilde{\mathrm{C}}_{t_{1}}^{\text {step }}$ are correct keys for $\widetilde{\mathrm{C}}_{t_{1}+1}^{\text {step }}$. Hence, the lemma holds since $\widetilde{\mathrm{C}}_{t_{1}}^{\text {step }}$ outputs stateKeys ${ }_{\text {state }}{ }^{\dagger}$, where state ${ }^{\prime}=y_{1}$.

Lemma 7.13. Consider the program $P_{i}$ and the database $D_{i}$ of the client $Q_{i}$ and the input $x$ of the server. Consider the execution $P_{i}^{D_{i}}\left(y_{i-1}\right)$ execution in the clear as $\left(\right.$ state $\left._{\tau}, \mathrm{rData}_{\tau}\right)$ as the values on which $\mathrm{C}^{\text {step }}$ is executed. Also, let ( $\hat{\mathrm{D}}_{i}^{\tau}$, digest ${ }_{\tau}$ ) denote the Hash $\left(\mathrm{crs}, D_{i}^{\tau}\right)$, where $D_{i}^{\tau}$ is the database at beginning of step $\tau$. Then, while decryption, $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$ is executed on inputs ( state $_{\tau}, \mathrm{rData}_{\tau}$, digest $_{\tau}$ ). Moreover, the encoded database held by the server before step $\tau$ is $\hat{\mathrm{D}}_{i}^{\tau}$.

Proof. The lemma follows via induction on number of clients where the base case is proved in Lemma 7.11. The rest of the proof follows simply via induction similar to Lemma 7.11 where at the end of each program and beginning of a new program we prove Lemma 7.12. This proves both the properties 1 and 2 .

### 7.3.3 Extending to Multiple Executions

Recall that for correctness we also proved that the after one execution, the resulting garbled database $\widetilde{D}=\left(\hat{\mathrm{D}}\right.$, digest) corresponds to the output of Hash(crs, $\left.D^{\prime}\right)$, where $D^{\prime}$ is the correct database resulting after the execution in the clear (See Property 2, Section 7.3.2).

Given this invariant after the first execution, the next execution happens identically as the first execution with minor differences. To run the algorithm Eval, the clients need the updated digest of their respective databases. The updated digests of all the clients taking part in an execution would be sent by the server to the first client on that execution path, and would be passed along with each ciphertext. Also, to ensure that no PRF output is used twice, each PRF invocation would take the session identifier sid as an additional input. With these changes, the second execution is identical to the first execution and hence, its correctness follows in a straight-forward manner.

Also, this does not affect the UMA-security because the simulator of the ideal world is given the databases as well as memory access pattern of the honest clients as input as well.

Moreover, note that this generalizes to the scenario when different subset of clients take part in different executions. Only the digests of the relevant client are passed around by the server. ${ }^{14}$

### 7.3.4 Security Proof

Server privacy follows receiver privacy of oblivious transfer. For ease of exposition, we prove client UMA privacy for the setting of a single honest client $Q_{i}$ for a single execution. At the end of this section we will show that the proof can be extended for the case of multiple honest clients and multiple executions as well.

In the following, we prove that there exists a PPT simulator ihopSim such that, for any set of databases $\left\{D_{j}\right\}_{j \in[n]}$, any sequence of compatible programs $P_{1}, \cdots, P_{n}$ running time $t_{1}, \cdots, t_{n}$ and input $x$, the outputs of the following two experiments are computational indistinguishable: Real experiment

- $\left(\mathrm{pk}_{j}, \mathrm{sk}_{j}\right) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}\right)$ for $\forall j \in[n]$.
- $\widetilde{D}_{j}=\left(\right.$ digest $\left._{j}, \hat{\mathrm{D}}_{j}\right) \leftarrow \operatorname{EncData}\left(\operatorname{crs}, D_{j}\right)$ for $\forall j \in[n]$.
- $\left(\right.$ ct $_{0}$, x_secret $) \leftarrow \operatorname{InpEnc}($ crs, $x)$.
- $\mathrm{ct}_{j} \leftarrow \operatorname{Eval}\left(\mathrm{crs}, j,\left\{\mathrm{pk}_{k}\right\}_{k=j+1}^{n}, \mathrm{ct}_{j-1}, \mathrm{sk}_{j},\left(P_{j}, t_{j}\right)\right.$, digest $\left._{j}\right)$ for $\forall j \in[n]$.
- Output $\mathrm{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}$.

Simulated experiment

- $\left(\mathrm{pk}_{i}, \mathrm{sk}_{i}\right) \leftarrow \operatorname{ihopSim}\left(1^{\lambda}, i\right)$.
- $\left(\mathrm{pk}_{j}, \mathrm{sk}_{j}\right) \leftarrow \operatorname{KeyGen}\left(1^{\lambda} ; r_{j}\right)$ for $\forall j \in[n] \backslash\{i\}$. Here, $r_{j}$ are uniform random coins.
- $\left(\mathrm{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}\right) \leftarrow \operatorname{ihopSim}\left(\operatorname{crs}, x,\left\{\mathrm{pk}_{j}, \mathrm{sk}_{j}, D_{j}, t_{j}\right\}_{j \in[n]},\left\{P_{j}, r_{j}\right\}_{j \in[n] \backslash\{i\}}\right.$, MemAccess $\left._{i}, y_{i}\right)$, where $y_{i}=P_{i}^{D_{i}}\left(\cdots\left(P_{1}{ }^{D_{1}}(x)\right) \cdots\right)$.
- Output ct ${ }_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}$.

[^10]The above definition can be made semi-malicious by allowing the adversary to pick random coins $r_{j}$ adversarially given the public key $\mathrm{pk}_{i}$ of honest client as follows: $\left\{r_{j}\right\}_{j \in[n] \backslash\{i\}} \leftarrow \mathcal{A}\left(1^{\lambda}, \mathrm{crs}^{\prime}, \mathrm{pk}_{i}\right)$ that will be used to define $\left(\mathrm{pk}_{j}, \mathrm{sk}_{j}\right)$ in Step 2. Our proof would also support this stronger setting as well.

Construction of ihopSim: We describe the two phases of ihopSim. In the first phase, ihopSim generates the keys of honest client $Q_{i}$ as $\left(\mathrm{pk}_{i}, \mathrm{sk}_{i}\right) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}\right)$.

In the second phase, ihopSim is described in Figure 22. At a high level, ihopSim generates everything honestly except $\mathrm{ct}_{i}$. When generating $\mathrm{ct}_{i}$, it simulates the step circuits one by one from the last to the first using the output $y_{i}$ and memory access MemAccess $_{i}$. In particular, since ihopSim takes $D_{i}$ and MemAccess ${ }_{i}$ as input, it can compute $D_{i}$ and digest ${ }_{i}$ before every step circuit, and use that to compute the output of every step circuit. Security follows from security of re-randomization of SFE, namely re-randomized garbled circuits are indistinguishable from freshly generated ones and that freshly generated garbled circuits are indistinguishable from simulated ones.

Now we give a series of hybrids such that the first hybrid outputs $\left(\operatorname{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}\right)$ in the real execution, and the last hybrid is the output of ihopSim. Notice that the only difference between the real and ideal experiments is $\mathrm{ct}_{i}$, so all the hybrids generate everything in the same way except $\mathrm{ct}_{i}$.

- $\hat{H}_{0}$ : Output in the real experiment.
- $\mathrm{H}_{0}$ : In this hybrid, replace $F_{s_{i}}(\cdot)$ with a truly random function $F$. In particular, when computing $\mathrm{ct}_{i}$ as in Figure 16, in steps 1a and 4, use the values generated by $F$; in step 3a embed labels corresponding to the values from $F$. The indistinguishability of this hybrid with $\hat{\mathrm{H}}_{0}$ follows from the pseudo-randomness of $F_{s_{i}}(\cdot)$ and privacy of oblivious transfer $\left(s_{i}\right.$ is hidden in $\mathrm{pk}_{i}$ ).
- $\mathrm{H}_{m}\left(m \in\left[T_{i}\right]\right):$ Next we consider a sequence of hybrids $\mathrm{H}_{1}, \cdots, \mathrm{H}_{T_{i}}$. The description of $\mathrm{H}_{m}$ is in Figure 23. Notice that $\mathrm{ct}_{i}$ consists of $T_{i}$ step circuits with corresponding PRF circuits. In hybrid $\mathbf{H}_{m}$, the step circuits from 1 to $m$ are simulated while the remaining step circuits ( $m+1$ to $T_{i}$ ) are generated honestly. Given all the programs and secret keys, the intermediate outputs as well as input/output labels of every step circuit can all be computed. Given the correct output of circuit $C_{m}^{\text {step }}$, the step circuits from 1 to $m$ can be simulated one by one from the $m$-th to the first similarly as in ihopSim.
To show $\mathrm{H}_{m}$ is indistinguishable from $\mathrm{H}_{m-1}$, first notice that they are the same except $\left(\widetilde{\mathrm{C}}_{m}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{m, j}^{\mathrm{PRF}}, \mathrm{L}_{m, j}\right\}_{j=i+1}^{n}\right)$ in $\mathrm{ct}_{i}$. Consider an intermediate hybrid $\hat{\mathrm{H}}_{m}$ which is the same as $\mathrm{H}_{m}$ except that in step 2 f when $\tau=m$, follow the steps in Figure 24. In particular, when $\tau=m, \hat{\mathrm{H}}_{m}$ computes the output of $\mathrm{C}_{m}^{\text {step }}$ and uses that output to simulate $\widetilde{\mathrm{C}}_{m}^{\text {step }}$ by CircSim and $\mathrm{OT}_{2}$. The output of $\widetilde{\mathrm{C}}_{m}^{\text {step }}$ is the same for $\mathrm{H}_{m-1}$ and $\hat{\mathrm{H}}_{m}$. The indistinguishibility of $\hat{\mathrm{H}}_{m}$ and $\mathbf{H}_{m-1}$ follows from the security of garbled circuits directly when $m \in\left[T_{i-1}+1, T_{i}\right]$. When $m \in\left[T_{i-1}\right]$, it follows from the security of garbled circuits and re-randomization. More precisely, the re-randomized garbled circuit is indistinguishable from a freshly generated garbled circuit, which is indistinguishable from a simulated one. Notice that the random coins used in re-randomization for $\stackrel{\widetilde{\mathrm{C}}}{m}_{\text {step }}$ is $F\left(\mathrm{GC}_{-} \star \| m\right)$, which is not used anywhere else in $\mathrm{H}_{m-1}$, so it can be treated as truly random coins.
$\left(\operatorname{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}\right) \leftarrow \operatorname{ihopSim}\left(\operatorname{crs}, x,\left\{\mathrm{pk}_{j}, \mathrm{sk}_{j}, D_{j}, t_{j}\right\}_{j \in[n]},\left\{P_{j}\right\}_{j \in[n] \backslash\{i\}}\right.$, MemAccess $\left._{i}, y_{i}\right)$.

1. Compute $\widetilde{D}_{j}=\left(\right.$ digest $\left._{j}, \hat{\mathrm{D}}_{j}\right) \leftarrow \operatorname{EncData}\left(\operatorname{crs}, D_{j}\right)$ for $\forall j \in[n]$.

Compute $\left(\mathrm{ct}_{0}\right.$, x_secret $) \leftarrow \operatorname{InpEnc}(x)$.
Compute $\mathrm{ct}_{j} \leftarrow \operatorname{Eval}\left(j,\left\{\mathrm{pk}_{k}\right\}_{k=j+1}^{n}, \mathrm{ct}_{j-1}, \mathrm{sk}_{j},\left(P_{j}, t_{j}\right)\right.$, digest $\left.{ }_{j}\right)$ for every $j \in[i-1]$.
Pick a random function $F$ (in the following use random values for $F(\cdot)$ ).
2. Let $T_{j}:=\sum_{k \in[j]} t_{k}$. Generate $\mathrm{ct}_{i}$ as follows:
(a) Run the program $P_{i-1}{ }^{D_{i-1}}\left(\cdots\left(P_{1}^{D_{1}}(x)\right) \cdots\right)$ to obtain (R/W ${ }^{\tau}$, $L^{\tau}$, wData ${ }^{\tau}$ ) for every CPU step $\tau \in\left[T_{i-1}\right]$. Obtain (R/W ${ }^{\tau}, L^{\tau}$, wData $\left.^{\tau}\right)$ for $\tau \in\left[T_{i-1}+1, T_{i}\right]$ from MemAccess ${ }_{i}$.
(b) $\left(\right.$ stateKeys $^{T_{i}+1}$, dataKeys $^{T_{i}+1}$, digestKeys $\left.^{T_{i}+1}\right) \leftarrow$ SampleKeys $\left(0^{*}\right)$.

Compute stateLabels ${ }^{T_{i}+1}$ using $y_{i}$; compute dataLabels ${ }^{T_{i}+1}$, digestLabels ${ }^{T_{i}+1}$ using ( $D_{i}$, digest $_{i}, \mathrm{R} / \mathrm{W}, L$, wData) of the last CPU step.
(c) For $\tau=T_{i}$ downto 1, do the following:
$(\mathrm{R} / \mathrm{W}, L, \mathrm{wData}):=\left(\mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}, \mathrm{wData}^{\tau}\right)$.
Define $j$ s.t. $\tau \in\left[T_{j-1}+1, T_{j}\right]$.
Let $D$ be the database of $Q_{j}$ before step $\tau$.

$$
\begin{aligned}
& \left(\text { stateLabels }^{\tau+1}, \text { dataLabels }^{\tau+1}, \text { digestLabels }^{\tau+1}\right) \\
& \left.\leftarrow \text { Transform }^{\left(\left(\text {stateLabels }^{\tau+1}, \text { dataLabels }^{\tau+1},\right.\right. \text { digestLabels }}{ }^{\tau+1}\right) \\
& \left.\quad\left\{F_{s_{i+1}}\left(\text { GC }_{-} \star \| \tau+1\right)\right\}_{\star \in\{\text { STATE,DATA }, \text { DIGEST }\}}\|\cdots\|\left\{F_{s_{n}}\left(\text { GC_}_{-} \star \| \tau+1\right)\right\}_{\star \in\{\text { STATE,DATA }, \text { DIGEST }\}}\right)
\end{aligned}
$$

```
if \(R / W=\) read then
\(\mathrm{e}_{\text {data }} \leftarrow \ell\) OTSim \(\left(\right.\) crs \(, D, L\), dataLabels \(\left.{ }^{\tau+1}\right)\).
            X\leftarrow(\mp@subsup{\mathrm{ stateLabels }}{}{\tau+1},\mp@subsup{\textrm{e}}{\mathrm{ data }}{},\mathrm{ digestLabels }}\mp@subsup{}{}{\tau+1})
```

else
$\mathrm{e}_{\text {digest }} \leftarrow \ell \mathrm{OTSimWrite}\left(\right.$ crs, $D, L$, wData, digestLabels $\left.{ }^{\tau+1}\right)$.
$X \leftarrow\left(\right.$ stateLabels $^{\tau+1}$, dataLabels ${ }^{\tau+1}$, edigest , wData $)$
$\left(\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\text {PRF }}\right\}_{j=i+1}^{n}\right\}, \operatorname{Labels}^{\tau}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathcal{U},(X, \mathrm{R} / \mathrm{W}, L)\right)$ such that the output labels of $\widetilde{\mathrm{C}}_{\tau, j}^{\mathrm{PRF}}$ are the same as input labels of RdLabs ${ }^{\tau, j}$ for $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$.
Parse Labels $^{\tau}=\left(\right.$ stateLabels $^{\tau}$, dataLabels $^{\tau}$, digestLabels $\left.^{\tau},\left\{\text { PLabels }^{\tau, j}\right\}_{j \in[i+1, n]}\right)$.
$\mathrm{L}_{\tau, j} \leftarrow \mathrm{OT}_{2}\left(\mathrm{pk}_{j},\left(\mathrm{PLabels}^{\tau, j}\right.\right.$, PLabels $\left.\left.^{\tau, j}\right)\right)$ for every $j \in[i+1, n]$.
if $\tau=T_{j-1}+1$ then
Embed stateLabels ${ }^{\tau}$ and digestLabels ${ }^{\tau}$ in $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$.
(stateKeys ${ }^{\tau}$, dataKeys ${ }^{\tau}$, digestKeys $\left.{ }^{\tau}\right) \leftarrow$ Transform (SampleKeys $\left(0^{*}\right)$,

$$
\left.\left\{F_{s_{j}}\left(\mathrm{GC}_{-} \star \| \tau\right)\right\}\|\cdots\|\left\{F_{s_{i-1}}\left(\mathrm{GC}_{-} \star \| \tau\right)\right\} \|\left\{F\left(\mathrm{GC}_{-} \star \| \tau\right)\right\}_{\star \in\{\text { STATE,DATA }, \mathrm{DIGEST}\}}\right)
$$

Compute (dataLabels ${ }^{\tau}$, digestLabels ${ }^{\tau}$ ) using ( $D_{j-1}$, digest $_{j-1}$, R/W, $L$, wData) at step $\tau-1$.
(d) $\mathrm{L}_{0} \leftarrow \mathrm{OT}_{2}\left(\right.$ ct $_{0},\left(\right.$ stateLabels $^{1}$, stateLabels $\left.\left.{ }^{1}\right)\right)$.
(e) $\mathrm{ct}_{i}:=\left(\mathrm{L}_{0},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\mathrm{PRF}}, \mathrm{L}_{\tau, j}\right\}_{j=i+1}^{n}\right\}_{\tau \in\left[T_{i}\right]}\right)$.
3. Compute $\mathrm{ct}_{j} \leftarrow \operatorname{Eval}\left(j,\left\{\mathrm{pk}_{k}\right\}_{k=j+1}^{n}, \mathrm{ct}_{j-1}, \mathrm{sk}_{j},\left(P_{j}, t_{j}\right)\right.$, digest $\left._{j}\right)$ for every $j \in[i+1, n]$.
4. Output $\mathrm{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}$.

Figure 22: Simulat $\oint^{3}$ of multi-hop RAM

1. Compute $\widetilde{D}_{j}=\left(\right.$ digest $\left._{j}, \hat{\mathrm{D}}_{j}\right) \leftarrow \operatorname{EncData}\left(\operatorname{crs}, D_{j}\right)$ for $\forall j \in[n]$.

Compute $\left(\right.$ ct $_{0}$, x_secret $) \leftarrow \operatorname{InpEnc}(x)$.
Compute $\mathrm{ct}_{j} \leftarrow \operatorname{Eval}\left(j,\left\{\mathrm{pk}_{k}\right\}_{k=j+1}^{n}, \mathrm{ct}_{j-1}, \mathrm{sk}_{j},\left(P_{j}, t_{j}\right)\right.$, digest $\left.{ }_{j}\right)$ for every $j \in[i-1]$.
Pick a random function $F$ (in the following use random values for $F(\cdot)$ ).
2. Let $T_{j}:=\sum_{k \in[j]} t_{k}$. Generate $\mathrm{ct}_{i}$ as follows:
(a) Compute $\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\mathrm{PRF}}, \mathrm{L}_{\tau, j}\right\}_{j=i+1}^{n}\right\}_{\tau \in\left[m+1, T_{i}\right]}$ honestly as in Figure 16.
(b) Run the program $P_{i}{ }^{D_{i}}\left(\cdots\left(P_{1}{ }^{D_{1}}(x)\right) \cdots\right)$ to obtain (state $\left.{ }^{\tau}, \mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}, \mathrm{wData}^{\tau}\right)$ for every $\tau \in\left[T_{i}\right]$.
(c) Define $j$ s.t. $m \in\left[T_{j-1}+1, T_{j}\right]$.
(d) Set stateKeys ${ }^{m+1}$, dataKeys ${ }^{m+1}$, digestKeys ${ }^{m+1}$ as SampleKeys $\left(F_{s_{j}}\left(\right.\right.$ GC_ $\left._{-} \star \| m+1\right)$ ) where $\star \in\{$ STATE, DATA, DIGEST $\}$, respectively.
If $m=T_{j}$, then set stateKeys ${ }^{m+1}$, dataKeys $^{m+1}$, digestKeys ${ }^{m+1}$ to SampleKeys $\left(0^{*}\right)$.
If $j<i$, then (stateKeys ${ }^{m+1}$, dataKeys ${ }^{m+1}$, digestKeys ${ }^{m+1}$ )

$$
\left.\begin{array}{l}
\leftarrow \text { Transform }\left(\left(\text { stateKeys }^{m+1}, \text { dataKeys }^{m+1}, \text { digestKeys }^{m+1}\right),\right. \\
\left.\left\{F_{s_{j+1}}\left(\text { GC }_{-} \star \| m+1\right)\right\}\|\cdots\|\left\{F_{s_{i-1}}\left(\text { GC }_{-} \star \| m+1\right)\right\} \|\left\{F\left(\text { GC }_{-} \star \| m+1\right)\right\}_{\star \in\{\text { STATET,DATA }} \text {,DigesT }\right\}
\end{array}\right) .
$$

(e) Compute (stateLabels ${ }^{m+1}$, dataLabels ${ }^{m+1}$, digestLabels $^{m+1}$ ) using ( state $^{m}, \mathrm{R} / \mathrm{W}^{m}, L^{m}$, wData $^{m}$ ) and $\left(D_{j}\right.$, digest $\left._{j}\right)$ at step $m$.
(f) For $\tau=m$ downto 1 , do the following: Follow the same steps as in Figure 22 step 2c.
(g) $\mathrm{L}_{0} \leftarrow \mathrm{OT}_{2}\left(\right.$ ct $_{0},\left(\right.$ stateLabels $^{1}$, stateLabels $\left.\left.{ }^{1}\right)\right)$.
(h) $\mathrm{ct}_{i}:=\left(\mathrm{L}_{0},\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {ste }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\mathrm{PRF}}, \mathrm{L}_{\tau, j}\right\}_{j=i+1}^{n}\right\}_{\tau \in\left[T_{i}\right]}\right)$.
3. Compute $\mathrm{ct}_{j} \leftarrow \operatorname{Eval}\left(j,\left\{\mathrm{pk}_{k}\right\}_{k=j+1}^{n}, \mathrm{ct}_{j-1}, \mathrm{sk}_{j},\left(P_{j}, t_{j}\right)\right.$, digest $\left.{ }_{j}\right)$ for every $j \in[i+1, n]$.
4. Output $\mathrm{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}$.

Figure 23: Hybrid $\mathrm{H}_{m}$
(R/W, $L$, wData) $:=\left(\mathrm{R} / \mathrm{W}^{\tau}, L^{\tau}, \mathrm{wData}^{\tau}\right)$.
Define $j$ s.t. $\tau \in\left[T_{j-1}+1, T_{j}\right]$.
Let ( $D$, digest) be the database and digest of $Q_{j}$ before step $\tau$.
Let state' be the CPU state after step $\tau$.
$\left(\right.$ stateKeys $^{\tau+1}{ }^{\tau}$ dataKeys $^{\tau+1}$, digestKeys $\left.{ }^{\tau+1}\right)$
$\leftarrow$ Transform $\left(\left(\right.\right.$ stateKeys $^{\tau+1}$, dataKeys $^{\tau+1}$, digestKeys $\left.^{\tau+1}\right)$,

$$
\left.\left\{F_{s_{i+1}}\left(\mathrm{GC}_{-} \star \| \tau+1\right)\right\}_{\star \in\{\operatorname{STATE}, \mathrm{DATA}, \mathrm{DIGEST}\}}\|\cdots\|\left\{F_{s_{n}}\left(\mathrm{GC}_{-} \star \| \tau+1\right)\right\}_{\star \in\{\operatorname{STATE}, \mathrm{DATA}, \mathrm{DIGEST}\}}\right)
$$

if $R / W=$ read then
$\mathrm{e}_{\text {data }} \leftarrow \operatorname{Send}\left(\right.$ crs, digest, $L$, dataKeys ${ }^{\tau+1} ; F(\mathrm{PSI} \| \tau) \oplus \bigoplus_{j>i} F_{s_{j}}($ LACONIC_OT $\left.\| \tau)\right)$.
$X \leftarrow\left(\right.$ stateKeys $_{\text {state }}{ }^{\tau+1}$, e $_{\text {data }}$, digestKeys $\left._{\text {digest }}^{\tau+1}\right)$.
else
$\mathrm{e}_{\text {digest }} \leftarrow$ SendWrite $\left(\right.$ crs, digest, $L$, wData, digestKeys ${ }^{\tau+1} ; F(\mathrm{PSI} \| \tau) \oplus \bigoplus_{j>i} F_{s_{j}}($ LACONIC_OT $\left.\| \tau)\right)$.
$X \leftarrow\left(\right.$ stateKeys $_{\text {state }^{\prime}}^{\tau+1}$, dataKeys $_{0}^{\tau+1}, \mathrm{e}_{\text {digest }}$, wData $)$.
$\left(\left\{\widetilde{\mathrm{C}}_{\tau}^{\text {step }},\left\{\widetilde{\mathrm{C}}_{\tau, j}^{\text {PRF }}\right\}_{j=i+1}^{n}\right\}\right.$, Labels $\left.^{\tau}\right) \leftarrow \operatorname{CircSim}\left(1^{\lambda}, \mathcal{U},(X, \mathrm{R} / \mathrm{W}, L)\right)$ such that the output labels of $\widetilde{\mathrm{C}}_{\tau, j}^{\text {PRF }}$ are the same as input labels of $\mathrm{RdLabs}^{\tau, j}$ for $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$.
Parse Labels ${ }^{\tau}=\left(\right.$ stateLabels $^{\tau}$, dataLabels $^{\tau}$, digestLabels $\left.^{\tau},\left\{\text { PLabels }^{\tau, j}\right\}_{j \in[i+1, n]}\right)$.
$\mathrm{L}_{\tau, j} \leftarrow \mathrm{OT}_{2}\left(\mathrm{pk}_{j},\left(\right.\right.$ PLabels $^{\tau, j}$, PLabels $\left.\left.^{\tau, j}\right)\right)$ for every $j \in[i+1, n]$.
if $\tau=T_{j-1}+1$ then
Embed stateLabels ${ }^{\tau}$ and digestLabels ${ }^{\tau}$ in $\widetilde{\mathrm{C}}_{\tau}^{\text {step }}$.
(stateKeys ${ }^{\tau}$, dataKeys ${ }^{\tau}$, digestKeys $\left.^{\tau}\right) \leftarrow$ Transform $\left(\right.$ SampleKeys $\left(0^{*}\right)$,

$$
\left.\left\{F_{s_{j}}\left(\operatorname{GC}_{-} \star \| \tau\right)\right\}\|\cdots\|\left\{F_{s_{i-1}}\left(\text { GC }_{-} \star \| \tau\right)\right\} \|\left\{F\left(\text { GC }_{-} \star \| \tau\right)\right\}_{\star \in\{\text { STATE,DATA,DIGEST }\}}\right)
$$

Compute (dataLabels ${ }^{\tau}$, digestLabels $^{\tau}$ ) using ( $D_{j-1}$, digest $_{j-1}$, R/W, $L$, wData) at step $\tau-1$.
Figure 24: Difference of $\mathrm{H}_{m}$ and $\hat{\mathrm{H}}_{m}$.

To switch from $\hat{\mathbf{H}}_{m}$ to $\mathrm{H}_{m}$, we replace $X$ in Figure 24 with simulated $\mathrm{e}_{\text {data }}$ and $\mathrm{e}_{\text {digest }}$ for CircSim and $\mathrm{OT}_{2}$. The indistinguishability follows from sender privacy of updatable laconic OT and that Send and SendWrite both take random coins $F($ PSI $\| m)$.

- $\hat{\mathrm{H}}_{T_{i}}$ : Output in the simulated experiment. This hybrid is the same as $\mathrm{H}_{T_{i}}$.

Extension. The above proof can be naturally extended to provide security for multiple clients and many executions. For example in the case of two clients $Q_{i_{1}}$ and $Q_{i_{2}}$, ihopSim first computes $\left(\mathrm{ct}_{0},\left\{\widetilde{D}_{j}\right\}_{j \in[n]},\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{1}-1\right]}\right)$ honestly, then computes $\mathrm{ct}_{i_{1}}$ same as in Figure 22 step 2. It then computes $\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{1}+1, i_{2}-1\right]}$ from $\mathrm{ct}_{i_{1}}$ by Eval, and computes $\mathrm{ct}_{i_{2}}$ same as in Figure 22 step 2. ${ }^{15}$ Finally it computes $\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{2}+1, n\right]}$ from $\mathrm{ct}_{i_{2}}$ by Eval. To show this is indistinguishable from the real execution, we consider the following hybrids:

- $\mathrm{H}_{0}$ : Output in the real experiment.
- $\mathrm{H}_{1}$ : First compute $\left(\mathrm{ct}_{0},\left\{\widetilde{D}_{j}\right\}_{j \in[n]},\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{2}-1\right]}\right)$ honestly, and then compute $\mathrm{ct}_{i_{2}}$ same as in Figure 22 step 2. Finally it computes $\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{2}+1, n\right]}$ from $\mathrm{ct}_{i_{2}}$ honestly by Eval.
- $\mathrm{H}_{2}$ : Output in the simulated experiment.

The above hybrids are indistinguishable because an honestly generated $\mathrm{ct}_{i_{1}}$ or $\mathrm{ct}_{i_{2}}$ is indistinguishable from a simulated one, as we have shown in the single-client case.

To simulate multiple executions, ihopSim can simply repeat the procedure for every execution. We note that there is no connection between executions except the digest, so they can be simulated separately given the initial digests of every execution. The hybrids go from the real experiment to the simulated experiment by replacing all the honestly generated ct's in one execution by simulated ones, one execution per hybrid.

### 7.3.5 UMA to Full Security for Multi-hop RAM Scheme

In this section we provide a fully secure multi-hop RAM scheme. We first review Oblivious RAM (ORAM), which was first introduced by Goldreich [Gol87, GO96] and Ostrovsky [Ost90, Ost92, GO96]. We then use ORAM as a compiler to encode the memory and program into a special format that does not reveal the access pattern or data contents during an execution.

Definition 7.14 (Oblivious RAM). An Oblivious RAM scheme consists of two procedures (OData, OProg) with syntax:

- $\left(D^{*}, s^{*}\right) \leftarrow \operatorname{OData}\left(1^{\lambda}, D\right):$ Given a security parameter $\lambda$ and memory $D \in\{0,1\}^{M}$ as input, OData outputs the encoded memory $D^{*}$ and encoding key $s^{*}$.
- $P^{*} \leftarrow \operatorname{OProg}\left(1^{\lambda}, 1^{\log M}, 1^{t}, P\right):$ Given a security parameter $\lambda$, a memory size $M$, and a program $P$ that runs in time $t$, OProg outputs an oblivious program $P^{*}$ that can access $D^{*}$ as $R A M$ and takes two inputs $x$ and $s^{*}$.

[^11]Efficiency. We require that the run-time of OData should be $M \cdot \operatorname{polylog}(M) \cdot \operatorname{poly}(\lambda)$, and the run-time of OProg should be $t \cdot$ poly $(\lambda) \cdot$ polylog $(M)$. Finally, the oblivious program $P^{*}$ itself should run in time $t^{\prime}=t \cdot \operatorname{poly}(\lambda) \cdot \operatorname{polylog}(M)$. Both the new memory size $M^{\prime}=\left|D^{*}\right|$ and the running time $t^{\prime}$ should be efficiently computable from $M, t$, and $\lambda$.

Correctness. Let $P_{1}, \ldots, P_{\ell}$ be programs running in polynomial times $t_{1}, \ldots, t_{\ell}$ on memory $D$ of size $M$. Let $x_{1}, \ldots, x_{\ell}$ be the inputs and $\lambda$ be a security parameter. Then we require that:

$$
\operatorname{Pr}\left[\left(P_{1}^{*}\left(x_{1}, s^{*}\right), \ldots, P_{\ell}^{*}\left(x_{\ell}, s^{*}\right)\right)^{D^{*}}=\left(P_{1}\left(x_{1}\right), \ldots, P_{\ell}\left(x_{\ell}\right)\right)^{D}\right]=1
$$

where $\left(D^{*}, s^{*}\right) \leftarrow \operatorname{OData}\left(1^{\lambda}, D\right), P_{i}^{*} \leftarrow \operatorname{OProg}\left(1^{\lambda}, 1^{\log M}, 1^{t}, P_{i}\right)$ and $\left(P_{1}^{*}\left(x_{1}, s^{*}\right), \ldots, P_{\ell}^{*}\left(x_{\ell}, s^{*}\right)\right)^{D^{*}}$ indicates running the ORAM programs on $D^{*}$ sequentially.

Security. For security, we require that there exists a PPT simulator $S$ such that for any sequence of programs $P_{1}, \ldots, P_{\ell}$, initial memory data $D \in\{0,1\}^{M}$, and inputs $x_{1}, \ldots, x_{\ell}$ we have that:

$$
\left(D^{*}, \text { MemAccess }\right) \stackrel{c}{\approx} S\left(1^{\lambda}, 1^{M},\left\{1^{t_{i}}, y_{i}\right\}_{i=1}^{\ell}\right)
$$

where $\left(y_{1}, \ldots, y_{\ell}\right)=\left(P_{1}\left(x_{1}\right), \ldots, P_{\ell}\left(x_{\ell}\right)\right)^{D},\left(D^{*}, s^{*}\right) \leftarrow \operatorname{OData}\left(1^{\lambda}, D\right)$, and MemAccess corresponds to the access pattern of the CPU-step circuits during the sequential execution of the oblivious programs $\left(P_{1}^{*}\left(x_{1}, s^{*}\right), \ldots, P_{\ell}^{*}\left(x_{\ell}, s^{*}\right)\right)^{D^{*}}$.

We prove the following theorem.
Theorem 7.15. Assume there exists a UMA-secure multi-hop RAM scheme and an ORAM scheme. Then there exists a fully secure multi-hop RAM scheme. Moreover, we give a black-box construction of one given a UMA-secure multi-hop RAM and ORAM scheme.

Proof. We first give the construction of the scheme itself and then provide a construction of an appropriate simulator to prove security. Let (Setup, KeyGen, EncData, InpEnc, Eval, Dec) be a UMAsecure multi-hop RAM scheme and let (OData, OProg) be an ORAM scheme. We construct a new multi-hop RAM scheme ( $\widehat{\text { Setup }}, \widehat{\text { KeyGen, EncData }}, \widehat{\text { InpEnc }}, \widehat{\text { Eval }}, \widehat{\text { Dec }})$ as follows:

- $\widehat{\operatorname{Setup}}\left(1^{\lambda}\right):$ Generate crs same as Setup.
- $\widehat{\operatorname{KeyGen}}\left(1^{\lambda}\right)$ : Generate (pk, sk) same as KeyGen.
- EncData $(\operatorname{crs}, D)$ : Execute $\left(D^{*}, s^{*}\right) \leftarrow \operatorname{OData}\left(1^{\lambda}, D\right)$ followed by $\widetilde{D} \leftarrow \operatorname{EncData}\left(1^{\lambda}, D^{*}\right)$.
- $\widehat{\operatorname{lnpEnc}}(x):$ Execute $\left(c t, \mathrm{x} \_\right.$secret $) \leftarrow \operatorname{InpEnc}(x)$.
- $\widehat{\text { Eval }}\left(i,\left\{\mathrm{pk}_{j}\right\}_{j=i+1}^{n}, \mathrm{ct}, \mathrm{sk},(P, t)\right.$, digest $)$ : Execute $\left(P^{*}, t^{*}\right) \leftarrow \operatorname{OProg}\left(1^{\lambda}, 1^{\log M}, 1^{t}, P\right)$ followed by Eval $\left(i,\left\{\mathrm{pk}_{j}\right\}_{j=i+1}^{n}, \mathrm{ct}, \mathrm{sk},\left(P^{*}\left[s^{*}\right], t^{*}\right)\right.$, digest $)$, where $P^{*}$ has $s^{*}$ hard-coded inside it.
- $\widehat{\operatorname{Dec}}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}$ (x_secret, ct): Output $\operatorname{Dec}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}$ (x_secret, ct).

We prove that the construction above given by ( $\widehat{\text { Setup }}, \widehat{\text { KeyGen, }} \widehat{\text { EncData, }} \widehat{\operatorname{lnpEnc}}, \widehat{\text { Eval }}, \widehat{\text { Dec }})$ is a fully secure multi-hop RAM scheme.

Correctness in a single execution. First we prove correctness for a single execution, and then we will generalize to multiple executions. In a single execution, our goal is to demonstrate that

$$
\operatorname{Pr}\left[\widehat{\operatorname{Dec}}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}\left(\text { x_secret, } \mathrm{ct}_{n}\right)=P_{n}^{D_{n}}\left(\cdots\left(P_{1}^{D_{1}}(x)\right) \cdots\right)\right]=1,
$$

where $\widetilde{D}_{i} \leftarrow \widehat{\operatorname{EncData}}\left(1^{\lambda}, D_{i}\right),\left(\right.$ ct $_{0}, \mathrm{x} \_$secret $) \leftarrow \widehat{\operatorname{lnpEnc}}(x), \mathrm{ct}_{i} \leftarrow \widehat{\operatorname{Eval}}\left(i,\left\{\mathrm{pk}_{j}\right\}_{j=i+1}^{n}, \mathrm{ct}_{i-1}\right.$, sk, $P_{i}, t_{i}$, digest $\left._{i}\right)$.
By definition, $\widehat{\operatorname{Dec}}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}\left(\mathrm{x} \_\right.$secret, $\left.\mathrm{ct}_{n}\right)=\operatorname{Dec}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}\left(\mathrm{x} \_\right.$secret, $\left.\mathrm{ct}_{n}\right)$. By the correctness of the UMA-secure multi-hop RAM scheme, we have that $\operatorname{Dec}^{\widetilde{D}_{1}, \cdots, \widetilde{D}_{n}}\left(\mathrm{x}_{\text {_secret }}, \mathrm{ct}_{n}\right)=P_{n}^{*}\left[s_{n}^{*}\right]^{D_{n}^{*}}\left(\cdots\left(P_{1}^{*}\left[s_{1}^{*}\right]^{D_{1}^{*}}(x)\right) \cdots\right)$. Finally, by the correctness of the ORAM scheme, $P_{n}^{*}\left[s_{n}^{*}\right]^{D_{n}^{*}}\left(\cdots\left(P_{1}^{*}\left[s_{1}^{*}\right]^{D_{1}^{*}}(x)\right) \cdots\right)=P_{n}{ }^{D_{n}}\left(\cdots\left(P_{1}{ }^{D_{1}}(x)\right) \cdots\right)$.

Correctness in multiple executions. To prove correctness in multiple executions, we need to show that

$$
\operatorname{Pr}\left[\widehat{\operatorname{Dec}}^{\widetilde{D}_{1}^{\text {(sid) })}, \cdots, \widetilde{D}_{n}^{\text {sid) }}}\left(\mathrm{x} \_ \text {_secret }^{(\text {sid })}, \mathrm{ct}_{n}^{(\text {sid })}\right)=P_{n}^{(\text {sid })} D_{n}^{(\text {sid })}\left(\cdots\left(P_{1}^{(\text {sid })} D_{1}^{(\text {sid })}\left(x^{(\text {sid })}\right)\right) \cdots\right)\right]=1,
$$

where $\widetilde{D}_{i}^{\text {(sid) }}$ is the resulting garbled database after executing sid- 1 homomorphic evaluations, $\left(\mathrm{ct}_{0}^{(\text {sid })}, \mathrm{x}_{2} \operatorname{secret}{ }^{(\text {sid })}\right) \leftarrow \widehat{\operatorname{mpEnc}}\left(x^{(\text {sid })}\right), \mathrm{ct}_{i}^{(\text {sid })} \leftarrow \widehat{\operatorname{Eval}}\left(i,\left\{\mathrm{pk}_{j}\right\}_{j=i+1}^{n}, \mathrm{ct}_{i-1}^{(\text {sid })}\right.$, sk, $P_{i}^{(\text {sid })}, t_{i}^{\text {(sid) }}$, digest $\left._{i}^{(\text {sid })}\right)$.

Recall that for correctness of the UMA-secure multi-hop RAM scheme we proved that the after every execution, the resulting garbled database $\widetilde{D}_{i}^{\text {(sid) }}$ corresponds to the output of Hash(crs, $D_{i}^{(\text {sid })}$ ), where $D_{i}^{(\text {sid })}$ is the correct $D_{i}^{*}$ resulting after previous sid-1 executions in the clear (see Property 2, Section 7.3.2). By correctness of ORAM and the underlying UMA-secure multi-hop RAM scheme, we conclude the correctness in multiple executions.

Security for a single client in a single execution. Server privacy is follows by receiver privacy of oblivious transfer. Now we prove client privacy for a single honest client $Q_{i}$ in a single execution. More precisely, we prove that there exists a PPT simulator ihopSim such that, for any set of databases $\left\{D_{j}\right\}_{j \in[n]}$, any sequence of compatible programs $P_{1}, \cdots, P_{n}$ running time $t_{1}, \cdots, t_{n}$ and input $x$, the outputs of the following two experiments are computational indistinguishable: Real experiment

- $\left(\mathrm{pk}_{j}, \mathrm{sk}_{j}\right) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}\right)$ for $\forall j \in[n]$.
- $\widetilde{D}_{j}=\left(\right.$ digest $\left._{j}, \hat{\mathrm{D}}_{j}\right) \leftarrow \operatorname{EncData}\left(\operatorname{crs}, D_{j}\right)$ for $\forall j \in[n]$.
- $\left(\mathrm{ct}_{0}\right.$, x_secret $) \leftarrow \operatorname{lnpEnc}(x)$.
- $\mathrm{ct}_{j} \leftarrow \operatorname{Eval}\left(j,\left\{\mathrm{pk}_{k}\right\}_{k=j+1}^{n}, \mathrm{ct}_{j-1}, \mathrm{sk}_{j},\left(P_{j}, t_{j}\right)\right.$, digest $\left._{j}\right)$ for $\forall j \in[n]$.
- Output ct ${ }_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}$.

Simulated experiment

- $\left(\mathrm{pk}_{i}, \mathrm{sk}_{i}\right) \leftarrow \operatorname{ihopSim}\left(1^{\lambda}, i\right)$.
- $\left(\mathrm{pk}_{j}, \mathrm{sk}_{j}\right) \leftarrow \operatorname{KeyGen}\left(1^{\lambda} ; r_{j}\right)$ for $\forall j \in[n] \backslash\{i\}$. Here, $r_{j}$ are uniform random coins.
- $\left(\operatorname{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}\right) \leftarrow \operatorname{ihopSim}\left(\operatorname{crs}, x,\left\{\mathrm{pk}_{j}, \mathrm{sk}_{j}, t_{j}\right\}_{j \in[n]},\left\{D_{j}, P_{j}, r_{j}\right\}_{j \in[n] \backslash\{i\}}, 1^{M_{i}}, y_{i}\right)$, where $y_{i}=P_{i}{ }^{D_{i}}\left(\cdots\left(P_{1}{ }^{D_{1}}(x)\right) \cdots\right)$.
- Output ct $0_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}$.

We let OSim be the ORAM simulator, and USim be the simulator for the UMA-secure multihop RAM scheme. We describe the two phases of ihopSim. In the first phase, ihopSim generates the keys of honest client $Q_{i}$ as $\left(\mathrm{pk}_{i}, \mathrm{sk}_{i}\right) \leftarrow \operatorname{KeyGen}\left(1^{\lambda}\right)$. In the second phase, ihopSim proceeds as follows.

1. Compute $\left(D_{i}^{*}\right.$, MemAccess $\left._{i}\right) \leftarrow \operatorname{OSim}\left(1^{\lambda}, 1^{M_{i}}, 1^{t_{i}}, y_{i}\right)$.
2. Compute $\left(D_{j}^{*}, s_{j}^{*}\right)$ from EncData and $P_{j}^{*}$ from Eval for every $j \in[n] \backslash\{i\}$.
3. Compute $\left(\mathrm{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}\right) \leftarrow \operatorname{USim}\left(\mathrm{crs}^{2}, x,\left\{\mathrm{pk}_{j}, \mathrm{sk}_{j}, t_{j}^{*}\right\}_{j \in[n]},\left\{D_{j}^{*}, P_{j}^{*}\left[s_{j}^{*}\right], r_{j}\right\}_{j \in[n] \backslash\{i\}}, D_{i}^{*}\right.$, MemAccess $_{i}$, $y_{i}$ ).
4. Output $\left(\mathrm{ct}_{0},\left\{\widetilde{D}_{j}, \mathrm{ct}_{j}\right\}_{j \in[n]}\right)$.

We now prove the output of the simulator is computationally indistinguishable from the real distribution.

- $\mathrm{H}_{0}$ : Output of the real experiment.
- $\mathrm{H}_{1}$ : Compute $\left(D_{j}^{*}, s_{j}^{*}\right)$ from EncData and $P_{j}^{*}$ from Eval for every $j \in[n] \backslash\{i\}$. Use the honestly generated $\left(D_{i}^{*}, s_{i}^{*}\right)$ from EncData and $P_{i}^{*}$ from Eval to execute the program $P_{i}^{*}\left[s_{i}^{*}\right]^{D_{i}^{*}}\left(\cdots\left(P_{1}^{*}\left[s_{1}^{*}\right]^{D_{1}^{*}}(x)\right) \cdots\right)$ and obtain $y_{i}$ and a sequence of memory accesses MemAccess ${ }_{i}$. Run $\left(\operatorname{ct}_{0},\left\{\widetilde{D}_{j}, \operatorname{ct}_{j}\right\}_{j \in[n]}\right) \leftarrow$ $\operatorname{USim}\left(\mathrm{crs}, x,\left\{\mathrm{pk}_{j}, \mathrm{sk}_{j}, t_{j}^{*}\right\}_{j \in[n]},\left\{D_{j}^{*}, P_{j}^{*}\left[s_{j}^{*}\right], r_{j}\right\}_{j \in[n] \backslash i\}}, D_{i}^{*}\right.$, MemAccess ${ }_{i}, y_{i}$ ) and output.
Since ( $D_{i}^{*}$, MemAccess $_{i}$ ) is the same as the real execution, the indistinguishability of this hybrid and $\mathrm{H}_{0}$ follows from UMA-security of the underlying multi-hop RAM scheme.
- $\mathrm{H}_{2}$ : Output of the simulated experiment. The only thing that differs in $\mathrm{H}_{1}$ and $\mathrm{H}_{2}$ is how we generate $D_{i}^{*}$ and MemAccess ${ }_{i}$. In $\mathrm{H}_{1}$ they are generated honestly and in $\mathrm{H}_{2}$ they are generated by OSim. $\mathrm{H}_{1} \stackrel{\mathrm{c}}{\approx} \mathrm{H}_{2}$ follows from the security of ORAM.

Security for multiple clients and multiple executions. Similar as in the proof of UMA security, the above proof can be naturally extended to provide security for multiple clients and many executions. For example in the case of two clients $Q_{i_{1}}$ and $Q_{i_{2}}$, ihopSim first computes $\left(\operatorname{ct}_{0},\left\{\widetilde{D}_{j}\right\}_{j \in[n]},\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{1}-1\right]}\right)$ honestly, then simulates $\mathrm{ct}_{i_{1}}$ same as above as if there were only one honest client $Q_{i_{1}}$. It then computes $\left\{\operatorname{ct}_{j}\right\}_{j \in\left[i_{1}+1, i_{2}-1\right]}$ from $\mathrm{ct}_{i_{1}}$ by Eval, and simulates $\mathrm{ct}_{i_{2}}$ same as above as if there were only one honest client $Q_{i_{2}}$. Notice that when simulating $\mathrm{ct}_{i_{2}}$, similar as in the UMA-secure scenario, ihopSim cannot generate ( $D_{i_{1}}^{*}, s_{i_{1}}^{*}, P_{i_{1}}^{*}$ ) honestly. Instead it will use the simulated ( $D_{i_{1}}^{*}$, MemAccess $i_{1}$ ) generated from OSim, and that is enough for the simulation. Finally it computes $\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{2}+1, n\right]}$ from $\mathrm{ct}_{i_{2}}$ by $\widehat{\text { Eval. }}$. To show this is indistinguishable from the real execution, we consider the following hybrids:

- $\mathrm{H}_{0}$ : Output in the real experiment.
- $\mathrm{H}_{1}$ : First compute $\left(\mathrm{ct}_{0},\left\{\widetilde{D}_{j}\right\}_{j \in[n]},\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{2}-1\right]}\right)$ honestly, and then compute $\mathrm{ct}_{i_{2}}$ same as above as if there were only one honest client $Q_{i_{2}}$. Finally it computes $\left\{\mathrm{ct}_{j}\right\}_{j \in\left[i_{2}+1, n\right]}$ from $\mathrm{ct}_{i_{2}}$ honestly by Eval.
- $\mathrm{H}_{2}$ : Output in the simulated experiment.

The above hybrids are indistinguishable because an honestly generated $\mathrm{ct}_{i_{1}}$ or $\mathrm{ct}_{i_{2}}$ is indistinguishable from a simulated one, as we have shown in the single-client case.

To simulate multiple executions, ihopSim should first use OSim to simulate ( $D_{i}^{*}$, MemAccess ${ }_{i}$ ) for every honest client $Q_{i}$ in all executions, and then repeat the above procedure for every execution. In the hybrids, we start from the real execution and first replace the honestly generated ct's by simulated ones while using honestly generated ( $D_{i}^{*}$, MemAccess ${ }_{i}$ ), and this step follows from the UMA-security of the underlying multi-hop RAM scheme. Afterwards we replace the honestly generated ( $D_{i}^{*}$, MemAccess ${ }_{i}$ ) by the output of OSim, and this step follows from the security of ORAM supporting multiple executions.

## Acknowledgement

We thank the anonymous reviewers of CRYPTO 2017 for their helpful suggestions in improving this paper. We also thank Yuval Ishai for useful discussions.

## References

[ADT11] Giuseppe Ateniese, Emiliano De Cristofaro, and Gene Tsudik. (If) size matters: Sizehiding private set intersection. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, PKC 2011: 14th International Conference on Theory and Practice of Public Key Cryptography, volume 6571 of Lecture Notes in Computer Science, pages 156-173, Taormina, Italy, March 6-9, 2011. Springer, Heidelberg, Germany.
[AIKW13] Benny Applebaum, Yuval Ishai, Eyal Kushilevitz, and Brent Waters. Encoding functions with constant online rate or how to compress garbled circuits keys. In Ran Canetti and Juan A. Garay, editors, Advances in Cryptology - CRYPTO 2013, Part II, volume 8043 of Lecture Notes in Computer Science, pages 166-184, Santa Barbara, CA, USA, August 18-22, 2013. Springer, Heidelberg, Germany.
[AIR01] William Aiello, Yuval Ishai, and Omer Reingold. Priced oblivious transfer: How to sell digital goods. In Birgit Pfitzmann, editor, Advances in Cryptology - EUROCRYPT 2001, volume 2045 of Lecture Notes in Computer Science, pages 119-135, Innsbruck, Austria, May 6-10, 2001. Springer, Heidelberg, Germany.
[ALSZ13] Gilad Asharov, Yehuda Lindell, Thomas Schneider, and Michael Zohner. More efficient oblivious transfer and extensions for faster secure computation. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, ACM CCS 13: 20th Conference
on Computer and Communications Security, pages 535-548, Berlin, Germany, November 4-8, 2013. ACM Press.
[BCCT12] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In Shafi Goldwasser, editor, ITCS 2012: 3rd Innovations in Theoretical Computer Science, pages 326-349, Cambridge, MA, USA, January 8-10, 2012. Association for Computing Machinery.
[Bea96] Donald Beaver. Correlated pseudorandomness and the complexity of private computations. In 28th Annual ACM Symposium on Theory of Computing, pages 479-488, Philadephia, PA, USA, May 22-24, 1996. ACM Press.
$\left[\mathrm{BGL}^{+} 15\right]$ Nir Bitansky, Sanjam Garg, Huijia Lin, Rafael Pass, and Sidharth Telang. Succinct randomized encodings and their applications. In Rocco A. Servedio and Ronitt Rubinfeld, editors, 47th Annual ACM Symposium on Theory of Computing, pages 439-448, Portland, OR, USA, June 14-17, 2015. ACM Press.
[BHHO08] Dan Boneh, Shai Halevi, Michael Hamburg, and Rafail Ostrovsky. Circular-secure encryption from decision Diffie-Hellman. In David Wagner, editor, Advances in Cryptology - CRYPTO 2008, volume 5157 of Lecture Notes in Computer Science, pages 108-125, Santa Barbara, CA, USA, August 17-21, 2008. Springer, Heidelberg, Germany.
[BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, ACM CCS 12: 19th Conference on Computer and Communications Security, pages 784-796, Raleigh, NC, USA, October 16-18, 2012. ACM Press.
[BPMW16] Florian Bourse, Rafaël Del Pino, Michele Minelli, and Hoeteck Wee. FHE circuit privacy almost for free. In Matthew Robshaw and Jonathan Katz, editors, Advances in Cryptology - CRYPTO 2016, Part II, volume 9815 of Lecture Notes in Computer Science, pages 62-89, Santa Barbara, CA, USA, August 14-18, 2016. Springer, Heidelberg, Germany.
[ $\left.\mathrm{BSCG}^{+} 13\right]$ Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Ran Canetti and Juan A. Garay, editors, Advances in Cryptology - CRYPTO 2013, Part II, volume 8043 of Lecture Notes in Computer Science, pages 90-108, Santa Barbara, CA, USA, August 18-22, 2013. Springer, Heidelberg, Germany.
[BV11a] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, 52nd Annual Symposium on Foundations of Computer Science, pages 97-106, Palm Springs, CA, USA, October 22-25, 2011. IEEE Computer Society Press.
[BV11b] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ringLWE and security for key dependent messages. In Phillip Rogaway, editor, Advances in Cryptology - CRYPTO 2011, volume 6841 of Lecture Notes in Computer Science,
pages 505-524, Santa Barbara, CA, USA, August 14-18, 2011. Springer, Heidelberg, Germany.
[CHJV15] Ran Canetti, Justin Holmgren, Abhishek Jain, and Vinod Vaikuntanathan. Succinct garbling and indistinguishability obfuscation for RAM programs. In Rocco A. Servedio and Ronitt Rubinfeld, editors, 47 th Annual ACM Symposium on Theory of Computing, pages 429-437, Portland, OR, USA, June 14-17, 2015. ACM Press.
[CHK04] Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In Christian Cachin and Jan Camenisch, editors, Advances in Cryptology - EUROCRYPT 2004, volume 3027 of Lecture Notes in Computer Science, pages 207-222, Interlaken, Switzerland, May 2-6, 2004. Springer, Heidelberg, Germany.
[COV15] Melissa Chase, Rafail Ostrovsky, and Ivan Visconti. Executable proofs, input-size hiding secure computation and a new ideal world. In Elisabeth Oswald and Marc Fischlin, editors, Advances in Cryptology - EUROCRYPT 2015, Part II, volume 9057 of Lecture Notes in Computer Science, pages 532-560, Sofia, Bulgaria, April 26-30, 2015. Springer, Heidelberg, Germany.
[CS98] Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Hugo Krawczyk, editor, Advances in Cryptology - CRYPTO'98, volume 1462 of Lecture Notes in Computer Science, pages 13-25, Santa Barbara, CA, USA, August 23-27, 1998. Springer, Heidelberg, Germany.
[CS02] Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Lars R. Knudsen, editor, Advances in Cryptology - EUROCRYPT 2002, volume 2332 of Lecture Notes in Computer Science, pages 45-64, Amsterdam, The Netherlands, April 28 - May 2, 2002. Springer, Heidelberg, Germany.
[CV12] Melissa Chase and Ivan Visconti. Secure database commitments and universal arguments of quasi knowledge. In Reihaneh Safavi-Naini and Ran Canetti, editors, Advances in Cryptology - CRYPTO 2012, volume 7417 of Lecture Notes in Computer Science, pages 236-254, Santa Barbara, CA, USA, August 19-23, 2012. Springer, Heidelberg, Germany.
[DG17] Nico Döttling and Sanjam Garg. Identity-based encryption from the diffie hellman assumption. CRYPTO 2017 (to appear), 2017.
[DS16] Léo Ducas and Damien Stehlé. Sanitization of FHE ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, Advances in Cryptology - EUROCRYPT 2016, Part I, volume 9665 of Lecture Notes in Computer Science, pages 294-310, Vienna, Austria, May 8-12, 2016. Springer, Heidelberg, Germany.
[FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In 31st Annual Symposium
on Foundations of Computer Science, pages 308-317, St. Louis, Missouri, October 2224, 1990. IEEE Computer Society Press.
[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, 41st Annual ACM Symposium on Theory of Computing, pages 169-178, Bethesda, MD, USA, May 31 - June 2, 2009. ACM Press.
[GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, Advances in Cryptology - EUROCRYPT 2013, volume 7881 of Lecture Notes in Computer Science, pages 1-17, Athens, Greece, May 26-30, 2013. Springer, Heidelberg, Germany.
[GGH ${ }^{+}$13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In 54th Annual Symposium on Foundations of Computer Science, pages 4049, Berkeley, CA, USA, October 26-29, 2013. IEEE Computer Society Press.
[GGMP16] Sanjam Garg, Divya Gupta, Peihan Miao, and Omkant Pandey. Secure multiparty RAM computation in constant rounds. In Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I, pages 491-520, 2016.
[GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, 45th Annual ACM Symposium on Theory of Computing, pages 467-476, Palo Alto, CA, USA, June 1-4, 2013. ACM Press.
[GHL $\left.{ }^{+} 14\right]$ Craig Gentry, Shai Halevi, Steve Lu, Rafail Ostrovsky, Mariana Raykova, and Daniel Wichs. Garbled RAM revisited. In Phong Q. Nguyen and Elisabeth Oswald, editors, Advances in Cryptology - EUROCRYPT 2014, volume 8441 of Lecture Notes in Computer Science, pages 405-422, Copenhagen, Denmark, May 11-15, 2014. Springer, Heidelberg, Germany.
[GHRW14] Craig Gentry, Shai Halevi, Mariana Raykova, and Daniel Wichs. Outsourcing private RAM computation. In 55th Annual Symposium on Foundations of Computer Science, pages 404-413, Philadelphia, PA, USA, October 18-21, 2014. IEEE Computer Society Press.
[GHV10] Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. i-Hop homomorphic encryption and rerandomizable Yao circuits. In Tal Rabin, editor, Advances in Cryptology CRYPTO 2010, volume 6223 of Lecture Notes in Computer Science, pages 155-172, Santa Barbara, CA, USA, August 15-19, 2010. Springer, Heidelberg, Germany.
$\left[\mathrm{GKK}^{+} 12\right]$ S. Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Fernando Krell, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure two-party computation in sublinear (amortized) time. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, ACM CCS 12: 19th Conference on Computer and Communications Security, pages 513-524, Raleigh, NC, USA, October 16-18, 2012. ACM Press.
[GKP $\left.{ }^{+} 13\right]$ Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. How to run turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, Advances in Cryptology - CRYPTO 2013, Part II, volume 8043 of Lecture Notes in Computer Science, pages 536-553, Santa Barbara, CA, USA, August 18-22, 2013. Springer, Heidelberg, Germany.
[GLO15] Sanjam Garg, Steve Lu, and Rafail Ostrovsky. Black-box garbled RAM. In Venkatesan Guruswami, editor, 56th Annual Symposium on Foundations of Computer Science, pages 210-229, Berkeley, CA, USA, October 17-20, 2015. IEEE Computer Society Press.
[GLOS15] Sanjam Garg, Steve Lu, Rafail Ostrovsky, and Alessandra Scafuro. Garbled RAM from one-way functions. In Rocco A. Servedio and Ronitt Rubinfeld, editors, 47 th Annual ACM Symposium on Theory of Computing, pages 449-458, Portland, OR, USA, June 14-17, 2015. ACM Press.
[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, 19th Annual ACM Symposium on Theory of Computing, pages 218-229, New York City, NY, USA, May 25-27, 1987. ACM Press.
[GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. J. ACM, 43(3):431-473, 1996.
[Gol87] Oded Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In Alfred Aho, editor, 19th Annual ACM Symposium on Theory of Computing, pages 182-194, New York City, NY, USA, May 25-27, 1987. ACM Press.
[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Non-interactive zaps and new techniques for NIZK. In Cynthia Dwork, editor, Advances in Cryptology - CRYPTO 2006, volume 4117 of Lecture Notes in Computer Science, pages 97-111, Santa Barbara, CA, USA, August 20-24, 2006. Springer, Heidelberg, Germany.
[GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, Advances in Cryptology - CRYPTO 2013, Part I, volume 8042 of Lecture Notes in Computer Science, pages 75-92, Santa Barbara, CA, USA, August 18-22, 2013. Springer, Heidelberg, Germany.
[HK12] Shai Halevi and Yael Tauman Kalai. Smooth projective hashing and two-message oblivious transfer. Journal of Cryptology, 25(1):158-193, January 2012.
[HW15] Pavel Hubacek and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In Tim Roughgarden, editor, ITCS 2015: 6th Innovations in Theoretical Computer Science, pages 163-172, Rehovot, Israel, January 11-13, 2015. Association for Computing Machinery.
[IKNP03] Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In Dan Boneh, editor, Advances in Cryptology - CRYPTO 2003, volume

2729 of Lecture Notes in Computer Science, pages 145-161, Santa Barbara, CA, USA, August 17-21, 2003. Springer, Heidelberg, Germany.
$\left[\mathrm{IKO}^{+}\right.$11] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, Manoj Prabhakaran, and Amit Sahai. Efficient non-interactive secure computation. In Kenneth G. Paterson, editor, Advances in Cryptology - EUROCRYPT 2011, volume 6632 of Lecture Notes in Computer Science, pages 406-425, Tallinn, Estonia, May 15-19, 2011. Springer, Heidelberg, Germany.
[IP07] Yuval Ishai and Anat Paskin. Evaluating branching programs on encrypted data. In Salil P. Vadhan, editor, TCC 2007: 4th Theory of Cryptography Conference, volume 4392 of Lecture Notes in Computer Science, pages 575-594, Amsterdam, The Netherlands, February 21-24, 2007. Springer, Heidelberg, Germany.
[IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In David Wagner, editor, Advances in Cryptology CRYPTO 2008, volume 5157 of Lecture Notes in Computer Science, pages 572-591, Santa Barbara, CA, USA, August 17-21, 2008. Springer, Heidelberg, Germany.
[Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In 20th Annual ACM Symposium on Theory of Computing, pages 20-31, Chicago, IL, USA, May 2-4, 1988. ACM Press.
[KK13] Vladimir Kolesnikov and Ranjit Kumaresan. Improved OT extension for transferring short secrets. In Ran Canetti and Juan A. Garay, editors, Advances in Cryptology CRYPTO 2013, Part II, volume 8043 of Lecture Notes in Computer Science, pages 54-70, Santa Barbara, CA, USA, August 18-22, 2013. Springer, Heidelberg, Germany.
[KLW15] Venkata Koppula, Allison Bishop Lewko, and Brent Waters. Indistinguishability obfuscation for turing machines with unbounded memory. In Rocco A. Servedio and Ronitt Rubinfeld, editors, 47 th Annual ACM Symposium on Theory of Computing, pages 419-428, Portland, OR, USA, June 14-17, 2015. ACM Press.
[LNO13] Yehuda Lindell, Kobbi Nissim, and Claudio Orlandi. Hiding the input-size in secure two-party computation. In Kazue Sako and Palash Sarkar, editors, Advances in Cryptology - ASIACRYPT 2013, Part II, volume 8270 of Lecture Notes in Computer Science, pages 421-440, Bengalore, India, December 1-5, 2013. Springer, Heidelberg, Germany.
[LO13] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In Thomas Johansson and Phong Q. Nguyen, editors, Advances in Cryptology - EUROCRYPT 2013, volume 7881 of Lecture Notes in Computer Science, pages 719-734, Athens, Greece, May 2630, 2013. Springer, Heidelberg, Germany.
[LP09] Yehuda Lindell and Benny Pinkas. A proof of security of Yao's protocol for two-party computation. Journal of Cryptology, 22(2):161-188, April 2009.
[MRK03] Silvio Micali, Michael O. Rabin, and Joe Kilian. Zero-knowledge sets. In 44 th Annual Symposium on Foundations of Computer Science, pages 80-91, Cambridge, MA, USA, October 11-14, 2003. IEEE Computer Society Press.
[NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In S. Rao Kosaraju, editor, 12th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 448-457, Washington, DC, USA, January 7-9, 2001. ACM-SIAM.
[OPP14] Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, Advances in Cryptology - CRYPTO 2014, Part I, volume 8616 of Lecture Notes in Computer Science, pages 536-553, Santa Barbara, CA, USA, August 17-21, 2014. Springer, Heidelberg, Germany.
[OPWW15] Tatsuaki Okamoto, Krzysztof Pietrzak, Brent Waters, and Daniel Wichs. New realizations of somewhere statistically binding hashing and positional accumulators. In Tetsu Iwata and Jung Hee Cheon, editors, Advances in Cryptology - ASIACRYPT 2015, Part I, volume 9452 of Lecture Notes in Computer Science, pages 121-145, Auckland, New Zealand, November 30 - December 3, 2015. Springer, Heidelberg, Germany.
[OS97] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In 29th Annual ACM Symposium on Theory of Computing, pages 294-303, El Paso, TX, USA, May 4-6, 1997. ACM Press.
[Ost90] Rafail Ostrovsky. Efficient computation on oblivious RAMs. In 22nd Annual ACM Symposium on Theory of Computing, pages 514-523, Baltimore, MD, USA, May 1416, 1990. ACM Press.
[Ost92] Rafail Ostrovsky. Software Protection and Simulation On Oblivious RAMs. PhD thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 1992.
[Rab81] Michael O. Rabin. How to exchange secrets with oblivious transfer, 1981.
[Vil12] Jorge Luis Villar. Optimal reductions of some decisional problems to the rank problem. In Xiaoyun Wang and Kazue Sako, editors, Advances in Cryptology - ASIACRYPT 2012, volume 7658 of Lecture Notes in Computer Science, pages 80-97, Beijing, China, December 2-6, 2012. Springer, Heidelberg, Germany.
[WHC $\left.{ }^{+} 14\right]$ Xiao Shaun Wang, Yan Huang, T.-H. Hubert Chan, Abhi Shelat, and Elaine Shi. SCORAM: Oblivious RAM for secure computation. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, ACM CCS 14: 21st Conference on Computer and Communications Security, pages 191-202, Scottsdale, AZ, USA, November 3-7, 2014. ACM Press.
[Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In 23rd Annual Symposium on Foundations of Computer Science, pages 160-164, Chicago, Illinois, November 3-5, 1982. IEEE Computer Society Press.


[^0]:    ${ }^{*}$ Research supported by a postdoc fellowship of the German Academic Exchange Service (DAAD).
    ${ }^{\dagger}$ Research supported in part from 2017 AFOSR YIP Award, DARPA/ARL SAFEWARE Award W911NF15C0210, AFOSR Award FA9550-15-1-0274, NSF CRII Award 1464397, and research grants by the Okawa Foundation, Visa Inc., and Center for Long-Term Cybersecurity (CLTC, UC Berkeley). The views expressed are those of the author and do not reflect the official policy or position of the funding agencies.
    ${ }^{\ddagger}$ Work done while at University of California, Berkeley.
    ${ }^{\text {§ }}$ Part of the work done while visiting University of California, Berkeley. Research supported in part the National Science Foundation under Grant No. 1617676, IBM under Agreement 4915013672, and the Packard Foundation under Grant 2015-63124.

[^1]:    ${ }^{1}$ We remark that related prior works on OT extension [Bea96, IKNP03, KK13, ALSZ13] makes the number of public key operations performed during protocol executions independent of the receiver's input size. However, the communication complexity of receivers in these protocols still grows with the input size of the receiver.

[^2]:    ${ }^{2}$ We remark that solutions for this problem based on fully-homomorphic encryption (FHE) [Gen09, LNO13], unlike our result, reduce the communication cost of both the sender's and the receiver's messages to be independent of the size of $D$, but require additional rounds of interaction.
    ${ }^{3}$ The communication cost of the receiver's message can be reduced to depend only on the running time of the program by allowing round complexity to grow with the running time of the program (using Merkle Hashing). Analogous to the circuit case, we remark that FHE-based solutions can make the communication of both the sender and the receiver small, but at the cost of extra rounds. Moreover, in the setting of RAM programs FHE-based solutions additionally incur an increased computational cost for the receiver. In particular, the receiver's computational cost grows with the size of its database.
    ${ }^{4}$ The receiver is required to keep the output of the computation private.

[^3]:    ${ }^{5}$ We finally note that relaxing to the weaker notion of indistinguishability-based security we can expect to obtain the best of both worlds, i.e. a non-interactive solution while making only a black-box use of the adversary (a.k.a. avoiding the use of extractability assumptions). We leave this open for future work.
    ${ }^{6}$ We remark that in an orthogonal work of Hubacek and Wichs [HW15] obtain constructions where the communication cost is independent of the length of the output of the computation using indistinguishability obfuscation $\left[\mathrm{GGH}^{+} 13 \mathrm{~b}\right]$.

[^4]:    ${ }^{7}$ One method for realizing homomorphic encryption for RAM programs [GKP ${ }^{+}$13, GHRW14, CHJV15, BGL ${ }^{+} 15$, KLW15] would be to use obfuscation $\left[\mathrm{GGH}^{+} 13 \mathrm{~b}\right]$ based on multilinear maps [GGH13a]. However, in this paper we focus on basing homomorphic RAM computation on DDH and defer the work on obfuscation to future work.
    ${ }^{8}$ Using NIZKs alone does not solve the problem, because locations accessed during computation are dynamically decided.

[^5]:    ${ }^{9}$ We note that the above idea of using laconic OT also gives a conceptually very simple solution for UMA secure garbled RAM scheme [LO13]. Moreover, there is a general transformation [GHL ${ }^{+}$14] that converts any UMA secure garbled RAM into one with full security via the usage of symmetric key encryption and oblivious RAM. This would give a simplified construction of fully secure garbled RAM under DDH assumption.

[^6]:    ${ }^{10}$ In general, the distinction between what to include in the program $P$, the memory data $D$ and the short input $x$ can be somewhat arbitrary. However as motivated by our applications we will typically be interested in a setting where the data $D$ is large while the size of the program $|P|$ and input length $m$ is small.

[^7]:    ${ }^{11}$ In fact, each OT message for keys of a wire can be randomized consistently just given the randomness used for that wire.

[^8]:    ${ }^{12}$ We do this by keeping track of the randomness used in randomizing the input wires for each garbled circuit.

[^9]:    ${ }^{13}$ Note that randomization of garbled circuits preserves the functionality. Since the keys for the next circuit are transferred using the laconic OT, we need to feed in correct keys into the Send functions of laconic OT.

[^10]:    ${ }^{14}$ It can also extended to the setting, when a client $Q_{i}$ occurs multiple times in the chain of clients in an execution. To handle this setting, the digest of $D_{i}$ is passed along all the programs between two instances of this client as additional state.

[^11]:    ${ }^{15}$ Notice that different from Figure 22, the simulator here doesn't take $P_{i_{1}}$ as input, but the simulation can still obtain (R/W $\left.{ }^{\tau}, L^{\tau}, \mathrm{wData}^{\tau}\right)$ for every step $\tau \in\left[T_{i_{1}-1}+1, T_{i_{1}}\right]$ given $D_{i_{1}}$ and MemAccess $i_{1}$, and that is enough for the simulation.

