

Adaptive Garbled RAM from Laconic Oblivious Transfer

Sanjam Garg*
University of California, Berkeley
sanjamg@berkeley.edu

Rafail Ostrovsky†
UCLA
rafail@cs.ucla.edu

Akshayaram Srinivasan
University of California, Berkeley
akshayaram@berkeley.edu

Abstract

We give a construction of an adaptive garbled RAM scheme. In the adaptive setting, a client first garbles a “large” persistent database which is stored on a server. Next, the client can provide garbling of multiple adaptively and adversarially chosen RAM programs that execute and modify the stored database arbitrarily. The garbled database and the garbled program should reveal nothing more than the running time and the output of the computation. Furthermore, the sizes of the garbled database and the garbled program grow only linearly in the size of the database and the running time of the executed program respectively (up to poly logarithmic factors). The security of our construction is based on the assumption that laconic oblivious transfer (Cho et al., CRYPTO 2017) exists. Previously, such adaptive garbled RAM constructions were only known using indistinguishability obfuscation or in random oracle model. As an additional application, we note that this work yields the first constant round secure computation protocol for persistent RAM programs in the malicious setting from standard assumptions. Prior works did not support persistence in the malicious setting.

1 Introduction

Over the years, garbling methods [Yao86, LP09, AIK04, BHR12b, App17] have been extremely influential and have engendered an enormous number of applications in cryptography. Informally, garbling a function f and an input x , yields the function encoding \hat{f} and the input encoding \hat{x} . Given \hat{f} and \hat{x} , there exists an efficient decoding algorithm that recovers $f(x)$. The security property requires that \hat{f} and \hat{x} do not reveal anything about f or x except $f(x)$. By now, it is well established that realizing garbling schemes [BHR12b, App17] is an important cryptographic goal.

*Research supported in part from DARPA/ARL SAFEWARE Award W911NF15C0210, AFOSR Award FA9550-15-1-0274, AFOSR YIP Award, DARPA and SPAWAR under contract N66001-15-C-4065, a Hellman Award 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.

†Research supported in part by NSF grant 1619348, DARPA SPAWAR contract N66001-15-1C-4065, US-Israel BSF grant 2012366, OKAWA Foundation Research Award, IBM Faculty Research Award, Xerox Faculty Research Award, B. John Garrick Foundation Award, Teradata Research Award, and Lockheed-Martin Corporation Research Award. The views expressed are those of the authors and do not reflect position of the Department of Defense or the U.S. Government.

One shortcoming of standard garbling techniques has been that the size of the function encoding grows linearly in the size of the circuit computing the function and thus leads to large communication costs. Several methods have been devised to overcome this constraint.

- Lu and Ostrovsky [LO13] addressed the question of garbling RAM program execution on a persistent garbled database. Here, the efficiency requirement is that the size of the function encoding grows only with the running time of the RAM program. This work has led to fruitful line of research [GHL⁺14, GLOS15, GLO15, LO17] that reduces the communication cost to grow linearly with running times of the programs executed, rather than the corresponding circuit sizes. A key benefit of this approach is that it has led to constructions based on one-way functions.
- Goldwasser, Kalai, Popa, Vaikuntanathan, and Zeldovich [GKP⁺13] addressed the question of reducing the communication cost by reusing the encodings. Specifically, they provided a construction of reusable garbled circuits based on standard assumptions (namely learning-with-errors). However, their construction needs input encoding to grow with the depth of the circuit being garbled.
- Finally, starting with Gentry, Halevi, Raykova, and Wichs [GHRW14], a collection of works [CHJV15, BGL⁺15, KLV15, CH16, CCHR16, ACC⁺16] have attempted to obtain garbling schemes where the size of the function encoding only grows with its description size and is otherwise independent of its running time on various inputs. However, these constructions are proven secure only assuming indistinguishability obfuscation [BGI⁺01, GGH⁺13].

A recurring theme in all the above research efforts has been the issue of *adaptivity*: Can the adversary adaptively choose the input after seeing the function encoding?

This task is trivial if one reveals both the function encoding and the input encoding together after the input is specified. However, this task becomes highly non-trivial if we require the size of the input encoding to only grow with the size of the input and independent of the complexity of computing f . The first solution to this problem was provided by Bellare, Hoang and Rogaway [BHR12a] for the case of circuits in the random oracle model [BR93]. Subsequently, several adaptive circuit garbling schemes have been obtained in the standard model from (i) one-way functions [HJO⁺16, JW16, JKK⁺17],¹ or (ii) using laconic OT [GS18a] which relies on public-key assumptions [CDG⁺17, DG17, DGHM18, BLSV18].

However, constructing adaptively secure schemes for more communication constrained settings has proved much harder. In this paper, we focus on the case of RAM programs. More specifically, adaptively secure garbled RAM is known only using random oracles (e.g. [LO13, GLOS15]) or under very strong assumptions such as indistinguishability obfuscation [CCHR16, ACC⁺16]. In this work, we ask:

Can we realize adaptively secure garbled RAM from standard assumptions?

Further motivating the above question, is the tightly related application of constructing *constant round* secure RAM computation over a persistent database in the malicious setting. More specifically, as shown by Beaver, Micali and Rogaway [BMR90] garbling techniques can be used

¹A drawback of these works is that the size of the input encoding grows with the width/depth of the circuit computing f

to realize constant round secure computation [Yao82, GMW87] constructions. Similarly, above-mentioned garbling schemes for RAM programs also yield constant round, communication efficient secure computation solutions [HY16, Mia16, GGMP16, KY18]. However, preserving persistence of RAM programs in the malicious setting requires the underlying garbling techniques to provide adaptive security.²

1.1 Our Results

In this work, we obtain a construction of adaptively secure garbled RAM based on the assumption that laconic oblivious transfer [CDG⁺17] exists. Laconic oblivious transfer can be based on a variety of public-key assumptions such as (i) Computation Diffie-Hellman Assumption [DG17], (ii) Factoring Assumption [DG17], or (iii) Learning-With-Errors Assumption [BLSV18, DGHM18]. In our construction, the size of the garbled database and the garbled program grow only linearly in the size of the database and the running time of the executed program respectively (up to poly logarithmic factors). The main result in our paper is:

Theorem 1.1 (Informal) *Assuming either the Computational Diffie-Hellman assumption or the Factoring assumption or the Learning-with-Errors assumption, there exists a construction of adaptive garbled RAM scheme where the time required to garble a database, a program and an input grows linearly (upto poly logarithmic factors) with the size of the database, running time of the program and length of the input respectively.*³

Additionally, plugging our adaptively secure garbled RAM scheme into a malicious secure constant round secure computation protocol yields a maliciously secure constant round secure RAM computation protocol [IKO⁺11, ORS15, BL18, GS18b] for a persistent database. Again, this construction is based on the assumption that laconic OT exists and the underlying assumptions needed for the constant round protocol.

2 Our Techniques

In this section, we outline the main challenges and the techniques used in our construction of adaptive garbled RAM.

Starting Point. In a recent result, Garg and Srinivasan [GS18a] gave a construction of adaptively secure garbled circuit transfer where the size of the input encoding grows only with the input and the output length. The main idea behind their construction is a technique to “linearize” a garbled circuit. Informally, a garbled circuit is said to be linearized if the simulation of particular garbled gate depends only on simulating one other gate (or in other words, the simulation dependency graph is a line). In order to linearize a garbled circuit, their work transforms a circuit into a sequence of CPU step circuits that can make read and write accesses at *fixed* locations in an external memory. The individual step circuits are garbled using a (plain) garbling scheme and the

²We note that adaptive security is not essential for obtaining protocols with round complexity that grows with the running time of the executed programs [OS97, GKK⁺12, WHC⁺14].

³As in the case of adaptively secure garbled circuits, the size of the input encoding must also grow with the output length of the program. Here, we implicitly assume that the input and the outputs have the same length.

access to the memory is mediated using a laconic OT.⁴ The use of laconic OT enables the above mentioned garbling scheme to have “linear” structure wherein the simulation of a particular CPU step depends only on simulating the previous step circuit.

A Generalization. Though the approach of Garg and Srinivasan shares some similarities with a garbling a RAM program (like garbling a sequence of CPU step circuits), there are some crucial differences.

1. The first difference is that unlike a circuit, the locations that are accessed by a RAM program are dynamically chosen depending on the program’s input.
2. The second difference is that the locations that are accessed might leak information about the program and the input and a garbled RAM scheme must protect against such leakages.

The first step we take in constructing an adaptive garbled RAM scheme is to generalize the above approach of Garg and Srinivasan [GS18a] to construct an adaptively secure garbled RAM scheme with weaker security guarantees. The security that we achieve is that of unprotected memory access [GHL⁺14]. Informally, a garbled RAM scheme is said to have unprotected memory access if both the contents of the database and the memory locations that are accessed are revealed in the clear. This generalization is given in Section 4.

In the non-adaptive setting, there are standard transformations (outlined in [GHL⁺14]) from a garbled RAM with unprotected memory access to a standard garbled RAM scheme where both the memory contents and the access patterns are hidden. This transformation involves the additional use of an ORAM scheme. Somewhat surprisingly, these transformations fail in the adaptive setting! The details follow.

Challenges. To understand the main challenges, let us briefly explain how the security proof goes through in the work of Garg and Srinivasan [GS18a]. In a typical construction of a garbled RAM program, using a sequence of garbled circuits, one would expect that the simulation of garbled circuits would be done from the first CPU step to the last CPU step. However, in [GS18a] proof, the simulation is done in a rather unusual manner, from the last CPU step to the first CPU step. Of course, it is not possible to simulate the last CPU step directly. Thus, the process of simulating the last CPU step itself involves a sequence of hybrids that simulate and “un-simulate” the garbling of the previous CPU steps. Extending this approach so that the memory contents and the access patterns are both hidden faces the following two main challenges.

- **Challenge 1:** In the Garg and Srinivasan construction [GS18a], memory contents were encrypted using one-time pads. Since the locations that each CPU step (for a circuit) reads from and write to are fixed, the one-time pad corresponding to that location could be hardwired to those CPU steps. On the other hand, in the case of RAM programs the locations being accessed are dynamically chosen and thus it is not possible to hard-wire the entire one-time pad into each CPU step as this would blow up the size of these CPU steps.

⁴A laconic OT scheme allows to compress a large database/memory to a small digest. The digest in some sense binds the entire database. In particular, given the digest there exists efficient algorithms that can read/update particular memory locations. The time taken by these algorithms grow only logarithmically with the size of the database.

It is instructive to note that encrypting the memory using an encryption scheme and decrypting the read memory contents does not suffice. See more on this in preliminary attempt below.

- **Challenge 2:** In the non-adaptive setting, it is easy to amplify unprotected memory access security to the setting where memory accesses are hidden using an oblivious RAM scheme [Gol87, Ost90, GO96]. However, in the adaptive setting this transformation turns out to be tricky. In a bit more detail, the Garg and Srinivasan [GS18a] approach of simulating CPU step circuits from the last to the first ends up in conflict with the security of the ORAM scheme where the simulation is typically done from the first to the last CPU steps. We note here that the techniques of Canetti et al. [CCHR16] and Ananth et al. [ACC⁺16], though useful, do not apply directly to our setting. In particular, in the Canetti et al. [CCHR16] and Ananth et al. [ACC⁺16] constructions, CPU steps were obfuscated using an indistinguishability obfuscation scheme. Thus, in their scheme the obfuscation for any individual CPU step could be changed independently. For example, the PRF key used in any CPU step could be punctured independent of the other CPU steps. On the other hand, in our construction, in spite of each CPU step being garbled separately, its input labels are hardwired in the previous garbled circuit. Therefore, a change in hardwired secret value (like a puncturing a key) in a CPU step needs an intricate sequence of hybrids for making this change. For instance, in the case of the example above, it is not possible to puncture the PRF key hardwired in a particular CPU step in one simple hybrid step. Instead any change in this CPU step must change the CPU step before it and so on. In summary, in our case, any such change would involve a new and intricate hybrid argument.

2.1 Solving Challenge 1

In this subsection, we describe our techniques to solve challenge 1.

Preliminary Attempt. A very natural approach to encrypting external memory would be to use a pseudorandom function to encrypt memory content in each location. More precisely, a data value d in location L is encrypted using the key $\text{PRF}_K(L)$ where K is the PRF key. The key K for this pseudorandom function is hardwired in each CPU step so that it first decrypts the ciphertext that is read from the memory and uses the underlying data for further processing. This approach to solving Challenge 1 was in fact used in the works of Canetti et al. [CCHR16] and Ananth et al. [ACC⁺16] (and several other prior works) in a similar context. However, in order to use the security of this PRF, we must first remove the hardwired key from each of the CPU steps. This is easily achieved if we rely on indistinguishability obfuscation. Indeed, a single hybrid change is sufficient to have the punctured key to be hardwired in each of the CPU steps. However, in our setting this does not work! In particular, we need to puncture the PRF key in each of the CPU step circuits by simulating them individually and the delicate dependencies involved in garbling each CPU step blows up the size of the garbled input to grow with the running time of the program.⁵ Due to the same reason, the approaches of encrypting the memory by maintaining a tree of secret keys [GLOS15, GLO15] do not work.

⁵For the readers who are familiar with [GS18a], the number of CPU steps that have to be maintained in the input dependent simulation for puncturing the PRF key grows with the number of CPU steps that last wrote to this location and this could be as large as the running time of the program.

Our New Idea: A Careful Timed Encryption Mechanism. From the above attempts, the following aspect of secure garbled RAM arise. Prior approaches for garbling RAM programs use PRF keys that in some sense “decrease in power”⁶ as hybrids steps involve sequential simulation of the CPU steps starting with the first CPU step and ending in the last CPU step. However, in the approach of [GS18a], the hybrids do a backward pass, from the last CPU step circuit to the first CPU step circuit. Therefore, we need a mechanism wherein the hardwired key for encryption in some sense “strengthens” along the first to the last CPU step.

Location vs. Time. In almost all garbled RAM constructions, the data stored at a particular location is encrypted using a location dependent key (e.g. [GLOS15]). This was not a problem when the keys are being weakened across CPU steps. However, in our case we need the key to be strengthened in power across CPU steps. Thus, we need a special purpose encryption scheme where the keys are derived based on time rather than the locations. Towards this goal, we construct a special purpose encryption scheme called as a *timed encryption* scheme. Let us explain this in more detail.

Timed Encryption. A timed encryption scheme is just like any (plain) symmetric key encryption except that every message is encrypted with respect to a timestamp. Additionally, there is a special key constrain algorithm that constrains a key to only decrypt ciphertexts that are encrypted within a specific timestamp. The security requirement is that the constrained key does not help in distinguishing ciphertexts of two messages that are encrypted with respect to some future timestamp. We additionally require the encryption using a key constrained with respect to a timestamp time to have the same distribution as an encryption using an unconstrained key as long as the timestamp to which we are encrypting is less than or equal to time. For efficiency, we require that the size of the constrained key to grow only with the length of the binary representation of the timestamp.

Solving Challenge 1. Timed encryption provides a natural approach to solving challenge 1. In every CPU step, we hardwire a time constrained key that allows that CPU step to decrypt all the memory updates done by the prior CPU steps. The last CPU step in some sense has the most powerful key hardwired, i.e., it can decrypt all the updates made by all the prior CPU steps and the first CPU step has the least powerful key hardwired. Thus, the hardwired secret key strengthens from the first CPU step to the last CPU step. In the security proof, a backward pass of simulating the last CPU step to the first CPU step conforms well with the semantics and security properties of a timed encryption scheme. This is because we remove the most powerful keys first and the rest of the hardwired secret keys in the previous CPU steps do not help in distinguishing between encryptions of the actual value that is written and some junk value. We believe that the notion timed encryption might have other applications and be of independent interest.

Constructing Timed Encryption. We give a construction of a timed encryption scheme from any one-way function. Towards this goal, we introduce a notion called as *range constrained PRF*. A range constrained PRF is a special constrained PRF [BW13] where the PRF key can be constrained to evaluate input points that fall within a particular range. The ranges that we will be interested in are of the form $[0, x]$. That is, the constrained key can be used to evaluate the PRF on any $y \in [0, x]$. For efficiency, we require that the size of the constrained key to only grow with the binary representation of x . Given such a PRF, we can construct a timed encryption scheme as follows. The key generation samples a range constrained PRF key. The encryption of a message m with respect

⁶The tree-based approaches of storing the secret keys use the mechanism wherein the hardwired secret keys decrease in power in subsequent CPU steps. In particular, the secret key corresponding to the root can decrypt all the locations, the secret keys corresponding to its children can only decrypt a part of the database and so on.

to a timestamp time proceeds by evaluating the PRF on time to derive sk and then using sk as a key for symmetric encryption scheme to encrypt the message m . The time constraining algorithm just constrains the PRF key with respect to the range $[0, \text{time}]$. Thus, the goal of constructing a timed encryption scheme reduces to the goal of constructing a range constrained PRF. In this work, we give a construction of range constrained PRF by adding a range constrain algorithm to the tree-based PRF scheme of Goldreich, Goldwasser and Micali [GGM86].

2.2 Solving Challenge 2

Challenge 1 involves protecting the contents of the memory whereas challenge 2 involves protecting the access pattern. As mentioned before, in the non-adaptive setting, this problem is easily solved using an oblivious RAM scheme. However, in our setting we need an oblivious RAM scheme with some special properties.

The works of Canetti et al. [CCHR16] and Ananth et al. [ACC⁺16] define a property of an ORAM scheme as *strong localized randomness property* and then use this property to hide their access patterns. Informally, an ORAM scheme is said to have a strong localized randomness property if the locations of the random tape accessed by an oblivious program in simulating each memory access are disjoint. Further, the number of locations touched for simulating each memory access must be poly logarithmic in the size of the database. These works further proved that the Chung-Pass ORAM scheme [CP13] satisfies the strong localized randomness property. Unfortunately, this strong localized randomness property alone is not sufficient for our purposes. Let us give the details.

To understand why the strong localized randomness property alone is not sufficient, we first recall the details of the Chung-Pass ORAM (henceforth, denoted as CP ORAM) scheme. The CP ORAM is a tree-based ORAM scheme where the leaves of this tree are associated with the actual memory. A position map associates each data block in the memory with a random leaf node. Accessing a memory location involves first reading the position map to get the address of the leaf where this data block resides. Then, the path from the root to this particular leaf is traversed and the content of the this data block is read. It is guaranteed that the data block is located somewhere along the path from the root to leaf node. The read data block is then placed in the root and the position map is updated so that another random leaf node is associated with this data block. To balance the memory, an additional flush is performed but for the sake of this introduction we ignore this step. The CP ORAM scheme has strong localized randomness as the randomness used in each memory accesses involves choosing a random leaf to update the position map. Let us now explain why this property alone is not sufficient for our purpose.

Recall that in the security proof of [GS18a], the CPU steps are simulated from the last step to the first. A simulation of a CPU step involves changing the bit written by the step to some junk value and the changing the location accessed to a random location. We can change the bit to be written to a junk value using the security of the timed encryption scheme, however changing the location accessed to random is problematic. Note that the location that is being accessed in the CP ORAM is a random root to leaf path. However, the address of this leaf is stored in the memory via the position map. Therefore, to simulate a particular CPU step, we must first change the contents of the position map. This change must be performed in those CPU steps that last updated this memory location. Unfortunately, timed encryption is not useful in this setting as we can use its security only after removing all the secret keys that are hardwired in the future time steps. However, in our case, the CPU steps that last updated this particular location might be so far into the past that removing all the intermediate encryption keys might blow up the cost of the

input encoding to be as large as the program running time.

To solve this issue, we modify the Chung-Pass ORAM to additionally have the CPU steps to encrypt the data block that is written using a puncturable PRF. Unlike the previous approaches of encrypting the data block with respect to the location, we encrypt it with respect to the time step that modifies the location. This helps in circumventing the above problem as we can first puncture the PRF key (which in turn involves a careful set of hybrids) and use its security to change the position map to contain an encryption of the junk value instead of the actual address of the leaf node.⁷ Once this change is done, the locations that the concerned CPU step is accessing is a random root to leaf path.

3 Preliminaries

Let λ denote the security parameter. A function $\mu(\cdot) : \mathbb{N} \rightarrow \mathbb{R}^+$ is said to be negligible if for any polynomial $\text{poly}(\cdot)$ there exists $\lambda_0 \in \mathbb{N}$ such that for all $\lambda > \lambda_0$ we have $\mu(\lambda) < \frac{1}{\text{poly}(\lambda)}$. For a probabilistic algorithm A , we denote $A(x; r)$ to be the output of A on input x with the content of the random tape being r . When r is omitted, $A(x)$ denotes a distribution. For a finite set S , we denote $x \leftarrow S$ as the process of sampling x uniformly from the set S . We will use PPT to denote Probabilistic Polynomial Time. We denote $[a]$ to be the set $\{1, \dots, a\}$ and $[a, b]$ to be the set $\{a, a + 1, \dots, b\}$ for $a \leq b$ and $a, b \in \mathbb{Z}$. For a binary string $x \in \{0, 1\}^n$, we will denote the i^{th} bit of x by x_i . We assume without loss of generality that the length of the random tape used by all cryptographic algorithms is λ . We will use $\text{negl}(\cdot)$ to denote an unspecified negligible function and $\text{poly}(\cdot)$ to denote an unspecified polynomial function.

3.1 Puncturable PRF

We recall the notion of puncturable PRF [GGM86, BW13, BGI14, KPTZ13, SW14].

Definition 3.1 *A puncturable pseudorandom function is a tuple of PPT algorithms (PP.KeyGen, PP.Eval, PP.Punc) with the following properties:*

- **Functionality is preserved under puncturing:** For all λ , for all $x \in \{0, 1\}^n$ (where n is the input length) and $\forall y \neq x$,

$$\Pr[\text{PP.Eval}(K[x], y) = \text{PP.Eval}(K, y)] = 1$$

where $K \leftarrow \text{PP.KeyGen}(1^\lambda)$ and $K[x] \leftarrow \text{PP.Punc}(K, x)$.

- **Pseudorandomness at punctured points:** For all λ , for all $x \in \{0, 1\}^n$, and for all poly sized adversaries \mathcal{A}

$$|\Pr[\mathcal{A}(K[x], \text{PP.Eval}(K, x)) = 1] - \Pr[\mathcal{A}(K[x], r) = 1]| \leq \text{negl}(\lambda)$$

where $K \leftarrow \text{PP.KeyGen}(1^\lambda)$ and $K[x] \leftarrow \text{PP.Punc}(K, x)$ and r is a random string having the same length as $\text{PP.Eval}(K, x)$.

Theorem 3.2 ([GGM86, BW13, BGI14, KPTZ13, SW14]) *Assuming the existence of one-way functions, there exists a construction of puncturable pseudorandom functions.*

⁷Unlike in the location based encryption scheme, it is sufficient to change the encryption only in the CPU steps that last modified this location.

3.2 Garbled Circuits

Below we recall the definition of garbling scheme for circuits [Yao82, AIK04, AIK05] with selective security (see Lindell and Pinkas [LP09] and Bellare et al. [BHR12b] for a detailed proof and further discussion). A garbling scheme for circuits is a tuple of PPT algorithms (GarbleCkt , EvalCkt). Very roughly, GarbleCkt is the circuit garbling procedure and EvalCkt the corresponding evaluation procedure. We use a formulation where input labels for a garbled circuit are provided as input to the garbling procedure rather than generated as output. (This simplifies the presentation of our construction.) More formally:

- $\tilde{C} \leftarrow \text{GarbleCkt}(1^\lambda, C, \{\text{lab}_{w,b}\}_{w \in n, b \in \{0,1\}})$: GarbleCkt takes as input a security parameter λ , a circuit C , and input labels $\text{lab}_{w,b}$ where $w \in n$ (n is the set of input wires to the circuit C) and $b \in \{0,1\}$. This procedure outputs a *garbled circuit* \tilde{C} . We assume that for each w, b , $\text{lab}_{w,b}$ is chosen uniformly from $\{0,1\}^\lambda$.
- $y \leftarrow \text{EvalCkt}(\tilde{C}, \{\text{lab}_{w,x_w}\}_{w \in n})$: Given a garbled circuit \tilde{C} and a sequence of input labels $\{\text{lab}_{w,x_w}\}_{w \in n}$ (referred to as the garbled input), EvalCkt outputs a string y .

Correctness. For correctness, we require that for any circuit C , input $x \in \{0,1\}^{|n|}$ and input labels $\{\text{lab}_{w,b}\}_{w \in n, b \in \{0,1\}}$ we have that:

$$\Pr \left[C(x) = \text{EvalCkt}(\tilde{C}, \{\text{lab}_{w,x_w}\}_{w \in n}) \right] = 1$$

where $\tilde{C} \leftarrow \text{GarbleCkt}(1^\lambda, C, \{\text{lab}_{w,b}\}_{w \in n, b \in \{0,1\}})$.

Selective Security. For security, we require that there exists a PPT simulator Sim_{Ckt} such that for any circuit C and input $x \in \{0,1\}^{|n|}$, we have that

$$\left\{ \tilde{C}, \{\text{lab}_{w,x_w}\}_{w \in n} \right\} \stackrel{c}{\approx} \left\{ \text{Sim}_{\text{Ckt}}(1^\lambda, 1^{|C|}, C(x), \{\text{lab}_{w,x_w}\}_{w \in n}), \{\text{lab}_{w,x_w}\}_{w \in n} \right\}$$

where $\tilde{C} \leftarrow \text{GarbleCkt}(1^\lambda, C, \{\text{lab}_{w,b}\}_{w \in n, b \in \{0,1\}})$ and for each $w \in n$ and $b \in \{0,1\}$ we have $\text{lab}_{w,b} \leftarrow \{0,1\}^\lambda$. Here $\stackrel{c}{\approx}$ denotes that the two distributions are computationally indistinguishable.

Theorem 3.3 ([Yao86, LP09]) *Assuming the existence of one-way functions, there exists a construction of garbling scheme for circuits.*

3.3 Range Constrained PRF

We now define a type of constrained PRF called as Range Constrained PRF. Informally, a range constrained PRF key allows to evaluate the PRF on any point in the domain $\{0,1\}^n$ that falls in a particular range. We will specifically be interested in ranges of the form $[0, T]$ for any $T \in [0, 2^n - 1]$. We now give the formal definition of this primitive.

Syntax. A Range Constrained PRF consists of the following PPT algorithms

- $\text{RC.KeyGen}(1^\lambda)$: It is a PPT algorithm that takes the security parameter (encoded in unary) as input and outputs a PRF key K . We implicitly assume that the key K defines an efficiently computable function $\text{PRF}_K : \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^\lambda$. For brevity, we will use n to denote $n(\lambda)$ respectively.
- $\text{RC.Constrain}(K, T)$: It is a deterministic algorithm that takes a PRF key K as input and a value $T \in [0, 2^n - 1]$ (encoded in binary) and outputs a constrained key $K[T]$.
- $\text{RC.Eval}(K[T], x)$: It is a deterministic algorithm that takes as input a constrained key $K[T]$ and an input x and outputs either a string y or \perp .

Correctness. We say that a range constrained PRF to be *correct* if for all K in the support of RC.KeyGen , $T \in [0, 2^n - 1]$ and $x \in [0, T]$ we have,

$$\text{RC.Eval}(K[T], x) = \text{PRF}_K(x)$$

where $K[T] := \text{RC.Constrain}(K, T)$.

Security. For security, we require that for any PPT adversaries \mathcal{A} , any $T_1, T_2, \dots, T_\ell \in [0, 2^n - 1]$ any $y \notin [0, T_i]$ for any $i \in [\ell]$,

$$|\Pr[\mathcal{A}(\{K[T_i]\}_{i \in [\ell]}, \text{PRF}_K(y)) = 1] - \Pr[\mathcal{A}(\{K[T_i]\}_{i \in [\ell]}, r) = 1]| \leq \text{negl}(\lambda)$$

where $K \leftarrow \text{RC.KeyGen}(1^\lambda)$, $K[T_i] := \text{RC.Constrain}(K, T_i)$ for all $i \in [\ell]$ and $r \leftarrow \{0, 1\}^\lambda$.

We give the proof of the following theorem in the Appendix A.

Theorem 3.4 *Assuming the existence of one-way functions, there exists a construction of range constrained PRFs.*

3.4 Updatable Laconic Oblivious Transfer

In this subsection, we recall the definition of updatable laconic oblivious transfer from [CDG⁺17].

We give the formal definition below from [CDG⁺17]. We generalize their definition to work for blocks of data instead of bits. More precisely, the reads and the updates happen at the block-level rather than at the bit-level.

Definition 3.5 ([CDG⁺17]) *An updatable laconic oblivious transfer consists of the following algorithms:*

- $\text{crs} \leftarrow \text{crsGen}(1^\lambda, 1^N)$: It takes as input the security parameter 1^λ (encoded in unary) and a block size N and outputs a common reference string crs .
- $(\text{d}, \widehat{D}) \leftarrow \text{Hash}(\text{crs}, D)$: It takes as input the common reference string crs and database $D \in \{\{0, 1\}^N\}^*$ as input and outputs a digest d and a state \widehat{D} . We assume that the state \widehat{D} also includes the database D .

- $e \leftarrow \text{Send}(\text{crs}, d, L, \{m_{i,0}, m_{i,1}\}_{i \in [N]})$: It takes as input the common reference string crs , a digest d , and a location $L \in \mathbb{N}$ and set of messages $m_{i,0}, m_{i,1} \in \{0, 1\}^{p(\lambda)}$ for every $i \in [N]$ and outputs a ciphertext e .
- $(m_1, \dots, m_N) \leftarrow \text{Receive}^{\widehat{D}}(\text{crs}, e, L)$: This is a RAM algorithm with random read access to \widehat{D} . It takes as input a common reference string crs , a ciphertext e , and a location $L \in \mathbb{N}$ and outputs a set of messages m_1, \dots, m_N .
- $e_w \leftarrow \text{SendWrite}(\text{crs}, d, L, \{b_i\}_{i \in [N]}, \{m_{j,0}, m_{j,1}\}_{j=1}^{|\mathbf{d}|})$: It takes as input the common reference string crs , a digest d , and a location $L \in \mathbb{N}$, bits $b_i \in \{0, 1\}$ for each $i \in [N]$ to be written, and $|\mathbf{d}|$ pairs of messages $\{m_{j,0}, m_{j,1}\}_{j=1}^{|\mathbf{d}|}$, where each $m_{j,c}$ is of length $p(\lambda)$ and outputs a ciphertext e_w .
- $\{m_j\}_{j=1}^{|\mathbf{d}|} \leftarrow \text{ReceiveWrite}^{\widehat{D}}(\text{crs}, L, \{b_i\}_{i \in [N]}, e_w)$: This is a RAM algorithm with random read/write access to \widehat{D} . It takes as input the common reference string crs , a location L , a set of bits $b_1, \dots, b_N \in \{0, 1\}$ and a ciphertext e_w . It updates the state \widehat{D} (such that $D[L] = b_1 \dots b_N$) and outputs messages $\{m_j\}_{j=1}^{|\mathbf{d}|}$.

We require an updatable laconic oblivious transfer to satisfy the following properties.

Correctness: We require that for any database D of size at most $M = \text{poly}(\lambda)$, any memory location $L \in [M]$, any set of messages $(m_{i,0}, m_{i,1}) \in \{0, 1\}^{p(\lambda)}$ for each $i \in [N]$ where $p(\cdot)$ is a polynomial that

$$\Pr \left[\forall i \in [N], m_i = m_{i,D[L,i]} \mid \begin{array}{l} \text{crs} \leftarrow \text{crsGen}(1^\lambda) \\ (\mathbf{d}, \widehat{D}) \leftarrow \text{Hash}(\text{crs}, D) \\ e \leftarrow \text{Send}(\text{crs}, d, L, \{m_{i,0}, m_{i,1}\}_{i \in [N]}) \\ (m_1, \dots, m_N) \leftarrow \text{Receive}^{\widehat{D}}(\text{crs}, e, L) \end{array} \right] = 1,$$

where $D[L, i]$ denotes the i^{th} bit in the L^{th} block of D .

Correctness of Writes: Let database D be of size at most $M = \text{poly}(\lambda)$ and let $L \in [M]$ be any two memory locations. Let D^* be a database that is identical to D except that $D^*[L, i] = b_i$ for all $i \in [N]$ some sequence of $\{b_j\} \in \{0, 1\}$. For any sequence of messages $\{m_{j,0}, m_{j,1}\}_{j \in [N]} \in \{0, 1\}^{p(\lambda)}$ we require that

$$\Pr \left[\begin{array}{l} m'_j = m_{j,d'_j} \\ \forall j \in [|\mathbf{d}|] \end{array} \mid \begin{array}{l} \text{crs} \leftarrow \text{crsGen}(1^\lambda, 1^N) \\ (\mathbf{d}, \widehat{D}) \leftarrow \text{Hash}(\text{crs}, D) \\ (\mathbf{d}^*, \widehat{D}^*) \leftarrow \text{Hash}(\text{crs}, D^*) \\ e_w \leftarrow \text{SendWrite}(\text{crs}, d, L, \{b_i\}_{i \in [N]}, \{m_{j,0}, m_{j,1}\}_{j=1}^{|\mathbf{d}|}) \\ \{m'_j\}_{j=1}^{|\mathbf{d}|} \leftarrow \text{ReceiveWrite}^{\widehat{D}}(\text{crs}, L, \{b_i\}_{i \in [N]}, e_w) \end{array} \right] = 1,$$

Sender Privacy: There exists a PPT simulator Sim_{OT} such that for any non-uniform PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function $\text{negl}(\cdot)$ s.t.,

$$\left| \Pr[\text{Expt}^{\text{real}}(1^\lambda, \mathcal{A}) = 1] - \Pr[\text{Expt}^{\text{ideal}}(1^\lambda, \mathcal{A}) = 1] \right| \leq \text{negl}(\lambda)$$

where $\text{Expt}^{\text{real}}$ and $\text{Expt}^{\text{ideal}}$ are described in Figure 1.

| $\text{Expt}^{\text{real}}[1^\lambda, \mathcal{A}]$ | $\text{Expt}^{\text{ideal}}[1^\lambda, \mathcal{A}]$ |
|---|--|
| <ol style="list-style-type: none"> 1. $\text{crs} \leftarrow \text{crsGen}(1^\lambda, 1^N)$. 2. $(D, L, \{m_{i,0}, m_{i,1}\}_{i \in [N]}, \text{st}) \leftarrow \mathcal{A}_1(\text{crs})$. 3. $(\text{d}, \widehat{D}) \leftarrow \text{Hash}(\text{crs}, D)$. 4. Output $\mathcal{A}_2(\text{st}, \text{Send}(\text{crs}, \text{d}, L, \{m_{i,0}, m_{i,1}\}_{i \in [N]}))$. | <ol style="list-style-type: none"> 1. $\text{crs} \leftarrow \text{crsGen}(1^\lambda)$. 2. $(D, L, \{m_{i,0}, m_{i,1}\}_{i \in [N]}, \text{st}) \leftarrow \mathcal{A}_1(\text{crs})$. 3. $(\text{d}, \widehat{D}) \leftarrow \text{Hash}(\text{crs}, D)$. 4. Output $\mathcal{A}_2(\text{st}, \text{Sim}_{\ell\text{OT}}(\text{crs}, D, L, \{m_{i,D[L,i]}\}_{i \in [N]}))$. |

Figure 1: Sender Privacy Security Game

Sender Privacy for Writes: *There exists a PPT simulator $\text{Sim}_{\ell\text{OTW}}$ such that for any non-uniform PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ there exists a negligible function $\text{negl}(\cdot)$ s.t.,*

$$\left| \Pr[\text{WriSenPrivExpt}^{\text{real}}(1^\lambda, \mathcal{A}) = 1] - \Pr[\text{WriSenPrivExpt}^{\text{ideal}}(1^\lambda, \mathcal{A}) = 1] \right| \leq \text{negl}(\lambda)$$

where $\text{WriSenPrivExpt}^{\text{real}}$ and $\text{WriSenPrivExpt}^{\text{ideal}}$ are described in Figure 2.

| $\text{WriSenPrivExpt}^{\text{real}}[1^\lambda, \mathcal{A}]$ | $\text{WriSenPrivExpt}^{\text{ideal}}[1^\lambda, \mathcal{A}]$ |
|--|--|
| <ol style="list-style-type: none"> 1. $\text{crs} \leftarrow \text{crsGen}(1^\lambda, 1^N)$. 2. $(D, L, \{b_i\}_{i \in [N]}, \{m_{j,0}, m_{j,1}\}_{j \in [\lambda]}, \text{st}) \leftarrow \mathcal{A}_1(\text{crs})$. 3. $(\text{d}, \widehat{D}) \leftarrow \text{Hash}(\text{crs}, D)$. 4. $e_w \leftarrow \text{SendWrite}(\text{crs}, \text{d}, L, \{b_i\}_{i \in [N]}, \{m_{j,0}, m_{j,1}\}_{j=1}^{ \text{d} })$. 5. Output $\mathcal{A}_2(\text{st}, e_w)$. | <ol style="list-style-type: none"> 1. $\text{crs} \leftarrow \text{crsGen}(1^\lambda, 1^N)$. 2. $(D, L, \{b_i\}_{i \in [N]}, \{m_{j,0}, m_{j,1}\}_{j \in [\lambda]}, \text{st}) \leftarrow \mathcal{A}_1(\text{crs})$. 3. $(\text{d}, \widehat{D}) \leftarrow \text{Hash}(\text{crs}, D)$. 4. $(\text{d}^*, \widehat{D}^*) \leftarrow \text{Hash}(\text{crs}, D^*)$ where D^* be a database that is identical to D except that $D^*[L, i] = b_i$ for each $i \in [N]$. 5. $e_w \leftarrow \text{Sim}_{\ell\text{OTW}}(\text{crs}, D, L, \{b_i\}_{i \in [N]}, \{m_{j,d_j^*}\}_{j \in [\lambda]})$. 6. Output $\mathcal{A}_2(\text{st}, e_w)$. |

Figure 2: Sender Privacy for Writes Security Game

Efficiency: *The algorithm Hash runs in time $|D|\text{poly}(\log |D|, \lambda)$. The algorithms Send, SendWrite, Receive, ReceiveWrite run in time $N \cdot \text{poly}(\log |D|, \lambda)$.*

Theorem 3.6 ([CDG⁺17, DG17, BLSV18, DGHM18]) *Assuming either the Computational Diffie-Hellman assumption or the Factoring assumption or the Learning with Errors assumption, there exists a construction of updatable laconic oblivious transfer.*

Remark 3.7 *We note that the security requirements given in Definition 3.5 is stronger than the one in [CDG⁺17] as we require the crs to be generated before the adversary provides the database D and the location L . However, the construction in [CDG⁺17] already satisfies this definition since in the proof, we can guess the location by incurring a $1/|D|$ loss in the security reduction.*

3.5 Somewhere Equivocal Encryption

We now recall the definition of Somewhere Equivocal Encryption from the work of [HJO⁺16]. Informally, a somewhere equivocal encryption allows to create a simulated ciphertext encrypting a message m with certain positions of the message being “fixed” and the other positions having a “hole.” The simulator can later fill these “holes” with arbitrary message values by deriving a suitable decryption key. The main efficiency requirement is that the size of the decryption key grows only with the number of “holes” and is otherwise independent of the message size. We give the formal definition below.

Definition 3.8 ([HJO⁺16]) *A somewhere equivocal encryption scheme with block-length s , message length n (in blocks) and equivocation parameter t (all polynomials in the security parameter) is a tuple of probabilistic polynomial algorithms $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{SimEnc}, \text{SimKey})$ such that:*

- $\text{key} \leftarrow \text{KeyGen}(1^\lambda)$: *It is a PPT algorithm that takes as input the security parameter (encoded in unary) and outputs a key key .*
- $\bar{c} \leftarrow \text{Enc}(\text{key}, m_1 \dots m_n)$: *It is a PPT algorithm that takes as input a key key and a vector of messages $\bar{m} = m_1 \dots m_n$ with each $m_i \in \{0, 1\}^s$ and outputs a ciphertext \bar{c} .*
- $\bar{m} \leftarrow \text{Dec}(\text{key}, \bar{c})$: *It is a deterministic algorithm that takes as input a key key and a ciphertext \bar{c} and outputs a vector of messages $\bar{m} = m_1 \dots m_n$.*
- $(\text{st}, \bar{c}) \leftarrow \text{SimEnc}((m_i)_{i \notin I}, I)$: *It is a PPT algorithm that takes as input a set of indices $I \subseteq [n]$ and a vector of messages $(m_i)_{i \notin I}$ and outputs a ciphertext \bar{c} and a state st .*
- $\text{key}' \leftarrow \text{SimKey}(\text{st}, (m_i)_{i \in I})$: *It is a PPT algorithm that takes as input the state information st and a vector of messages $(m_i)_{i \in I}$ and outputs a key key' .*

and satisfies the following properties:

Correctness. *For every $\text{key} \leftarrow \text{KeyGen}(1^\lambda)$, for every $\bar{m} \in \{0, 1\}^{s \times n}$ it holds that:*

$$\text{Dec}(\text{key}, \text{Enc}(\text{key}, \bar{m})) = \bar{m}$$

Simulation with No Holes. *We require that the distribution of (\bar{c}, key) computed via $(\text{st}, \bar{c}) \leftarrow \text{SimEnc}(\bar{m}, \emptyset)$ and $\text{key} \leftarrow \text{SimKey}(\text{st}, \emptyset)$ to be identical to $\text{key} \leftarrow \text{KeyGen}(1^\lambda)$ and $\bar{c} \leftarrow \text{Enc}(\text{key}, m_1 \dots m_n)$. In other words, simulation when there are no holes (i.e., $I = \emptyset$) is identical to honest key generation and encryption.*

Security. *For any PPT adversary \mathcal{A} , there exists a negligible function $\nu = \nu(\lambda)$ such that:*

$$\left| \Pr[\text{Exp}_{\mathcal{A}, \Pi}^{\text{simenc}}(1^\lambda, 0) = 1] - \Pr[\text{Exp}_{\mathcal{A}, \Pi}^{\text{simenc}}(1^\lambda, 1) = 1] \right| \leq \nu(\lambda)$$

where the experiment $\text{Exp}_{\mathcal{A}, \Pi}^{\text{simenc}}$ is defined as follows:

Experiment $\text{Exp}_{\mathcal{A}, \Pi}^{\text{simenc}}$

1. *The adversary \mathcal{A} on input 1^λ outputs a set $I \subseteq [n]$ s.t. $|I| < t$, a vector $(m_i)_{i \notin I}$, and a challenge $j \in [n] \setminus I$. Let $I' = I \cup \{j\}$.*

2.
 - If $b = 0$, compute \bar{c} as follows: $(\text{st}, \bar{c}) \leftarrow \text{SimEnc}((m_i)_{i \notin I}, I)$.
 - If $b = 1$, compute \bar{c} as follows: $(\text{st}, \bar{c}) \leftarrow \text{SimEnc}((m_i)_{i \notin I'}, I')$.
3. Send \bar{c} to the adversary \mathcal{A} .
4. The adversary \mathcal{A} outputs the set of remaining messages $(m_i)_{i \in I}$.
 - If $b = 0$, compute key as follows: $\text{key} \leftarrow \text{SimKey}(\text{st}, (m_i)_{i \in I})$.
 - If $b = 1$, compute key as follows: $\text{key} \leftarrow \text{SimKey}(\text{st}, (m_i)_{i \in I'})$.
5. Send key to the adversary.
6. \mathcal{A} outputs b' which is the output of the experiment.

Theorem 3.9 ([HJO⁺16]) *Assuming the existence of one-way functions, there exists a somewhere equivocal encryption scheme for any polynomial message-length n , black-length s and equivocation parameter t , having key size $t \cdot s \cdot \text{poly}(\lambda)$ and ciphertext of size $n \cdot s \cdot \text{poly}(\lambda)$ bits.*

3.6 Random Access Machine (RAM) Model of Computation

We start by describing the Random Access Machine (RAM) model of computation in Section 3.6. Most of this subsection is taken verbatim from [CDG⁺17].

Notation for the RAM Model of Computation. The RAM model consists of a CPU and a memory storage of M blocks where each block has length N . 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 $D \in \{\{0, 1\}^N\}^M$. Additionally, the program gets a “short” input $x \in \{0, 1\}^n$, which we alternatively think of as the initial state of the program. We use $|P|$ to denote the running time of program P . 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.⁸

We will also consider the case where several different programs are executed sequentially and the memory persists between executions. We denote this process as $(y_1, \dots, y_\ell) = (P_1(x_1), \dots, P_\ell(x_\ell))^D$ to indicate that first $P_1^D(x_1)$ is executed, resulting in some memory contents D_1 and output y_1 , then $P_2^{D_1}(x_2)$ is executed resulting in some memory contents D_2 and output y_2 etc. As an example, 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:⁹

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

⁸In 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 x is small.

⁹In the definition below, we model each C_{CPU} as a deterministic circuit. Later, we extend the definition to allow each C_{CPU} to have access to random coins.

This circuit takes as input the current CPU state state and $\text{rData} \in \{0, 1\}^N$. Looking ahead the data 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 R/W , the next location to read/write from $L \in [M]$, and data wData to write into that location ($\text{wData} = \perp$ when reading). The sequence of locations accessed during the execution of the program collectively form what is known as the *access pattern*, namely $\text{MemAccess} = \{(\text{R/W}^\tau, L^\tau) : \tau = 1, \dots, T\}$. We assume that the CPU state state contains information about the location that the previous CPU step requested to read from. In particular, $\text{lastLocation}(\text{state})$ outputs the location that the previous CPU step requested to read and it is \perp if the previous CPU step was a write.

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 writes to 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 $\text{state}_1 = x$. In each step $\tau \in \{1, \dots, T\}$, the computation proceeds as follows: If $\tau = 1$ or $\text{R/W}^{\tau-1} = \text{write}$, then $\text{rData}^\tau := \perp$; otherwise $\text{rData}^\tau := D[L^{\tau-1}]$. Next it executes the CPU-Step Circuit $C_{\text{CPU}}^{P, \tau}(\text{state}^\tau, \text{rData}^\tau) = (\text{state}^{\tau+1}, \text{R/W}^\tau, L^\tau, \text{wData}^\tau)$. If $\text{R/W}^\tau = \text{write}$, then set $D[L^\tau] = \text{wData}^\tau$. Finally, when $\tau = T$, then $\text{state}^{\tau+1}$ is the output of the program.

3.7 Oblivious RAM

In this subsection, we recall the definition of oblivious RAM [Gol87, Ost90, GO96].

Definition 3.10 (Oblivious RAM) *An Oblivious RAM scheme consists of two procedures (OProg, OData) with the following syntax:*

- $P^* \leftarrow \text{OProg}(1^\lambda, 1^{\log M}, 1^T, P)$: *Given a security parameter λ , a memory size M , a program P that runs in time T , OProg outputs an probabilistic oblivious program P^* that can access D^* as RAM. A probabilistic RAM program is modeled exactly as a deterministic program except that each step circuit additionally take random coins as input.*
- $D^* \leftarrow \text{OData}(1^\lambda, D)$: *Given the security parameter λ , the contents of the database $D \in \{\{0, 1\}^N\}^M$, outputs the oblivious database D^* . For convenience, we assume that OData works by compiling a program P that writes D to the memory using OProg to obtain P^* . It then evaluates the program P^* by using uniform random tape and outputs the contents of the memory as D^* .*

Efficiency. *We require that the run-time of OData should be $M \cdot N \cdot \text{poly}(\log(MN)) \cdot \text{poly}(\lambda)$, and the run-time of OProg should be $T \cdot \text{poly}(\lambda) \cdot \text{poly}(\log(MN))$. Finally, the oblivious program P^* itself should run in time $T' = T \cdot \text{poly}(\lambda) \cdot \text{poly}(\log(MN))$. Both the new memory size $M' = |D^*|$ and the running time T' should be efficiently computable from M, N, T , and λ .*

Correctness. *Let P_1, \dots, P_ℓ be programs running in polynomial times t_1, \dots, t_ℓ on memory D of size M . Let x_1, \dots, x_ℓ be the inputs and λ be a security parameter. Then we require that:*

$$\Pr[(P_1^*(x_1), \dots, P_\ell^*(x_\ell))^{D^*} = (P_1(x_1), \dots, P_\ell(x_\ell))^D] = 1$$

where $D^* \leftarrow \text{OData}(1^\lambda, D)$, $P_i^* \leftarrow \text{OProg}(1^\lambda, 1^{\log M}, 1^T, P_i)$ and $(P_1^*(x_1), \dots, P_\ell^*(x_\ell))^{D^*}$ indicates running the ORAM programs on D^* sequentially using an uniform random tape.

Security. For security, we require that there exists a PPT simulator Sim such that for any sequence of programs P_1, \dots, P_ℓ (running in time t_1, \dots, t_ℓ respectively), initial memory data $D \in \{\{0, 1\}^N\}^M$, and inputs x_1, \dots, x_ℓ we have that:

$$\text{MemAccess} \stackrel{s}{\approx} \text{Sim}(1^\lambda, \{1^{t_i}\}_{i=1}^\ell)$$

where $(y_1, \dots, y_\ell) = (P_1(x_1), \dots, P_\ell(x_\ell))^{D^*}$, $D^* \leftarrow \text{OData}(1^\lambda, 1^N, D)$, $P_i^* \leftarrow \text{OProg}(1^\lambda, 1^{\log M}, 1^T, P_i)$ and MemAccess corresponds to the access pattern of the CPU-step circuits during the sequential execution of the oblivious programs $(P_1^*(x_1), \dots, P_\ell^*(x_\ell))^{D^*}$ using an uniform random tape.

3.7.1 Strong Localized Randomness

For our construction of adaptively secure garbled RAM, we need an additional property called as strong localized randomness property [CCHR16] from an ORAM scheme. We need a slightly stronger formalization than the one given in [CCHR16] (refer to footnote 10).

Strong Localized Randomness. Let $D \in \{\{0, 1\}^N\}^M$ be any database and (P, x) be any program/input pair. Let $D^* \leftarrow \text{OData}(1^\lambda, 1^N, D)$ and $P^* \leftarrow \text{OProg}(1^\lambda, 1^{\log M}, 1^T, P)$. Further, let the step circuits of P^* be indicated by $\{C_{\text{CPU}}^{P^*, \tau}\}_{\tau \in [T']}$. Let R be the contents of the random tape used in the execution of P^* .

Definition 3.11 ([CCHR16]) *We say that an ORAM scheme has strong localized randomness property if there exists a sequence of efficiently computable values $\tau_1 < \tau_2 < \dots < \tau_m$ where $\tau_1 = 1$, $\tau_m = T'$ and $\tau_t - \tau_{t-1} \leq \text{poly}(\log MN)$ for all $t \in [2, m]$ such that:*

1. For every $j \in [m - 1]$ there exists an interval I_j (efficiently computable from j) of size $\text{poly}(\log MN, \lambda)$ s.t. for any $\tau \in [\tau_j, \tau_{j+1})$, the random tape accessed by $C_{\text{CPU}}^{P^*, \tau}$ is given by R_{I_j} (here, R_{I_j} denotes the random tape restricted to the interval I_j).
2. For every $j, j' \in [m - 1]$ and $j \neq j'$, $I_j \cap I_{j'} = \emptyset$.
3. Further, for every $j \in [m]$, there exists an $k < j$ such that given $R_{\setminus \{I_k \cup I_j\}}$ (where $R_{\setminus \{I_k \cup I_j\}}$ denotes the content of the random tape except in positions $I_j \cup I_k$) and the output of step circuits $C_{\text{CPU}}^{P^*, \tau}$ for $\tau \in [\tau_k, \tau_{k+1})$, the memory access made by step circuits $C_{\text{CPU}}^{P^*, \tau}$ for $\tau \in [\tau_j, \tau_{j+1})$ is computationally indistinguishable to random. This k is efficiently computable given the program P and the input x .¹⁰

In Appendix B, we show that the Chung-Pass ORAM scheme [CP13] where the contents of the database are encrypted using a special encryption scheme satisfies the above definition of strong localized randomness. We now give details on this special encryption scheme. The key generation samples a puncturable PRF key $K \leftarrow \text{PP.KeyGen}(1^\lambda)$. If the τ^{th} step-circuit has to write a value $w\text{Data}$ to a location L , it first samples $r \leftarrow \{0, 1\}^\lambda$ and computes $c = (\tau \parallel r, \text{PP.Eval}(K, \tau \parallel r) \oplus w\text{Data})$.

¹⁰Here, we require that the memory access to be indistinguishable to random even given the outputs of the step circuits $C_{\text{CPU}}^{P^*, \tau}$ for $\tau \in [\tau_k, \tau_{k+1})$. This is where we differ from the definition of [CCHR16].

It writes c to location L . The decryption algorithm uses K to first compute $\text{PP.Eval}(K, \tau || r)$ and uses it to compute $w\text{Data}$.

Remark 3.12 *For the syntax of the ORAM scheme to be consistent with this special encryption scheme, we will use a puncturable PRF to generate the random tape of P^* . This key will also be used implicitly to derive the key for this special encryption scheme.*

3.8 Adaptive Garbled RAM

We now give the definition of adaptive garbled RAM.

Definition 3.13 *An adaptive garbled RAM scheme GRAM consists of the following PPT algorithms satisfying the correctness, efficiency and security properties.*

- $\text{GRAM.Memory}(1^\lambda, D)$: It is a PPT algorithm that takes the security parameter 1^λ and a database $D \in \{0, 1\}^M$ as input and outputs a garbled database \tilde{D} and a secret key SK .
- $\text{GRAM.Program}(SK, i, P)$: It is a PPT algorithm that takes as input a secret key SK , a sequence number i , and a program P as input (represented as a sequence of CPU steps) and outputs a garbled program \tilde{P} .
- $\text{GRAM.Input}(SK, i, x)$: It is a PPT algorithm that takes as input a secret key SK , a sequence number i and a string x as input and outputs the garbled input \tilde{x} .
- $\text{GRAM.Eval}^{\tilde{D}}(\text{st}, \tilde{P}, \tilde{x})$: It is a RAM program with random read write access to \tilde{D} . It takes the state information st , garbled program \tilde{P} and the garbled input \tilde{x} as input and outputs a string y and updated database \tilde{D}' .

Correctness. We say that a garbled RAM GRAM is correct if for every database D , $t = \text{poly}(\lambda)$ and every sequence of program and input pair $\{(P_1, x_1), \dots, (P_t, x_t)\}$ we have that

$$\Pr[\text{Expt}_{\text{correctness}}(1^\lambda, \text{GRAM}) = 1] \leq \text{negl}(\lambda)$$

where $\text{Expt}_{\text{correctness}}$ is defined in Figure 5.

Adaptive Security. We say that GRAM satisfies adaptive security if there exists (stateful) simulators $(\text{SimD}, \text{SimP}, \text{SimIn})$ such that for all t that is polynomial in the security parameter λ and for all polynomial time (stateful) adversaries \mathcal{A} , we have that

$$\left| \Pr[\text{Expt}_{\text{real}}(1^\lambda, \text{GRAM}, \mathcal{A}) = 1] - \Pr[\text{Expt}_{\text{ideal}}(1^\lambda, \text{Sim}, \mathcal{A}) = 1] \right| \leq \text{negl}$$

where $\text{Expt}_{\text{real}}, \text{Expt}_{\text{ideal}}$ are defined in Figure 4.

- $(\tilde{D}, SK) \leftarrow \text{GRAM.Memory}(1^\lambda, D)$.
- Set $D_1 := D$, $\tilde{D}_1 := \tilde{D}$ and $\text{st} = \perp$.
- **for** every i from 1 to t
 - $\tilde{P}_i \leftarrow \text{GRAM.Program}(SK, i, P_i)$.
 - $\tilde{x}_i \leftarrow \text{GRAM.Input}(SK, i, x_i)$.
 - Compute $(y_i, D_{i+1}) := P_i^{D_i}(x_i)$ and $(\tilde{y}_i, \tilde{D}_{i+1}, \text{st}) := \text{UGRAM.Eval}^{\tilde{D}_i}(i, \text{st}, \tilde{P}_i, \tilde{x}_i)$.
- Output 1 if there exists an $i \in [t]$ such that $\tilde{y}_i \neq y_i$.

Figure 3: Correctness Experiment for GRAM

| $\text{Expt}_{\text{real}}[1^\lambda, \text{GRAM}, \mathcal{A}]$ | $\text{Expt}_{\text{ideal}}[1^\lambda, \text{Sim}, \mathcal{A}]$ |
|---|---|
| <ul style="list-style-type: none"> • $D \leftarrow \mathcal{A}(1^\lambda)$ where $D \in \{0, 1\}^M$ • $(\tilde{D}, SK) \leftarrow \text{GRAM.Memory}(1^\lambda, D)$. • for every i from 1 to t <ul style="list-style-type: none"> – $P_i \leftarrow \mathcal{A}(\tilde{D}, \{(\tilde{P}_1, \tilde{x}_1), \dots, (\tilde{P}_{i-1}, \tilde{x}_{i-1})\})$. – $\tilde{P}_i \leftarrow \text{GRAM.Program}(SK, i, P_i)$. – $x_i \leftarrow \mathcal{A}(\tilde{D}, \{(\tilde{P}_1, \tilde{x}_1), \dots, (\tilde{P}_{i-1}, \tilde{x}_{i-1})\}, \tilde{P}_i)$. – $\tilde{x}_i \leftarrow \text{GRAM.Input}(SK, i, x_i)$. • Output $\mathcal{A}(\{(\tilde{P}_1, \tilde{x}_1), \dots, (\tilde{P}_t, \tilde{x}_t)\})$. | <ul style="list-style-type: none"> • $D \leftarrow \mathcal{A}(1^\lambda)$ where $D \in \{0, 1\}^M$. • $(\tilde{D}, \text{st}) \leftarrow \text{SimD}(1^\lambda, 1^M)$. • for every i from 1 to t <ul style="list-style-type: none"> – $P_i \leftarrow \mathcal{A}(\tilde{D}, \{(\tilde{P}_1, \tilde{x}_1), \dots, (\tilde{P}_{i-1}, \tilde{x}_{i-1})\})$. – $(\tilde{P}_i, \text{st}) \leftarrow \text{SimP}(1^{ P_i }, \text{st})$. – $x_i \leftarrow \mathcal{A}(\tilde{D}, \{(\tilde{P}_1, \tilde{x}_1), \dots, (\tilde{P}_{i-1}, \tilde{x}_{i-1})\}, \tilde{P}_i)$. – $(y_i, D_{i+1}) := P_i^{D_i}(x_i)$ where $D_1 := D$. – $\tilde{x}_i \leftarrow \text{SimIn}(\text{st}, y_i)$. • Output $\mathcal{A}(\{(\tilde{P}_1, \tilde{x}_1), \dots, (\tilde{P}_t, \tilde{x}_t)\})$. |

Figure 4: Adaptive Security Experiment for GRAM

Efficiency. We require the following efficiency properties from a UGRAM scheme.

- The running time of GRAM.Memory should be bounded by $M \cdot \text{poly}(\log M) \cdot \text{poly}(\lambda)$.
- The running time of GRAM.Program should be bounded by $T \cdot \text{poly}(\log M) \cdot \text{poly}(\lambda)$ where T is the number of CPU steps in the description of the program P .
- The running time of GRAM.Input should be bounded by $|x| \cdot \text{poly}(\log M, \log T) \cdot \text{poly}(\lambda)$.
- The running time of GRAM.Eval should be bounded by $T \cdot \text{poly}(\log M) \cdot \text{poly}(\lambda)$ where T is the number of CPU steps in the description of the program P .

4 Adaptive Garbled RAM with Unprotected Memory Access

Towards our goal of constructing an adaptive garbled RAM, we first construct an intermediate primitive with weaker security guarantees. We call this primitive as *adaptive garbled RAM with unprotected memory access*. Informally, a garbled RAM scheme has unprotected memory access if both the contents of the database and the access to the database are revealed in the clear to the adversary. We differ from the security definition given in [GHL⁺14] in three aspects. Firstly, we give an indistinguishability style definition for security whereas [GHL⁺14] give a simulation style definition. The indistinguishability based definition makes it easier to get full-fledged adaptive security later. Secondly and most importantly, we allow the adversary to adaptively choose the inputs based on the garbled program. Thirdly, we also require the garbled RAM scheme to satisfy a special property called as *equivocability*. Informally, equivocability requires that the real garbling of a program P is indistinguishable to a simulated garbling where the simulator is not provided with the description of the step circuits for a certain number of time steps (this number is given by the equivocation parameter). Later, when the input is specified, the simulator is given the output of these step circuits and must come-up with an appropriate garbled input.

We now give the formal definition of this primitive.

Definition 4.1 *An adaptive garbled RAM scheme with unprotected memory access UGRAM consists of the following PPT algorithms satisfying the correctness, efficiency and security properties.*

- $\text{UGRAM.Memory}(1^\lambda, 1^n, D)$: It is a PPT algorithm that takes the security parameter 1^λ , an equivocation parameter n and a database $D \in \{\{0, 1\}^N\}^M$ as input and outputs a garbled database \tilde{D} and a secret key SK .
- $\text{UGRAM.Program}(SK, i, P)$: It is a PPT algorithm that takes as input a secret key SK , a sequence number i , and a program P as input (represented as a sequence of CPU steps) and outputs a garbled program \tilde{P} .
- $\text{UGRAM.Input}(SK, i, x)$: It is a PPT algorithm that takes as input a secret key SK , a sequence number i and a string x as input and outputs the garbled input \tilde{x} .
- $\text{UGRAM.Eval}^{\tilde{D}}(\text{st}, \tilde{P}, \tilde{x})$: It is a RAM program with random read write access to \tilde{D} . It takes the state information st , garbled program \tilde{P} and the garbled input \tilde{x} as input and outputs a string y and updated database \tilde{D}' .

Correctness. We say that a garbled RAM UGRAM is correct if for every database D , $t = \text{poly}(\lambda)$ and every sequence of program and input pair $\{(P_1, x_1), \dots, (P_t, x_t)\}$ we have that

$$\Pr[\text{Expt}_{\text{correctness}}(1^\lambda, \text{UGRAM}) = 1] \leq \text{negl}(\lambda)$$

where $\text{Expt}_{\text{correctness}}$ is defined in Figure 5.

Security. We require the following two properties to hold.

- **Equivocability.** There exists a simulator Sim such that for any non-uniform PPT stateful adversary \mathcal{A} and $t = \text{poly}(\lambda)$ we require that:

$$\left| \Pr[\text{Expt}_{\text{equiv}}(1^\lambda, \mathcal{A}, 0) = 1] - \Pr[\text{Expt}_{\text{equiv}}(1^\lambda, \mathcal{A}, 1) = 1] \right| \leq \text{negl}(\lambda)$$

- $(\tilde{D}, SK) \leftarrow \text{UGRAM.Memory}(1^\lambda, 1^n, D)$.
- Set $D_1 := D$, $\tilde{D}_1 := \tilde{D}$ and $\text{st} = \perp$.
- for every i from 1 to t
 - $\tilde{P}_i \leftarrow \text{UGRAM.Program}(SK, i, P_i)$.
 - $\tilde{x}_i \leftarrow \text{UGRAM.Input}(SK, i, x_i)$.
 - Compute $(y_i, D_{i+1}) := P_i^{D_i}(x_i)$ and $(\tilde{y}_i, \tilde{D}_{i+1}, \text{st}) := \text{UGRAM.Eval}^{\tilde{D}_i}(i, \text{st}, \tilde{P}_i, \tilde{x}_i)$.
- Output 1 if there exists an $i \in [t]$ such that $\tilde{y}_i \neq y_i$.

Figure 5: Correctness Experiment for UGRAM

where $\text{Expt}_{\text{equiv}}(1^\lambda, \mathcal{A}, b)$ is described in Figure 6.

- **Adaptive Security.** For any non-uniform PPT stateful adversary \mathcal{A} and $t = \text{poly}(\lambda)$ we require that:

$$\left| \Pr[\text{Expt}_{\text{UGRAM}}(1^\lambda, \mathcal{A}, 0) = 1] - \Pr[\text{Expt}_{\text{UGRAM}}(1^\lambda, \mathcal{A}, 1) = 1] \right| \leq \text{negl}(\lambda)$$

where $\text{Expt}_{\text{UGRAM}}(1^\lambda, \mathcal{A}, b)$ is described in Figure 7.

Efficiency. We require the following efficiency properties from a UGRAM scheme.

- The running time of UGRAM.Memory should be bounded by $MN \cdot \text{poly}(\log MN) \cdot \text{poly}(\lambda)$.
- The running time of UGRAM.Program should be bounded by $T \cdot \text{poly}(\log MN) \cdot \text{poly}(\lambda)$ where T is the number of CPU steps in the description of the program P .
- The running time of UGRAM.Input should be bounded by $n \cdot |x| \cdot \text{poly}(\log MN, \log T) \cdot \text{poly}(\lambda)$.
- The running time of UGRAM.Eval should be bounded by $T \cdot \text{poly}(\log MN, \log T) \cdot \text{poly}(\lambda)$ where T is the number of CPU steps in the description of the program P .

4.1 Construction

In this subsection, we give a construction of adaptive garbled RAM with unprotected memory access from updatable laconic oblivious transfer, somewhere equivocal encryption and garbling scheme for circuits with selective security using the techniques developed in the construction of adaptive garbled circuits [GS18a]. Our main theorem is:

Theorem 4.2 *Assuming the existence of updatable laconic oblivious transfer, somewhere equivocal encryption, a pseudorandom function and garbling scheme for circuits with selective security, there exists a construction of adaptive garbled RAM with unprotected memory access.*

1. $D \leftarrow \mathcal{A}(1^\lambda, 1^n)$.
2. \tilde{D} is computed as follows:
 - (a) If $b = 0$: $(\tilde{D}, SK) \leftarrow \text{UGRAM.Memory}(1^\lambda, 1^n, D)$.
 - (b) If $b = 1$: $\tilde{D} \leftarrow \text{Sim}(1^\lambda, 1^n, D)$.
3. **for** each i from t :
 - (a) $(P_i, I) \leftarrow \mathcal{A}(\tilde{D}, \{\tilde{P}_j, \tilde{x}_j\}_{j \in [i-1]})$ where $I \subset [|P_i|]$ and $|I| \leq n$.
 - (b) \tilde{P}_i is computed as follows:
 - i. If $b = 0$: $\tilde{P}_i \leftarrow \text{UGRAM.Program}(SK, i, P_i)$.
 - ii. If $b = 1$: $\tilde{P}_i \leftarrow \text{Sim}(\{C_{\text{CPU}}^{P_i, t}\}_{t \notin I})$
 - (c) $x_i \leftarrow \mathcal{A}(\{\tilde{P}_j, \tilde{x}_j\}_{j \in [i-1]}, \tilde{P}_i)$.
 - (d) \tilde{x}_i is computed as follows:
 - i. If $b = 0$: $\tilde{x}_i \leftarrow \text{UGRAM.Input}(SK, i, x_i)$
 - ii. If $b = 1$: $\tilde{x}_i \leftarrow \text{Sim}(x_i, \{y_t\}_{t \in I})$ where y_t is the o/p of $C_{\text{CPU}}^{P_i, t}$ when P_i is executed with x_i .
4. $b' \leftarrow \mathcal{A}(\{\tilde{P}_j, \tilde{x}_j\}_{j \in [t]})$.
5. Output b' .

Figure 6: $\text{Expt}_{\text{equiv}}(1^\lambda, \mathcal{A}, b)$

1. $D \leftarrow \mathcal{A}(1^\lambda, 1^n)$.
2. $(\tilde{D}, SK) \leftarrow \text{UGRAM.Memory}(1^\lambda, 1^n, D)$.
3. **for** x each i from t :
 - (a) $(P_{i,0}, P_{i,1}) \leftarrow \mathcal{A}(\tilde{D}, \{\tilde{P}_j, \tilde{x}_j\}_{j \in [i-1]})$.
 - (b) \tilde{P}_i is computed as follows:
 - i. If $b = 0$: $\tilde{P}_i \leftarrow \text{UGRAM.Program}(SK, i, P_{i,0})$.
 - ii. If $b = 1$: $\tilde{P}_i \leftarrow \text{UGRAM.Program}(SK, i, P_{i,1})$
 - (c) $x_i \leftarrow \mathcal{A}(\{\tilde{P}_j, \tilde{x}_j\}_{j \in [i-1]}, \tilde{P}_i)$.
 - (d) $\tilde{x}_i \leftarrow \text{UGRAM.Input}(i, SK, x_i)$
4. $b' \leftarrow \mathcal{A}(\{\tilde{P}_j, \tilde{x}_j\}_{j \in [t]})$.
5. Output b' if the output of each step circuit in $P_{i,0}^D(x_i)$ is same as $P_{i,1}^D(x_i)$ for every $i \in [t]$.

Figure 7: $\text{Expt}_{\text{UGRAM}}(1^\lambda, \mathcal{A}, b)$

Construction. We give the formal description of the construction in Figure 8. We use a somewhere equivocal encryption with block length set to $|\widetilde{\text{SC}}_\tau|$ where $\widetilde{\text{SC}}_\tau$ denotes the garbled version of the step circuit SC described in Figure 9, the message length to be T (which is the running time of the program P) and the equivocation parameter to be $t + \log T$ where t is the actual equivocation parameter for the UGRAM scheme.

Correctness. The correctness of the above construction follows from a simple inductive argument that for each step $\tau \in [|P|]$, the state and the database are updated correctly at the end of the execution of $\widetilde{\text{SC}}_\tau$. The base case is $\tau = 0$. In order to prove the inductive step for a step τ , observe that if the step τ outputs a read then labels recovered in Step 4.(c).(ii) of AdpEvalCkt correspond to data block in the location requested. Otherwise, the labels recovered in Step 4(b).(ii) of AdpEvalCkt corresponds to the updated value of the digest with the corresponding block written to the database.

Efficiency. The efficiency of our construction directly follows from the efficiency of updatable laconic oblivious transfer and the parameters set for the somewhere equivocal encryption. In particular, the running time of UGRAM.Memory is $D \cdot \text{poly}(\lambda)$, UGRAM.Program is $T \cdot \text{poly}(\log MN, \lambda)$ and that of UGRAM.Input is $n|x| \cdot \text{poly}(\log M, \log T, \lambda)$. The running time of UGRAM.Eval is $T \cdot \text{poly}(\log M, \log T, \lambda)$.

Security. In Appendix C, we argue the security of the construction.

5 Timed Encryption

In this section, we give the definition and construction of a timed encryption scheme. We will use a timed encryption scheme in the construction of adaptive garbled RAM in the next section.

A timed encryption scheme is a symmetric key encryption scheme with some special properties. In this encryption scheme, every message is encrypted with respect to a timestamp time. Additionally, there is a special algorithm called as constrain that takes an encryption key K and a timestamp time' as input and outputs a time constrained key $K[\text{time}']$. A time constrained key $K[\text{time}']$ can be used to decrypt any ciphertext that is encrypted with respect to timestamp $\text{time} < \text{time}'$. For security, we require that knowledge of a time constrained key does not help an adversary to distinguish between encryptions of two messages that are encrypted with respect to some future timestamp.

Definition 5.1 *A timed encryption scheme is a tuple of algorithms $(\text{TE.KeyGen}, \text{TE.Enc}, \text{TE.Dec}, \text{TE.Constrain})$ with the following syntax.*

- $\text{TE.KeyGen}(1^\lambda)$: It is a randomized algorithm that takes the security parameter 1^λ and outputs a key K .
- $\text{TE.Constrain}(K, \text{time})$: It is a deterministic algorithm that takes a key K and a timestamp $\text{time} \in [0, 2^\lambda - 1]$ and outputs a time-constrained key $K[\text{time}]$.
- $\text{TE.Enc}(K, \text{time}, m)$: It is a randomized algorithm that takes a key K , a timestamp time and a message m as input and outputs a ciphertext c or \perp .
- $\text{TE.Dec}(K, c)$: It is a deterministic algorithm that takes a key K and a ciphertext c as input and outputs a message m .

UGRAM.Memory($1^\lambda, 1^t, D$): On input a database $D \in \{\{0, 1\}^N\}^M$ do:

1. Sample $\text{crs} \leftarrow \text{crsGen}(1^\lambda, 1^N)$ and $K \leftarrow \text{PRFKeyGen}(1^\lambda)$ defining $\text{PRF}_K : \{0, 1\}^{2\lambda+1} \rightarrow \{0, 1\}^\lambda$.
2. For each $k \in [\lambda]$ and $b \in \{0, 1\}$, compute $\text{lab}_{k,b}^1 := \text{PRF}_K(1\|k\|b)$.
3. Compute $(d, \widehat{D}) = \text{Hash}(\text{crs}, D)$.
4. Output $\widehat{D}, \{\text{lab}_{k,d_k}^1\}_{k \in [\lambda]}$ as the garbled memory and (K, crs) as the secret key.

UGRAM.Program(SK, i, P): On input $SK = (K, \text{crs})$, sequence number i , and a program P (with T step-circuits) do:

1. For each step $\tau \in [2, T]$, $k \in [\lambda + n + N]$ and $b \in \{0, 1\}$,
 - (a) Sample $\text{lab}_{k,b}^\tau \leftarrow \{0, 1\}^\lambda$.
 - (b) Set $\text{lab}_{k,b}^1 := \text{PRF}_K(i\|k\|b)$ and $\text{lab}_{k,b}^{\tau+1} := \text{PRF}_K((i+1)\|k\|b)$.

We use $\{\text{lab}_{k,b}^\tau\}$ to denote $\{\text{lab}_{k,b}^\tau\}_{k \in [\lambda+n+N], b \in \{0,1\}}$.

2. **for** each τ from T down to **1** **do**:
 - (a) Compute $\widetilde{\text{SC}}_\tau \leftarrow \text{GarbleCkt}(1^\lambda, \text{SC}[\text{crs}, \tau, \{\text{lab}_{k,b}^{\tau+1}\}], \{\text{lab}_{k,b}^\tau\})$ where the step-circuit SC is described in Figure 9.
3. Compute $\text{key} = \text{KeyGen}(1^\lambda; \text{PRF}_K(i\|0^\lambda\|0))$
4. Compute $c \leftarrow \text{Enc}(\text{key}, \{\widetilde{\text{SC}}_\tau\}_{\tau \in [T]})$ and output $\widetilde{P} := c$.

UGRAM.Input(SK, i, x): On input the secret key $SK = (K, \text{crs})$, sequence number i and a string $x \in \{0, 1\}^n$ do:

1. For each $k \in [\lambda + n + N]$ and $b \in \{0, 1\}$, compute $\text{lab}_{k,b}^1 := \text{PRF}_K(i\|k\|b)$.
2. Compute $\text{key} = \text{KeyGen}(1^\lambda; \text{PRF}_K(i\|0^\lambda\|0))$.
3. Output $\widetilde{x} := (\text{key}, \{\text{lab}_{k,x_k}^1\}_{k \in [\lambda+1, \lambda+n]}, \{\text{lab}_{k,0}^1\}_{k \in [n+\lambda+1, n+\lambda+N]})$.

UGRAM.Eval $\widetilde{D}(i, \text{st}, \widetilde{P}, \widetilde{x})$: On input i , state st , the garbled program \widetilde{P} , and garbled input \widetilde{x} do:

1. Parse \widetilde{x} as $(\text{key}, \{\text{lab}_k\}_{k \in [\lambda+1, n+\lambda+N]})$ and \widetilde{P} as c .
2. If $i = 1$, obtain $\{\text{lab}_k\}_{k \in [\lambda]}$ from garbled memory; else, parse st as $\{\text{lab}_k\}_{k \in [\lambda]}$.
3. Compute $\{\widetilde{\text{SC}}_\tau\}_{\tau \in [T]} := \text{Dec}(\text{key}, c)$ and set $\overline{\text{lab}} := \{\text{lab}_k\}_{k \in [n+\lambda+N]}$.
4. **for** each τ from **1** to T **do**:
 - (a) Compute $(\text{R/W}, L, A, \{\text{lab}_k\}_{k \in [\lambda+1, \lambda+n]}, B) := \text{EvalCkt}(\widetilde{\text{SC}}_\tau, \overline{\text{lab}})$.
 - (b) If $\text{R/W} = \text{write}$,
 - i. Parse A as (e_w, wData) and B as $\{\text{lab}_k\}_{k \in [\lambda+1, n+\lambda+N]}$.
 - ii. $\{\text{lab}_k\}_{k \in [\lambda]} \leftarrow \text{ReceiveWrite}^{\widehat{D}}(\text{crs}, L, \text{wData}, e_w)$
 - (c) **else**,
 - i. Parse A as $\{\text{lab}_k\}_{k \in [n+\lambda]}$ and B as e .
 - ii. $\{\text{lab}_k\}_{k \in [n+\lambda+1, n+\lambda+N]} \leftarrow \text{Receive}^{\widehat{D}}(\text{crs}, L, e)$
 - (d) Set $\overline{\text{lab}} := \{\text{lab}_k\}_{k \in [n+\lambda+N]}$.
5. Parse $\overline{\text{lab}}$ as $\{\text{lab}_k\}_{k \in [n+\lambda+N]}$. Output $\{\text{lab}_k\}_{k \in [\lambda+1, n+\lambda]}$ and $\text{st} := \{\text{lab}_k\}_{k \in [\lambda]}$.

Figure 8: Adaptive Garbled RAM with Unprotected Memory Access

Step Circuit SC

Input: A digest d , state state and a block rData .

Hardcoded: The common reference string crs , the step number τ and a set of labels $\{\text{lab}_{k,b}\}$.

1. Compute $(\text{state}', R/W, L, \text{wData}) := C_{\text{CPU}}^{P,\tau}(\text{state}, \text{rData})$.
2. If $\tau = T$, reset $\text{lab}_{k,b} = b$ for all $k \in [\lambda + 1, \lambda + n]$ and $b \in \{0, 1\}$.
3. **if** $R/W = \text{write}$ **do**:
 - (a) Compute $e_w \leftarrow \text{SendWrite}(\text{crs}, d, L, \text{wData}, \{\text{lab}_{k,b}\}_{k \in [\lambda], b \in \{0,1\}})$.
 - (b) Output $(R/W, L, e_w, \text{wData}, \{\text{lab}_{k,\text{state}'_{k-\lambda}}\}_{k \in [\lambda+1, \lambda+n]}, \{\text{lab}_{k,0}\}_{k \in [n+\lambda+1, n+\lambda+N]})$.
4. **else**,
 - (a) Compute $e \leftarrow \text{Send}(\text{crs}, d, L, \{\text{lab}_{k,b}\}_{k \in [n+\lambda+1, n+\lambda+N], b \in \{0,1\}})$.
 - (b) $(R/W, L, \{\text{lab}_{k,d_k}\}_{k \in [\lambda]}, \{\text{lab}_{k,\text{state}'_{k-\lambda}}\}_{k \in [\lambda+1, \lambda+n]}, e)$.

Figure 9: Description of the Step Circuit

We require a timed encryption scheme to follow the following properties.

Correctness. We require that for all messages m and for all timestamps $\text{time}_1 \leq \text{time}_2$:

$$\Pr[\text{TE.Dec}(K[\text{time}_2], c) = m] = 1$$

where $K \leftarrow \text{TE.KeyGen}(1^\lambda)$, $K[\text{time}_2] := \text{TE.Constrain}(K, \text{time}_2)$ and $c \leftarrow \text{TE.Enc}(K, \text{time}_1, m)$.

Encrypting with Constrained Key. For any message m and timestamps $\text{time}_1 \leq \text{time}_2$, we require that:

$$\{\text{TE.Enc}(K, \text{time}_1, m)\} \approx \{\text{TE.Enc}(K[\text{time}_2], \text{time}_1, m)\}$$

where $K \leftarrow \text{TE.KeyGen}(1^\lambda)$, $K[\text{time}_2] := \text{TE.Constrain}(K, \text{time}_2)$ and \approx denotes that the two distributions are identical.

Security. For any two messages m_0, m_1 and timestamps $(\text{time}, \{\text{time}_i\}_{i \in [t]})$ where $\text{time}_i < \text{time}$ for all $i \in [t]$, we require that:

$$\{\{K[\text{time}_i]\}_{i \in [t]}, \text{TE.Enc}(K, \text{time}, m_0)\} \stackrel{c}{\approx} \{\{K[\text{time}_i]\}_{i \in [t]}, \text{TE.Enc}(K, \text{time}, m_1)\}$$

where $K \leftarrow \text{TE.KeyGen}(1^\lambda)$ and $K[\text{time}_i] := \text{TE.Constrain}(K, \text{time}_i)$ for every $i \in [t]$.

Theorem 5.2 Assuming the existence of one-way functions, there exists a construction of timed encryption.

Construction. Let $(\text{SK.KeyGen}, \text{SK.Enc}, \text{SK.Dec})$ be a symmetric key encryption scheme. We assume without loss of generality that $\text{SK.KeyGen}(1^\lambda)$ outputs an uniformly chosen random string in $\{0, 1\}^\lambda$. We describe the construction below.

- $\text{TE.KeyGen}(1^\lambda)$: Sample $K \leftarrow \text{RC.KeyGen}(1^\lambda)$ defining $\text{PRF}_K : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ and output K .
- $\text{TE.Enc}(K, \text{time}, m)$: Compute $SK := \text{PRF}_K(\text{time})$ and output $(\text{time}, \text{SK.Enc}(SK, m))$.
- $\text{TE.Dec}(K, c)$: Parse c as (time, ct) and compute $SK := \text{PRF}_K(\text{time})$ and output $\text{SK.Dec}(SK, ct)$.
- $\text{TE.Constrain}(K, \text{time})$: Output $\text{RC.Constrain}(K, \text{time})$.

We note that correctness follows directly from the correctness of range constrained PRF and symmetric key encryption scheme. The encryption with constrained key property follows from the observation that for any $\text{time} \leq \text{time}'$, $\text{PRF}_{K[\text{time}']}(\text{time}) = \text{PRF}_K(\text{time})$. The security property can be easily argued from the security of range constrained PRF and symmetric key encryption.

6 Construction of Adaptive Garbled RAM

In this section, we give a construction of adaptive garbled RAM. We make use of the following primitives.

- A timed encryption scheme $(\text{TE.KeyGen}, \text{TE.Enc}, \text{TE.Dec}, \text{TE.Constrain})$. Let N be the output length of TE.Enc when encrypting single bit messages.
- A puncturable pseudorandom function $(\text{PP.KeyGen}, \text{PP.Eval}, \text{PP.Punc})$.
- An oblivious RAM scheme $(\text{OData}, \text{OProg})$ with strong localized randomness.
- An adaptive garbled RAM scheme UGRAM with unprotected memory access.

The formal description of our construction appears in Figure 10.

Correctness. We give an informal argument for correctness. The only difference between UGRAM and the construction we give in Figure 10 is that we encrypt the database using a timed encryption scheme and encode it using a ORAM scheme. To argue the correctness of our construction, it is sufficient to argue that each step circuit SC faithfully emulates the corresponding step circuit of P^* . Let $\text{SC}^{i,\tau}$ be the step circuit that corresponds to the τ^{th} step of the i^{th} program P_i . We observe that any point in time the L^{th} location of the database \widehat{D} is an encryption of the actual data bit with respect to timestamp $\text{time} := (i' || \tau')$ where $\text{SC}^{i',\tau'}$ last wrote at the L^{th} location. It now follows from this invariant and the correctness of the timed encryption scheme that the hardwired constrained key $K[i || \tau]$ in $\text{SC}^{i,\tau}$ can be used to decrypt the read block X as the step that last modified this block has a timestamp that is less than $(i || \tau)$.

Efficiency. We note that setting the equivocation parameter $n = \text{poly}(\log MN)$, we obtain that the running time of GRAM.Input is $|x| \cdot \text{poly}(\lambda, \log MN)$. The rest of the efficiency criterion follow directly from the efficiency of adaptive garbled RAM with unprotected memory access.

GRAM.Memory($1^\lambda, D$): On input the database $D \in \{0, 1\}^M$:

1. Sample $K \leftarrow \text{TE.KeyGen}(1^\lambda)$ and $S \leftarrow \text{PRFKeyGen}(1^\lambda)$ defining $\text{PRF}_S : \{0, 1\}^\lambda \rightarrow \{0, 1\}^n$ (where n is the input length of each program).
2. Initialize an empty array \hat{D} of M blocks with block length N .
3. **for** each i from 1 to M **do**:
 - (a) Set $\hat{D}[i] \leftarrow \text{TE.Enc}(K, 0^\lambda, D[i])$.
4. $D^* \leftarrow \text{OData}(1^\lambda, 1^N, \hat{D})$.
5. $(\tilde{D}, SK) \leftarrow \text{UGRAM.Memory}(1^\lambda, 1^t, D^*)$ where $t = \text{poly}(\log MN)$.
6. Output \tilde{D} as the garbled memory and (K, S, SK) as the secret key.

GRAM.Program(SK', i, P): On input $SK' = (K, S, SK)$, sequence number i , and a program P :

1. Sample $K' \leftarrow \text{PP.KeyGen}(1^\lambda)$
2. $P^* \leftarrow \text{OProg}(1^\lambda, 1^{\log M}, 1^T, P)$ where P^* runs in time T' .
3. For each $\tau \in [T']$, compute $K[(i||\tau)] \leftarrow \text{TE.Constrain}(K, (i||\tau))$ where $(i||\tau)$ is expressed as a λ -bit string.
4. Compute $r := \text{PRF}_S(i)$.
5. Let τ_1, \dots, τ_m be the sequence of values guaranteed by strong localized randomness.
6. **for** each $\tau \in [T']$ **do**:
 - (a) Let $j \in [m - 1]$ be such that $\tau \in [\tau_j, \tau_{j+1})$.
 - (b) Let $C_{\text{CPU}}^\tau := \text{SC}_\tau[i, \tau, K[(i||\tau)], I_j, K', r']$ where $r' = r$ if $\tau = T'$, else $r' = \perp$. The step circuit SC is described in Figure 11
7. Construct a RAM program P' with step-circuits given by $\{C_{\text{CPU}}^\tau\}$.
8. $\tilde{P} \leftarrow \text{UGRAM.Program}(SK, i, P')$.
9. Output \tilde{P} .

GRAM.Input(SK', i, P): On input $SK' = (K, S, SK)$, i and x :

1. Compute $r = \text{PRF}_S(i)$
2. Compute $\hat{x} \leftarrow \text{UGRAM.Input}(SK, i, x)$
3. Output $\tilde{x} = (\hat{x}, r)$.

GRAM.Eval $\tilde{D}(i, \text{st}, \tilde{P}, \tilde{x})$: On input state st , the garbled program \tilde{P} , and garbled input \tilde{x} :

1. Compute $(y, \text{st}') \leftarrow \text{UGRAM.Eval}\tilde{D}(\text{st}, \tilde{P}, \hat{x})$ and update st to st' . Output $y \oplus r$.

Figure 10: Construction of Adaptive GRAM

Step Circuit SC_τ

Input: A ciphertext c_{CPU} and a data block $X \in \{0, 1\}^N$.

Hardcoded: The sequence number i , step number τ , the constrained key $K[(i||\tau)]$, the interval I_j , the key K' and a string r' .

1. Compute $\text{rData} := \text{TE.Dec}(K[(i||\tau)], X)$ and $\text{state} = \text{TE.Dec}(K[(i||\tau)], c_{\text{CPU}})$.
2. Compute $R_{I_j} = \text{PP.Eval}(K', I_j)$.
3. Compute $(\text{R/W}, L, \text{state}', \text{wData}) := C_{\text{CPU}}^{P^*, \tau}(\text{state}, \text{rData}, R_{I_j})$.
4. **if** $\tau = T'$, then output $c'_{\text{CPU}} = \text{state}' \oplus r'$; else $c'_{\text{CPU}} = \text{TE.Enc}(K[(i||\tau)], \text{state}')$.
5. **else if** $\text{R/W} = \text{write}$ **do**:
 - (a) Compute $X' \leftarrow \text{TE.Enc}(K[i||\tau], (i, \tau), \text{wData})$.
 - (b) Output $(c'_{\text{CPU}}, \text{R/W}, L, X')$.
6. **else if** $\text{R/W} = \text{read}$, output $(c'_{\text{CPU}}, \text{R/W}, L, \perp)$.

Figure 11: Description of the Step Circuit

Security. In Appendix D, we argue the security of our construction.

References

- [ACC⁺16] Prabhanjan Ananth, Yu-Chi Chen, Kai-Min Chung, Huijia Lin, and Wei-Kai Lin. Delegating RAM computations with adaptive soundness and privacy. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 3–30. Springer, Heidelberg, October / November 2016.
- [AIK04] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0 . In *45th FOCS*, pages 166–175. IEEE Computer Society Press, October 2004.
- [AIK05] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Computationally private randomizing polynomials and their applications. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity, CCC '05*, pages 260–274, Washington, DC, USA, 2005. IEEE Computer Society.
- [App17] Benny Applebaum. Garbled circuits as randomized encodings of functions: a primer. *IACR Cryptology ePrint Archive*, 2017:385, 2017.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2001.

- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014.
- [BGL⁺15] 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 ACM STOC*, pages 439–448. ACM Press, June 2015.
- [BHR12a] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 134–153. Springer, Heidelberg, December 2012.
- [BHR12b] 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*, pages 784–796. ACM Press, October 2012.
- [BL18] Fabrice Benhamouda and Huijia Lin. k-round multiparty computation from k-round oblivious transfer via garbled interactive circuits. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 500–532. Springer, Heidelberg, April / May 2018.
- [BLSV18] Zvika Brakerski, Alex Lombardi, Gil Segev, and Vinod Vaikuntanathan. Anonymous ibe, leakage resilience and circular security from new assumptions. To appear in Eurocrypt, 2018. <https://eprint.iacr.org/2017/967>.
- [BMR90] Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols (extended abstract). In *22nd ACM STOC*, pages 503–513. ACM Press, May 1990.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.
- [CCHR16] Ran Canetti, Yilei Chen, Justin Holmgren, and Mariana Raykova. Adaptive succinct garbled RAM or: How to delegate your database. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 61–90. Springer, Heidelberg, October / November 2016.
- [CDG⁺17] Chongwon Cho, Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Antigoni Polychroniadou. Laconic receiver oblivious transfer and applications. *To appear in Crypto*, 2017.
- [CH16] Ran Canetti and Justin Holmgren. Fully succinct garbled RAM. In Madhu Sudan, editor, *ITCS 2016*, pages 169–178. ACM, January 2016.

- [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, *47th ACM STOC*, pages 429–437. ACM Press, June 2015.
- [CP13] Kai-Min Chung and Rafael Pass. A simple oram. Cryptology ePrint Archive, Report 2013/243, 2013. <https://eprint.iacr.org/2013/243>.
- [DG17] Nico Döttling and Sanjam Garg. Identity based encryption from diffie-hellman assumptions. *To appear in Crypto*, 2017.
- [DGHM18] Nico Döttling, Sanjam Garg, Mohammad Hajiabadi, and Daniel Masny. New constructions of identity-based and key-dependent message secure encryption schemes. To appear in PKC, 2018. <https://eprint.iacr.org/2017/978>.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- [GGMP16] Sanjam Garg, Divya Gupta, Peihan Miao, and Omkant Pandey. Secure multiparty RAM computation in constant rounds. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 491–520. Springer, Heidelberg, October / November 2016.
- [GHL⁺14] Craig Gentry, Shai Halevi, Steve Lu, Rafail Ostrovsky, Mariana Raykova, and Daniel Wichs. Garbled RAM revisited. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 405–422. Springer, Heidelberg, May 2014.
- [GHRW14] Craig Gentry, Shai Halevi, Mariana Raykova, and Daniel Wichs. Outsourcing private RAM computation. In *55th FOCS*, pages 404–413. IEEE Computer Society Press, October 2014.
- [GKK⁺12] 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*, pages 513–524. ACM Press, October 2012.
- [GKP⁺13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 555–564. ACM Press, June 2013.
- [GLO15] Sanjam Garg, Steve Lu, and Rafail Ostrovsky. Black-box garbled RAM. In Venkatesan Guruswami, editor, *56th FOCS*, pages 210–229. IEEE Computer Society Press, October 2015.

- [GLOS15] Sanjam Garg, Steve Lu, Rafail Ostrovsky, and Alessandra Scafuro. Garbled RAM from one-way functions. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 449–458. ACM Press, June 2015.
- [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 ACM STOC*, pages 218–229. ACM Press, May 1987.
- [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 ACM STOC*, pages 182–194. ACM Press, May 1987.
- [GS18a] Sanjam Garg and Akshayaram Srinivasan. Adaptively secure garbling with near optimal online complexity. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, pages 535–565, 2018.
- [GS18b] Sanjam Garg and Akshayaram Srinivasan. Two-round multiparty secure computation from minimal assumptions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 468–499. Springer, Heidelberg, April / May 2018.
- [HJO⁺16] Brett Hemenway, Zahra Jafargholi, Rafail Ostrovsky, Alessandra Scafuro, and Daniel Wichs. Adaptively secure garbled circuits from one-way functions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 149–178. Springer, Heidelberg, August 2016.
- [HY16] Carmit Hazay and Avishay Yanai. Constant-round maliciously secure two-party computation in the RAM model. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 521–553. Springer, Heidelberg, October / November 2016.
- [IKO⁺11] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, Manoj Prabhakaran, and Amit Sahai. Efficient non-interactive secure computation. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 406–425. Springer, Heidelberg, May 2011.
- [JKK⁺17] Zahra Jafargholi, Chethan Kamath, Karen Klein, Ilan Komargodski, Krzysztof Pietrzak, and Daniel Wichs. Be adaptive, avoid overcommitting. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, pages 133–163, 2017.
- [JW16] Zahra Jafargholi and Daniel Wichs. Adaptive security of Yao’s garbled circuits. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of *LNCS*, pages 433–458. Springer, Heidelberg, October / November 2016.

- [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, *47th ACM STOC*, pages 419–428. ACM Press, June 2015.
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 669–684. ACM Press, November 2013.
- [KY18] Marcel Keller and Avishay Yanai. Efficient maliciously secure multiparty computation for ram. To appear in EUROCRYPT, 2018. <https://eprint.iacr.org/2017/981>.
- [LO13] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 719–734. Springer, Heidelberg, May 2013.
- [LO17] Steve Lu and Rafail Ostrovsky. Black-box parallel garbled RAM. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II*, pages 66–92, 2017.
- [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.
- [Mia16] Peihan Miao. Cut-and-choose for garbled RAM. Cryptology ePrint Archive, Report 2016/907, 2016. <http://eprint.iacr.org/2016/907>.
- [ORS15] Rafail Ostrovsky, Silas Richelson, and Alessandra Scafuro. Round-optimal black-box two-party computation. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 339–358. Springer, Heidelberg, August 2015.
- [OS97] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In *29th ACM STOC*, pages 294–303. ACM Press, May 1997.
- [Ost90] Rafail Ostrovsky. Efficient computation on oblivious RAMs. In *22nd ACM STOC*, pages 514–523. ACM Press, May 1990.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.
- [WHC⁺14] 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*, pages 191–202. ACM Press, November 2014.
- [Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164. IEEE Computer Society Press, November 1982.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.

A Constructing Range Constrained PRF

In this section, we describe a construction of range constrained PRF from one-way functions.

Construction. We now give a construction of range constrained PRF from any length doubling PRG : $\{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$. The construction is exactly same as the GGM construction [GGM86] with an additional constrain algorithm. We will denote the left half of the output of PRG by $\text{PRG}_0(\cdot)$ and the right half by $\text{PRG}_1(\cdot)$.

- $\text{RC.KeyGen}(1^\lambda)$: Sample $s \leftarrow \{0, 1\}^\lambda$ as the seed of the PRG and output s as the PRF key. The implicit PRF $_s$: $\{0, 1\}^n \rightarrow \{0, 1\}^\lambda$ for any $n = \text{poly}(\lambda)$ is defined as $\text{PRF}_s(x_0x_1 \dots x_{n-1}) := \text{PRG}_{x_{n-1}}(\text{PRG}_{x_{n-2}}(\dots \text{PRG}_{x_0}(s)))$.
- $\text{RC.Constrain}(s, T)$:
 - Parse the binary representation of T as $T_0 \dots T_{n-1}$.
 - Initialize an empty set S .
 - Set $t := s$.
 - **for** each i from 0 to $n - 1$ **do**:
 - * If T_i is 1, add $(i, \text{PRG}_0(t))$ to S and reset $t := \text{PRG}_1(t)$.
 - * Else, reset $t := \text{PRG}_0(t)$.
 - Add (v, t) to S (where v is a special symbol) and output S as the constrained key.
- $\text{RC.Eval}(S, x)$:
 - To evaluate PRF using the constrained key S on an input $x < T$, compute the least index i such that $x_i = 0$ and $T_i = 1$ (such a bit must exist since $x < T$). Recover (i, t_i) from S and output $\text{PRG}_{x_{n-1}}(\dots \text{PRG}_{x_{i+1}}(t_i))$ as the output.
 - To evaluate PRF using the constrained key S on an input T , obtain (v, t) from S and output t .

The correctness of the above construction is easy to observe. The security can be argued as follows. Let y be the challenge point and by definition $y > \max_k T_k$. Let us call $\max_k T_k$ as T . Let i be the least index such that $y_i = 1$ and $T_i = 0$ (such an index must exist since $y > T$). By construction, we have that $s' = \text{PRG}_{y_i}(\dots \text{PRG}_{y_0}(s))$ is not present in any of the constrained keys $\{K[T_k]\}_{k \in [\ell]}$. We can use the security of PRG repeatedly to replace this string with random. We can then invoke the PRG security again to show that $\text{PRG}_{y_{n-1}}(\text{PRG}(y_{n-2} \dots \text{PRG}_{y_{i+1}}(s')) \dots)$ is indistinguishable to random given the constrained key.

B Strong Localized Randomness of [CP13]

In this section, we show that the Chung-Pass ORAM scheme [CP13] instantiated with the following special encryption scheme satisfies the strong localized randomness property.

1. $\text{KeyGen}(1^\lambda)$: Sample $K \leftarrow \text{PP.KeyGen}(1^\lambda)$.

2. $\text{Enc}(K, \tau, v)$: The encryption algorithm takes as input the key K , the time step τ and the value v . It samples a random string $r \leftarrow \{0, 1\}^\lambda$ and outputs $(\tau \| r, \text{PP.Eval}(K, (\tau \| r)) \oplus v)$ as the ciphertext.
3. $\text{Dec}(K, c)$: The ciphertext is parsed as (c_1, c_2) and the value v is computed as $\text{PP.Eval}(K, c_1) \oplus c_2$.

We now recall the Chung-Pass ORAM scheme [CP13]. The following text is taken verbatim from [CP13].

In database D^* , is maintained as a complete binary tree Γ of depth $d = \log(M/\alpha)$; we index nodes in the tree by a binary string of length at most d , where the root is indexed by the empty string ε , and each node indexed by γ has left and right children indexed $\gamma 0$ and $\gamma 1$, respectively. Each memory cell r will be associated with a random leaf pos in the tree, specified by the position map Pos ; as we shall see shortly, the memory cell r will be stored at one of the nodes on the path from the root ε to the leaf pos . To ensure that the position map is smaller than the memory size, we assign a block of α consecutive memory cells to the same leaf; thus memory cell r corresponding to block $b = \lfloor r/\alpha \rfloor$ will be associated with leaf $Pos(b)$.

Each node in the tree is associated with a bucket which stores (at most) K tuples (b, pos, v) where v is the content of block b and pos is the leaf associated with the block b ; $K \in (\log M) \text{poly} \log(M)$ is a parameter that will determine the security of the ORAM (thus each bucket stores $K(\alpha + 2)$ words.) We assume that all registers and memory cells are initialized with a special symbol \perp .

We first specify how a datablock is read.

- **Fetch:** Let $b = \lfloor r/\alpha \rfloor$ be the block containing memory cell r , and let $i = r \bmod \alpha$ be the i th component in the block b . We first look up the position of the block b using the position map: $\text{pos} = Pos(b)$; if $Pos(b) = \perp$, let $\text{pos} \leftarrow \lfloor n/\alpha \rfloor$ to be a uniformly random leaf. Next, we traverse the tree from the root to the leaf pos , making exactly one read and one write operation for every memory cell associated with the nodes along the path. More precisely, we read the content once, and then we either write it back (unchanged), or we simply erase it (writing \perp) so as to implement the following task: search for a tuple of the form (b, pos, v) in any of the nodes during the traversal; if such a tuple is found, remove it, and otherwise let $v = \perp$. Finally return the i^{th} component of v .
- **Update Position Map:** Pick a uniformly random leaf $\text{pos}' \leftarrow \lfloor n/\alpha \rfloor$ and let $Pos(b) = \text{pos}'$.
- **Put Back:** Add the tuple (b, pos', v) to the root ε of the tree. If there is not enough space left in the bucket, abort outputting overflow.
- **Flush:** Pick a uniformly random leaf $\text{pos}^* \leftarrow \lfloor n/\alpha \rfloor$ and traverse the tree from the root to the leaf pos^* , making exactly one read and one write operation for every memory cell associated with the nodes along the path so as to implement the following task: push down each tuple (b'', pos'', v'') read in the nodes traversed as far as possible along the path to pos^* while ensuring that the tuple is still on the path to its associated leaf pos'' (that is, the tuple ends up in the node $\gamma = \text{longest common prefix of } \text{pos}'' \text{ and } \text{pos}^*$.) (Note that this operation can be done trivially as long as the CPU has sufficiently many work registers to load two whole buckets into memory; since the bucket size is polylogarithmic, this is possible.) If at any point some bucket is about to overflow, abort outputting overflow.

The process of writing a datablock val to a location proceeds identically to how a datablock is read, except that in the Put Back steps, we add the tuple (b, pos', v') where v' is the string v but the i^{th} component is set to val (instead of adding the tuple (b, pos', v) as in reading).

Modified ORAM scheme. The modification that we make to the Chung-Pass ORAM scheme is that every value v that is being written to the database by the τ^{th} step-circuit is encrypted using the special encryption scheme with respect to the timestep τ .

Strong Localized Randomness. We now argue why this ORAM construction satisfies the strong localized randomness property. Note that the random tape accessed by each memory access is independent of other memory accesses. This shows the first two properties in Definition 3.11. We now argue the third property. Note that the second root to leaf path that is being accessed is random and independent. The only issue is the first path that is being accessed is a root to leaf path where the leaf node is maintained in the position map. In order to make this access to be random and independent, we need to change how the Pos map is updated and that is where the special encryption scheme helps us. For a particular memory access to a location L , let τ_k to τ_{k+1} be the step circuits that last updated the Pos map for this location. Now, instead of pushing the correct leaf node in the Pos map, we use the security of the special encryption scheme to push a junk value. It now follows from that the first root to leaf path for accessing the location L is also random. We can even make the computational indistinguishability argument even when given the encryption key K punctured at $\tau || r_\tau$ for each $\tau \in [\tau_k, \tau_{k+1}]$.

C Security of UMA Adaptive Garbled RAM

C.1 Security

We now show that the construction given in Figure 8 satisfies the equivocability and the adaptive security properties described in Definition 4.1. Before we proceed to proving these two properties, we give the description of the simulator (for the equivocability property) in Figure 12. Recall that this simulator takes as input the set I that has to be equivocated along with the set of step circuits not in I . We denote this set by $P_{\setminus I}$. For the ease of exposition, we split the simulator into three sub-routines: $SimD$, $SimP$ and $SimIn$ that correspond to simulating the database, simulating the program and simulating the input respectively.

C.1.1 Equivocability

To prove equivocability, we need to show that the advantage of any adversary in the experiment described in Figure 6 is negligible. In order to show this, we argue that for every sequence number, $Sim(I, P)$ is computationally indistinguishable to the real world garbling of program P . The rest follows via a standard hybrid argument.

We now argue that $Sim(I, P_{\setminus I})$ is computationally indistinguishable to the real world garbling of program P . Let us assume that P has a running time of T or in other words, is described by T CPU steps. In this proof of indistinguishability, we assume that the labels for the first garbled circuit and the coins used in $KeyGen$ are generated randomly instead of using a PRF key to compute them. This assumption follows directly from the security of the PRF.

$\text{SimD}(1^\lambda, 1^n, D)$: On input $1^\lambda, 1^n$ and the database D , compute \tilde{D} exactly as in UGRAM. Memory and output the garbled database \tilde{D} .

$\text{SimP}(1^\lambda, i, I, P_{\setminus I}, \text{st})$: On input $SK = (K, \text{crs})$, sequence number i , the set I and $\{C_{\text{CPU}}^{P_i, t}\}_{t \notin I}$ do:

1. For step $\tau \in [2, T]$, $k \in [\lambda + n + N]$ and $b \in \{0, 1\}$,
 - (a) Sample $\text{lab}_{k,b}^\tau \leftarrow \{0, 1\}^\lambda$.
 - (b) Set $\text{lab}_{k,b}^1 := \text{PRF}_K(i \| k \| b)$ and $\text{lab}_{k,b}^{T+1} := \text{PRF}_K((i+1) \| k \| b)$.
2. **for** each τ from T down to 1 such that $\tau \notin I$ **do**:
 - (a) Compute $\tilde{\text{SC}}_\tau \leftarrow \text{GarbleCkt}\left(1^\lambda, \text{SC}[\text{crs}, \tau, \{\text{lab}_{k,b}^{\tau+1}\}], \{\text{lab}_{k,b}^\tau\}\right)$ where the step-circuit SC is described in Figure 9.
3. Compute $(\text{st}_1, c) \leftarrow \text{SimEnc}(I, \{\tilde{\text{SC}}_\tau\}_{\tau \notin I})$.
4. Output $\tilde{P} := c$.

$\text{SimIn}(\text{st}, i, x, \{y_\tau\}_{\tau \in I})$: On input the string $x \in \{0, 1\}^n$, and $\{y_\tau\}_{\tau \in I}$ do:

Notation: For every $\tau \in I$, we parse $y_\tau = (\text{state}^{i,\tau}, \text{R/W}^{i,\tau}, L^{i,\tau}, \text{wData}^{i,\tau})$ where $y_{i,\tau}$ is the output of the τ^{th} step circuit of program P_i . Let $D^{i,\tau}$ to be content of the database after the execution of the τ^{th} step in the i^{th} program. We let $\text{d}^{i,\tau}$ be the digest of $D^{i,\tau}$ (i.e., $(\text{d}^{i,\tau}, \cdot) := \text{Hash}(\text{crs}, D^{i,\tau})$). and $\text{inp}^{i,\tau}$ to be the input to SC_τ of the i^{th} program. We define $D^{i,0}, \text{inp}^{i,0}, \text{state}^{i,0}, \text{d}^{i,0}$ to be database, input to the first step circuit SC_1 , the CPU state and the digest before starting the execution of program i .

1. **for** each τ from T down to 1 such that $\tau \in I$:
 - (a) For each $k \in [\lambda + n + N]$ and $b \in \{0, 1\}$, compute $\text{lab}_{k,b}^1 := \text{PRF}_K(i \| k \| b)$ and $\text{lab}_{k,b}^{T+1} := \text{PRF}_K((i+1) \| k \| b)$.
 - (b) **if** $\text{R/W}^{i,\tau} = \text{write}$ **do**:
 - i. Compute $c \leftarrow \text{Sim}_{\text{OTW}}(\text{crs}, D^{i,\tau-1}, L^{i,\tau}, \text{wData}^{i,\tau}, \{\text{lab}_{k,\text{d}_k^{i,\tau+1}}^{\tau+1}\}_{k \in [\lambda]})$.
 - ii. Set $m := \left(\text{R/W}^{i,\tau}, L^{i,\tau}, c, \text{wData}, \{\text{lab}_{k+\lambda, \text{state}_k^{i,\tau}}^{\tau+1}\}_{k \in [n]}, \{\text{lab}_{k+n+\lambda, 0}^{\tau+1}\}_{k \in [N]}\right)$.
 - iii. Compute $\tilde{\text{SC}}_\tau \leftarrow \text{Sim}_{\text{Ckt}}\left(1^\lambda, 1^{|\text{SC}|}, m, \{\text{lab}_{k, \text{inp}_k^{i,\tau}}^\tau\}_{k \in [n+\lambda+N]}\right)$
 - (c) **else**,
 - i. Compute $c \leftarrow \text{Sim}_{\text{OT}}(\text{crs}, D^{i,\tau-1}, L^{i,\tau}, \{\text{lab}_{j, \text{inp}_j^{i,\tau+1}}\}_{j \in [n+\lambda+1, n+\lambda+N]})$.
 - ii. Set $m := \left(\text{R/W}^{i,\tau}, L^{i,\tau}, \{\text{lab}_{k, \text{d}_k^{i,\tau+1}}^{\tau+1}\}_{k \in [\lambda]}, \{\text{lab}_{k+\lambda, \text{state}_k^{i,\tau}}^{\tau+1}\}_{k \in [n]}, c\right)$.
 - iii. Compute $\tilde{\text{SC}}_\tau \leftarrow \text{Sim}_{\text{Ckt}}\left(1^\lambda, 1^{|\text{SC}|}, m, \{\text{lab}_{k, \text{inp}_k^{i,\tau}}^\tau\}_{k \in [n+\lambda+N]}\right)$
2. Compute $\text{key} \leftarrow \text{SimKey}(\text{st}_1, \{\tilde{\text{SC}}_\tau\}_{\tau \in I})$.
3. Output $\tilde{x} := (\text{key}, \{\text{lab}_{k+\lambda, x_k}^1\}_{k \in [n]}, \{\text{lab}_{k+n+\lambda, 0}\}_{k \in [N]})$.

Figure 12: $\text{Sim}(I, P_{\setminus I})$

In order to show the computational indistinguishability, we define a sequence of hybrids where the indistinguishability between the successive hybrids is shown via the following rule.

Rule A. This rule states that for any two sets $I, I' \subseteq [T]$ such that $I' = I \cup \{g^*\}$, $\text{Sim}(I, P_{\setminus I})$ is computationally indistinguishable to $\text{Sim}(I', P_{\setminus I'})$ if g^* is the first step or if $g^* - 1$ is in I .

Hybrid Sequence. Our sequence of hybrids that proves the indistinguishability is described via an optimal strategy for the following pebbling game. The rule A described above corresponds to the rule of our pebbling game below.

Pebbling Game. Consider the positive integer line $1, 2, \dots, T$. We are given a certain number pebbles. We can place the pebbles on this positive integer line according to the following rule:

Rule A: We can place or remove a pebble in position i if and only if there is a pebble in position $i - 1$. This restriction does not apply to position 1: we can always place or remove a pebble at position 1.

Note that the set of indexes where a pebble is present corresponds to the equivocated set I . Thus, any configuration of the pebbling game where there is a pebble in position $i \in I$ corresponds to the hybrid $\text{Sim}(I, P)$.

Optimization goal of the pebbling game. The goal is to pebble the line $[1, T]$ such that every position in I has a pebble while minimizing the number of pebbles that are present on the line at any point in time.

Optimal Pebbling Strategy. We provide an optimal pebbling strategy that uses $t + \log(T)$ pebbles. We first state the following lemma which was proved in [GS18a].

Lemma C.1 ([GS18a]) *For any integer $1 \leq p \leq 2^k - 1$, it is possible to make $O(p^{\log_2 3}) \approx O(p^{1.585})$ moves and get a pebble at position p using k pebbles.*

Using the above lemma, we now give an optimal strategy for our pebbling game.

Lemma C.2 *For any $T \in \mathbb{N}$, there exists a strategy for pebbling the line graph $[1, N]$ according to rule using at most $t + \log T$ pebbles and making $\text{poly}(T)$ moves.*

Proof The strategy is given below. For each $i \in I$ in the decreasing order **do**:

1. Use the strategy in Lemma C.1 to place a pebble in position i .
2. Recover all the other pebbles except that one in positions $\geq i$ by reversing the moves.

The correctness of this strategy follows by inspection and the number of moves is polynomial in T . ■

C.1.2 Implementing Rule A

Lemma C.3 (Rule A) *Let I and I' be two subsets of $[T]$ that satisfy the constraints given in rule A. Assuming the security of somewhere equivocal encryption, garbling scheme for circuits and updatable laconic oblivious transfer, we have that $\text{Sim}(I, P_{\setminus I}) \stackrel{c}{\approx} \text{Sim}(I', P_{\setminus I'})$.*

Proof We prove this via a hybrid argument.

- Hybrid_I: This is our starting hybrid and is distributed as $\text{Sim}(I, P_{\setminus I})$.
- Hybrid₁: In this hybrid, we change the generation of garbled program such as $(\text{st}_1, c) \leftarrow \text{SimEnc}(I', \{\widetilde{\text{SC}}_\tau\}_{\tau \notin I'})$ instead of $(\text{st}_1, c) \leftarrow \text{SimEnc}(I, \{\widetilde{\text{SC}}_\tau\}_{\tau \notin I})$

Computational indistinguishability between hybrid Hybrid_I and Hybrid₁ reduces directly to the security of somewhere equivocal encryption scheme.

- Hybrid₂: By conditions of the rule we have that $g^* - 1 \in I$. Thus, we have that $g^* - 1 \in I'$. Therefore, we note that the input labels $\{\text{lab}_{k,b}^{g^*}\}$ are not used in SimP but only in SimIn where it is used to generate $\widetilde{\text{SC}}_{g^*-1}$ and $\widetilde{\text{SC}}_{g^*}$. In this hybrid, we postpone the sampling of $\{\text{lab}_{k,b}^{g^*}\}$ and the generation of $\widetilde{\text{SC}}_{g^*}$ from SimP to SimIn.

The change in hybrid Hybrid₂ from Hybrid₁ is just syntactic and the distributions are identical.

- Hybrid₃: In this hybrid, we change the sampling of $\{\text{lab}_{k,b}^{g^*}\}$ and the generation of $\widetilde{\text{SC}}_{g^*}$. Specifically, we do not sample the entire set of labels $\{\text{lab}_{k,b}^{g^*}\}$ but a subset namely $\{\text{lab}_{k,\text{inp}_k^{g^*}}^{g^*}\}_k$ (where inp^{g^*} is the input to SC_{g^*}) and we generate $\widetilde{\text{SC}}_{g^*}$ from the simulated distribution. (Note that since $g^* - 1 \in I'$, we have that $\widetilde{\text{SC}}_{g^*-1}$ is also simulated and only $\{\text{lab}_{k,\text{inp}_k^{g^*}}^{g^*}\}_k$ are needed for its generation.) More formally, we generate

$$\widetilde{\text{SC}}_{g^*} \leftarrow \text{Sim}_{\text{ckt}}(1^\lambda, 1^{|\text{SC}|}, m, \{\text{lab}_{k,\text{inp}_k^{g^*}}^{g^*}\}_{k \in [\lambda]})$$

where m is the output of the step circuit SC_{g^*} .

The only change in hybrid Hybrid₃ from Hybrid₂ is in the generation of the garbled circuit $\widetilde{\text{SC}}_{g^*}$ and the security follows directly from the selective security of the garbling scheme.

- Hybrid₄: In this hybrid, we change how the output value m hardwired in $\widetilde{\text{SC}}_{g^*}$ is generated for the case where R/W^{g^*} is a read. In particular, we simulate the laconic OT ciphertext e . Computational indistinguishability between hybrids Hybrid₃ and Hybrid₄ follows directly from the sender privacy of the laconic OT scheme.

- Hybrid₆: In this hybrid, we change how m is generated for the case where R/W is a write. More specifically, we simulate the laconic OT write ciphertext e_w

Computational indistinguishability between hybrids Hybrid₄ and Hybrid₅ follows directly from the sender privacy for writes of the laconic OT scheme.

- Hybrid₇: In this hybrid, we reverse the changes made earlier with respect to sampling of $\{\text{lab}_{k,b}^{g^*}\}$. Specifically, we sample all values $\{\text{lab}_{k,b}^{g^*}\}$ and not just $\{\text{lab}_{k,\text{inp}_k^{g^*}}^{g^*}\}_k$. Additionally, this is done in SimP rather than SimIn.

Note that this change is syntactic and the hybrids Hybrid₆ to Hybrid₇ are identical. Finally, observe that hybrid Hybrid₇ is the same as $\text{Sim}(I', P_{\setminus I'})$.

This completes the proof of the lemma. We additionally note that the above sequence of hybrids is reversible. ■

C.1.3 Adaptive security

To prove adaptive security, we need to show that all PPT adversaries have negligible advantage in the experiment described in Figure 7. As in the previous case, we show that for any two program P, P' that have the same output in each step circuit for a particular input x , garbling of program P is computationally indistinguishable to garbling of P' . The rest follows via a standard hybrid argument. As in the previous case, we assume that the labels and the coins in KeyGen are generated randomly instead of using a PRF.

To show that real world garbling of program P is computationally indistinguishable to the real world garbling of program P' (where P and P' have the same outputs in each step circuit), we repeatedly use the following rule of indistinguishability in addition to rule A.

Rule B. Let us define a hybrid distribution Hybrid_{I, g^*} which is same as $\text{Sim}(I, P_{\setminus I})$ except that for all steps $\tau > g^*$, $C_{\text{CPU}}^{\tau, P'}$ is garbled instead of $C_{\text{CPU}}^{\tau, P}$. This rule states that for any g^* such that either $g^* - 1 \in I$ or $g^* = 1$, $\text{Hybrid}_{I, g^*} \stackrel{c}{\approx} \text{Hybrid}_{I, g^* - 1}$.

Our Hybrids, As in the previous case, our sequence of hybrids corresponds to an optimal strategy for the pebbling game described below. The rules A and B correspond to two types of moves in the pebbling game.

Pebbling Game. Consider the positive integer line $1, 2, \dots, T$. We are given a certain number pebbles **Gray** and **Black** pebbles. We can place the pebbles on this positive integer line according to the following rule:

Rule A: We can place or remove a **Gray** pebble in position i if and only if there is a **Gray** pebble in position $i - 1$. This restriction does not apply to position 1: we can always place or remove a **Gray** pebble at position 1.

Rule B: We can replace a **Gray** pebble in position i with a **Black** pebble if there is a **Gray** pebble in position $i - 1$ or if $i = 1$ and all positions $i' > i$ have **Black** pebbles.

Note that in the above game a position τ without a pebble corresponds to garbling $C_{\text{CPU}}^{\tau, P}$ and a position τ with a black pebble corresponds to garbling $C_{\text{CPU}}^{\tau, P'}$. A position with a gray pebble corresponds to simulating the step circuit implementing that time step.

Optimization goal of the pebbling game. The goal is to pebble the line $[1, T]$ such that every position in has a **Black** pebble while minimizing the number of **Gray** pebbles that are present on the line at any point in time.

The same pebbling game was considered in [GS18a] who gave an optimal strategy. We now state the following lemma from [GS18a].

Lemma C.4 ([GS18a]) *For any $T \in \mathbb{N}$, there exists a strategy for pebbling the line graph $[T]$ according to rules A and B by using at most $\log T$ **Gray** pebbles and making $\text{poly}(T)$ moves.*

Implementing rule A directly follows from Lemma C.3 and we now show how to implement rule B.

C.1.4 Implementing Rule B

Lemma C.5 *Let $I \subseteq [T]$ and $g^* \in [T]$ satisfy the constraints given in rule B. Assuming the security of selective garbled circuits and updatable laconic oblivious transfer, we have $\text{Hybrid}_{I,g^*-1} \stackrel{c}{\approx} \text{Hybrid}_{I,g^*}$.*

Proof We prove this via a hybrid argument starting with Hybrid_{I,g^*} and ending in hybrid Hybrid_{I,g^*-1} .

- Hybrid₁: In this hybrid, we change the generation of garbled program such as $(\text{st}_1, c) \leftarrow \text{SimEnc}(I \cup \{g^*\}, \{\widetilde{\text{SC}}_\tau\}_{\tau \notin I \cup \{g^*\}})$ instead of $(\text{st}_1, c) \leftarrow \text{SimEnc}(I, \{\widetilde{\text{SC}}_\tau\}_{\tau \notin I})$

Computational indistinguishability between hybrid Hybrid_1 and Hybrid_{I,g^*-1} reduces directly to the security of somewhere equivocal encryption scheme.

- Hybrid₂: By conditions of the rule we have that $g^* - 1 \in I$. Therefore, we note that the input labels $\{\text{lab}_{k,b}^{g^*}\}$ are not used in SimP but only in SimIn where it is used to generate $\widetilde{\text{SC}}_{g^*-1}$ and $\widetilde{\text{SC}}_{g^*}$. In this hybrid, we postpone the sampling of $\{\text{lab}_{k,b}^{g^*}\}$ and the generation of $\widetilde{\text{SC}}_{g^*}$ from SimP to SimIn .

The change in hybrid Hybrid_2 from Hybrid_1 is just syntactic and the distributions are identical.

- Hybrid₃: In this hybrid, we change the sampling of $\{\text{lab}_{k,b}^{g^*}\}$ and the generation of $\widetilde{\text{SC}}_{g^*}$. Specifically, we do not sample the entire set of labels $\{\text{lab}_{k,b}^{g^*}\}$ but a subset namely $\{\text{lab}_{k,\text{inp}_k}^{g^*}\}_k$ and we generate $\widetilde{\text{SC}}_{g^*}$ from the simulated distribution. (Note that since $g^* - 1 \in I$, we have that $\widetilde{\text{SC}}_{g^*-1}$ is also simulated and only $\{\text{lab}_{k,\text{inp}_k}^{g^*}\}_k$ are needed for its generation.) More formally, we generate

$$\widetilde{\text{SC}}_{g^*} \leftarrow \text{Sim}_{\text{ckt}}(1^\lambda, 1^{|\text{SC}|}, m, \{\text{lab}_{k,\text{inp}_k}^{g^*}\}_{k \in [\lambda]})$$

where m is the output of the step circuit SC_{g^*} .

The only change in hybrid Hybrid_3 from Hybrid_2 is in the generation of the garbled circuit $\widetilde{\text{SC}}_{g^*}$ and the security follows directly from the selective security of the garbling scheme.

- Hybrid₄: In this hybrid, we change how the output value m hardwired in $\widetilde{\text{SC}}_{g^*}$ is generated for the case where $\text{R}/\text{W}^{i,g^*}$ is a read. In particular, we simulate the laconic OT ciphertext e . Computational indistinguishability between hybrids Hybrid_3 and Hybrid_4 follows directly from the sender privacy of the laconic OT scheme.
- Hybrid₅: In this hybrid, we change how m is generated for the case where R/W is a write. More specifically, we simulate the laconic OT write ciphertext e_w . Computational indistinguishability between hybrids Hybrid_4 and Hybrid_5 follows directly from the sender privacy for writes of the laconic OT scheme.
- Hybrid₆: In this hybrid, we change how the output value m hardwired in $\widetilde{\text{SC}}_{g^*}$. In particular, we change it to be the output of C_{CPU}^{P,g^*} instead of C_{CPU}^{P',g^*} . Hybrid_5 and Hybrid_6 are identically distributed since the outputs of these two step circuits are exactly the same.

- Hybrid₇ to Hybrid₁₁: In this hybrid, we reverse the changes made earlier with respect to the hybrids Hybrid₅ to Hybrid₁ and use C_{CPU}^{P,g^*} to generate \widetilde{SC}_{g^*} .

Note that Hybrid₁₁ is distributed identically to Hybrid_{I,g*}. This completes the proof of the lemma. ■

D Proof of Indistinguishability

We assume without loss of generality, that the adversary first provides the number of queries that it makes to the garbled program and the garbled input oracle. This assumption is without loss of generality since given any adversary with running time t , we can turn it into an adversary that makes t queries (where some of them are for dummy programs and inputs) to the garbled program and the garbled input oracle.

We first give the description of the simulator in Figure 13.

D.1 Proof of Indistinguishability

We need to show that the advantage of the adversary in the experiment described in Figure 4 is negligible. To prove this, we show that the real world garbling procedure is computationally indistinguishable to the simulated garbling for a single program. The rest follows via a hybrid argument where we repeatedly use this argument from the last (program, input) tuple to the first (program, input) tuple. For the hybrid argument to go through, we additionally need the intermediate hybrid distributions to know the number of (program, input) tuples that the adversary queries for. We assumed without of generality that the adversary provides this before the start of the experiment.

In the rest of this subsection, we show that the real world garbling procedure is indistinguishable to the simulated garbling for a single (program, input) tuple. We assume that in the rest of the hybrids the output of the PRF_S is replaced with a random string. This assumption follows directly from the security of the PRF. We define an intermediate hybrid distribution Hybrid _{t} as follows.

Hybrid _{t} : In this hybrid, we generate the garbled memory, garbled program and the garbled input as follows.

Garbled Memory: On input the security parameter and the database D , the garbled memory is generated exactly as in `GRAM.Memory` given in Figure 10.

Garbled Program: On input the program P , the garbled program \widetilde{P} is generated as follows:

1. Sample $K' \leftarrow \text{PP.KeyGen}(1^\lambda)$
2. $P^* \leftarrow \text{OProg}(1^\lambda, 1^{\log M}, 1^T, P)$ where P^* runs in time T' .
3. For each $\tau \in [t-1]$, compute $K[(i||\tau)] \leftarrow \text{TE.Constrain}(K, (i||\tau))$ where $(i||\tau)$ is expressed as a λ -bit string.
4. Choose $r \leftarrow \{0, 1\}^n$.
5. **for** each $\tau \in [t, T']$ **do**:
 - (a) Choose $L^\tau \leftarrow [M']$ where M' is the number of blocks in \widetilde{D} .

SimD($1^\lambda, 1^N, 1^M$) : On input the security parameter, 1^N and 1^M do:

1. Sample $K \leftarrow \text{TE.KeyGen}(1^\lambda)$ and $S \leftarrow \text{PRFKeyGen}(1^\lambda)$ defining $\text{PRF}_S : \{0, 1\}^\lambda \rightarrow \{0, 1\}^n$.
2. Initialize an empty array \widehat{D} of M blocks with block length N .
3. **for** each i from 1 to M **do**:
 - (a) Set $\widehat{D}[i] \leftarrow \text{TE.Enc}(K, 0^\lambda, 0)$.
4. $D^* \leftarrow \text{OData}(1^\lambda, 1^N, \widehat{D})$.
5. $(\widetilde{D}, SK) \leftarrow \text{UGRAM.Memory}(1^\lambda, 1^n, D^*)$
6. Output \widetilde{D} as the garbled memory and **st** as (K, S, SK) .

SimP($1^\lambda, i, 1^{|P_i|}, \text{st}$): On **st** and the running time T do:

1. Let T' be the running time of the program after compiling with an ORAM scheme.
2. Compute $r = \text{PRF}(S, i)$.
3. **for** each $\tau \in [T']$ **do**:
 - (a) Choose $L^\tau \leftarrow [M']$ where M' is the number of blocks in \widetilde{D} .
 - (b) Sample $\text{R/W}^\tau \leftarrow \{\text{read}, \text{write}\}$.
 - (c) Choose $X^\tau \leftarrow \text{TE.Enc}(K, (i \parallel \tau), 0)$ and $c_{\text{CPU}}^\tau = \text{TE.Enc}(K, (i \parallel \tau), 0^n)$.
 - (d) Let $C_{\text{CPU}}^\tau := \text{SC}'_\tau[i, \tau, c_{\text{CPU}}^\tau, X^\tau, L^\tau, \text{R/W}^\tau, r']$ where $r' = r$ if $\tau = T'$; else $r' = \perp$. The step-circuit SC' is described below.
4. Construct a RAM program P' with step-circuits given by $\{C_{\text{CPU}}^\tau\}$.
5. $\widetilde{P} \leftarrow \text{UGRAM.Program}(SK, i, P')$.
6. Output \widetilde{P} .

SimIn(y, st) : On input y and **st** do:

1. Output $\widehat{x} \leftarrow \text{UGRAM.Input}(SK, i, 0^n)$ and $y \oplus \text{PRF}_S(i)$.

Step Circuit SC'_τ

Input: The CPU state **state** and a data block $X \in \{0, 1\}^N$.

Hardcoded: The sequence number i , step number τ , the ciphertext $c_{\text{CPU}}^\tau, X^\tau, \text{R/W}^\tau$, the location L^τ and the string r' .

1. **if** $\tau = T'$, then output $c'_{\text{CPU}} = r'$; else $c'_{\text{CPU}} = c_{\text{CPU}}^\tau$.
2. **if** $\text{R/W}^\tau = \text{write}$ **do**:
 - (a) Output $(\text{state}, \text{R/W}^\tau, L^\tau, X^\tau)$.
3. **else if** $\text{R/W}^\tau = \text{read}$, output $(\text{state}, \text{R/W}^\tau, L^\tau, \perp)$.

Figure 13: Simulator for Adaptive Security

- (b) Sample $R/W^\tau \leftarrow \{\text{read}, \text{write}\}$.
- (c) Choose $X^\tau \leftarrow \text{TE.Enc}(K, (i||\tau), 0)$ and $c_{\text{CPU}}^\tau = \text{TE.Enc}(K, (i||\tau), 0^n)$.
- (d) Let $C_{\text{CPU}}^\tau := \text{SC}'_\tau[i, \tau, c_{\text{CPU}}^\tau, X^\tau, L^\tau, R/W^\tau, r']$ where $r' = r$ if $\tau = T'$; else $r' = \perp$.
- 6. Let τ_1, \dots, τ_m be the sequence of values guaranteed by strong localized randomness.
- 7. **for** each $\tau \in [t-1]$ **do**:
 - (a) Let $j \in [m-1]$ be such that $\tau \in [\tau_j + 1, \tau_{j+1}]$.
 - (b) Let $C_{\text{CPU}}^\tau := \text{SC}_\tau[i, \tau, K[(i||\tau)], I_j, K']$
- 8. Construct a RAM program P' with step-circuits given by $\{C_{\text{CPU}}^\tau\}$.
- 9. $\tilde{P} \leftarrow \text{UGRAM.Program}(SK, i, P')$.
- 10. Output \tilde{P} .

Garbled Input: On input $SK' = (K, SK)$, i and x and the output y :

- 1. Compute $\hat{x} \leftarrow \text{UGRAM.Input}(SK, i, x)$.
- 2. Output \hat{x} and $y \oplus r$ if $t < T' + 1$. Otherwise, output \hat{x} and r .

Note that $\text{Hybrid}_{T'+1}$ is distributed identically to the real world garbling procedure. We now show that Hybrid_t is computationally indistinguishable to Hybrid_{t-1} for every $t \in [2, T' + 1]$.

Lemma D.1 *Assuming the security of adaptive garbled RAM with unprotected memory access, timed encryption and puncturable PRF we have $\text{Hybrid}_t \stackrel{c}{\approx} \text{Hybrid}_{t-1}$ for every $t \in [2, T' + 1]$.*

Proof We prove via a hybrid argument.

Hybrid_{t,1} : In this hybrid, we change how the garbled program and the garbled input are generated. In particular, we will use the simulator uSim for the equivocal security of UGRAM to generate them. More formally, we generate

$$\tilde{P} \leftarrow \text{uSim}(\{C_{\text{CPU}}^\tau\}_{\tau \in [T'] \setminus \{t-1\}})$$

$$\tilde{x} \leftarrow \text{uSim}(x, y_{t-1})$$

where y_{t-1} is the output of the C_{CPU}^{t-1} on input x .

The computational indistinguishability between Hybrid_t and $\text{Hybrid}_{t,1}$ follows from the equivocal security of adaptive garbled RAM with unprotected memory access.

Hybrid_{t,2} : In this hybrid, we change how y_{t-1} is generated. In particular, if R/W^{t-1} is write we choose $X' \leftarrow \text{TE.Enc}(K[i||(t-1)], (i, (t-1)), 0)$ in y_{t-1} instead of setting $X' \leftarrow \text{TE.Enc}(K[i||(t-1)], (i, (t-1)), \text{wData}^{t-1})$. The computational indistinguishability between $\text{Hybrid}_{t,1}$ and $\text{Hybrid}_{t,2}$ follows directly from the security of timed encryption scheme.

Hybrid_{t,3} : In this hybrid, we change how the ciphertext encrypting the state is generated in y_{t-1} . In particular, if $t < T' + 1$, we choose $c_{\text{CPU}} \leftarrow \text{TE.Enc}(K[i||(t-1)], (i, (t-1)), 0^n)$ in y_{t-1} instead of setting it to be $\text{TE.Enc}(K[(i||\tau)], \text{state}')$. If $t = T' + 1$, we change $c_{\text{CPU}} = r$ and output $y \oplus r$ as

part of the garbled input. The indistinguishability between $\text{Hybrid}_{t,2}$ to $\text{Hybrid}_{t,3}$ reduces directly to the security of the timed encryption scheme if $t < T' + 1$. If $t = T' + 1$, $\text{Hybrid}_{t,2}$ and $\text{Hybrid}_{t,3}$ are identically distributed.

$\text{Hybrid}_{t,4}$: In this hybrid, we change C_{CPU}^{t-1} to a dummy step circuit that generates encryption of the state and the block to be written as in $\text{Hybrid}_{t,3}$. However, the memory access (R/W, L) made by this step circuit is still computed using C_{CPU}^{t-1} . It follows from the equivocal security of the adaptive garbled RAM with unprotected memory access that $\text{Hybrid}_{t,3}$ is computationally indistinguishable to $\text{Hybrid}_{t,4}$.

$\text{Hybrid}_{t,5}$: Let $j \in [m - 1]$ be such that $\tau_j \leq t - 1 < \tau_{j+1}$. Depending on the program input, there exists a $k < j$ such that property 3 in the definition of strong localized randomness holds. In this hybrid, we guess such a k (we are successful with probability $1/m$) while generating the garbled program. In this hybrid, we change how the garbled program is generated. Instead of garbling the program as in $\text{Hybrid}_{t,4}$, we will construct another program that will have the same output in each step circuit as the program in $\text{Hybrid}_{t,4}$. Let us explain the construction of this new program.

We puncture the PRF key at I_k and I_j and hardwire the output of the PRF on input I_k in the step circuits $\tau_k \leq t - 1 < \tau_{k+1}$ and the output of the PRF on input I_j in the step circuits $\tau_j \leq t - 1 < \tau_{j+1}$. We use the punctured key in the rest of the step circuits. The step circuits for timesteps $\tau \in [\tau_j, \tau_{j+1})$ and $\tau \in [\tau_k, \tau_{k+1})$ additionally have output of the PRF on I_j and I_k hardwired. This value is used in the computation of the memory access (R/W, L).

Note that it follows from the correctness property of the puncturable PRF that the programs garbled in $\text{Hybrid}_{t,4}$ and $\text{Hybrid}_{t,5}$ have the same output in each step circuit for each input x . It now follows from the adaptive security of the garbled RAM with unprotected memory access that $\text{Hybrid}_{t,4}$ is computationally indistinguishable to $\text{Hybrid}_{t,5}$.

$\text{Hybrid}_{t,6}$: In this hybrid, we replace the hardwired PRF outputs on I_j and I_k with random strings. Computational indistinguishability between $\text{Hybrid}_{t,5}$ and $\text{Hybrid}_{t,6}$ follows from the security of puncturable PRFs.

$\text{Hybrid}_{t,7}$: In this hybrid, we use the uSim to simulate the step circuits for time steps $\tau \in [\tau_k, \tau_{k+1}) \cup [\tau_j, \tau_{j+1})$. The indistinguishability between $\text{Hybrid}_{t,6}$ and $\text{Hybrid}_{t,7}$ follows from the equivocal security.

$\text{Hybrid}_{t,8}$: In this hybrid, we change the memory access (i.e., (R/W, L)) in y_{t-1} to be random. It follows from the strong localized property of ORAM scheme that this change is computationally indistinguishable.

$\text{Hybrid}_{t,9}$: In this hybrid, we change C_{CPU}^{t-1} to a dummy circuit that outputs y_{t-1} as generated in $\text{Hybrid}_{t,8}$ and the rest of the circuits for the timesteps $\tau \in [\tau_k, \tau_{k+1}) \cup [\tau_j, \tau_{j+1})$ as in $\text{Hybrid}_{t,6}$. The indistinguishability of $\text{Hybrid}_{t,8}$ and $\text{Hybrid}_{t,9}$ follows from the equivocal security of adaptive garbled circuits with unprotected memory access.

$\text{Hybrid}_{t,10}$: In this hybrid, we reverse the change made in $\text{Hybrid}_{t,6}$, namely, we replace the hard-

wired random strings in the step circuits for time steps τ , where $\tau_k \leq \tau < \tau_{k+1}$ and $\tau_j \leq \tau < \tau_{j+1}$ with the PRF output on I_k and I_j . It follows from the security of puncturable PRF that $\text{Hybrid}_{t,9}$ is computationally indistinguishable to $\text{Hybrid}_{t,10}$.

$\text{Hybrid}_{t,11}$: This hybrid is distributed identically to Hybrid_{t-1} . Note that the only difference between $\text{Hybrid}_{t,10}$ and $\text{Hybrid}_{t,11}$ is that PRF key K' is punctured at intervals I_k and I_j in $\text{Hybrid}_{t,10}$ whereas it not punctured in $\text{Hybrid}_{t,11}$ and the PRF output is computed to obtain the random coins for the memory access. The indistinguishability between $\text{Hybrid}_{t,10}$ and $\text{Hybrid}_{t,11}$ follows directly from the adaptive security of garbled RAM with unprotected memory access since the two programs have the same output in each step circuit for each input x .

This completes the proof of the lemma. ■

Hybrid_0 : In this hybrid, we change how the garbled database is generated. In particular, instead of garbling the original database, we garble a dummy database whose initial contents are set to the all zeroes string. It follows from the security of the timed encryption scheme that Hybrid_0 is computationally indistinguishable to Hybrid_1 . Note that Hybrid_0 is identically distributed to the simulated distribution.