

Black-Box Garbled RAM

Sanjam Garg*

Steve Lu[†]Rafail Ostrovsky[‡]

Abstract

Garbled RAM, introduced by Lu and Ostrovsky, enables the task of garbling a RAM (Random Access Machine) program directly, there by avoiding the inefficient process of first converting it into a circuit. Garbled RAM can be seen as a RAM analogue of Yao's garbled circuit construction, except that known realizations of Garbled RAM make non-black-box use of the underlying cryptographic primitives.

In this paper we remove this limitation and provide the first black-box construction of Garbled RAM with polylogarithmic overhead. Our scheme allows for garbling multiple RAM programs being executed on a persistent database and its security is based only on the existence of one-way functions. We also obtain the first secure RAM computation protocol that is both constant round and makes only black-box use of one-way functions in the OT-hybrid model.

*University of California, Berkeley. Computer Science Division. sanjamg@berkeley.edu

[†]University of California, Los Angeles. Department of Computer Science. stevelu@cs.ucla.edu

[‡]University of California, Los Angeles. Department of Computer Science and Mathematics. rafail@cs.ucla.edu

1 Introduction

Alice wants to store a large private database D on the cloud in an encrypted form. Subsequently, Alice wants the cloud to be able to compute and learn the output of arbitrary dynamically chosen private programs P_1, P_2, \dots on private inputs x_1, x_2, \dots and the previously stored dataset, which gets updated as these programs are executed. Can we do this?

The above problem is a special case of the general problem of secure computation [Yao82, GMW87]. In the past three decades of work, both theoretical and practical improvements have been pushing the limits of the overall efficiency of such schemes. However most of these constructions work only for circuits and securely computing a RAM program involves the inefficient process of first converting it into a circuit. For example, Yao’s approach requires that the program be first converted to a circuit — the size of which will need to grow at least with the size of the input. Hence, in the example above, for each program that Alice wants the cloud to compute, it will need to send a message that grows with the size of the database. Using fully homomorphic encryption [Gen09] we can reduce the size of Alice’s message, but the cloud still needs to compute on the entire encrypted database. Consequently the work of the cloud still grows with the size of the database. These solutions can be prohibitive for various applications. For example, in the case of binary search the size of the database can be exponentially larger than execution path of the insecure solution. In other words security comes at the cost of an exponential overhead. We note that additionally even in settings where the size of the database is small, generic transformations from RAM programs with running time T result in a circuit of size $O(T^3 \log T)$ [CR73, PF79], which can be prohibitively inefficient.

Secure computation for RAM programs. Motivated by the above considerations, various secure computation techniques that work directly for RAM programs have been developed. However all known results have interesting theoretical bottlenecks that influence efficiency, either in terms of *round complexity* or in their *non-black-box* use of cryptographic primitives.

- For instance, Ostrovsky and Shoup [OS97] show how general secure RAM computation can be done using oblivious RAM techniques [Gol87, Ost90, GO96]. Subsequently, Gordon et al. [GKK⁺12] demonstrated an efficient realization based on specialized number-theoretic protocols. In follow up works, significant asymptotic and practical efficiency improvements have been obtained by Lu and Ostrovsky [LO13a] and by Wang et al. [WHC⁺14]). However, all these works require round complexity on the order of the running time of the program.
- In another recent line of work [LO13b, GHL⁺14, GLOS14], positive results on round efficient secure computation for RAM programs have been achieved. However all these constructions are inherently non-black-box in their use of cryptographic primitives.¹ These improvements are obtained by realizing the notion of garbled random-access machines (garbled RAMs) [LO13b] as a method to garble RAM programs directly, a RAM analogue of Yao’s garbled circuits [Yao82].

In particular, we use the notation $P^D(x)$ to denote the execution of some RAM program P on input x with initial memory D . A garbled RAM scheme should provide a mechanism to garble the data D into garbled data \tilde{D} , the program P into garbled program \tilde{P} and the input x into garbled input \tilde{x} such that given \tilde{D} , \tilde{P} and \tilde{x} allows for computing $P^D(x)$ and nothing more. Furthermore, up to only poly-logarithmic factors in the running time of the RAM $P^D(x)$ and the size of D , we require that the size of garbled data \tilde{D} is proportional only to the size of data D , the size of the garbled input \tilde{x} is proportional only to that of x and the size and the evaluation time of the garbled program \tilde{P} is proportional only to the running time of the RAM $P^D(x)$.

¹We note that several other cutting-edge results [GKP⁺13, GHRW14, CHJV14, BGT14, LP14] have been obtained in non-interactive secure computation over RAM programs but they all need to make non-black-box use of cryptographic primitives. Additionally some of these construction are based on strong computational assumptions such as [Reg05, GGH⁺13b, GGH13a]. We skip discussing this further and refer the reader to [GLOS14, Appendix A] for more details.

Starting with Impagliazzo-Rudich [IR90, IR89], researchers have been very interested in realizing cryptographic goals making just black-box use of underlying primitive. It has been the topic of many important recent works in cryptography [IKLP06, PW09, Wee10, GLOV12]. On the other hand, the problem of realizing black-box construction for various primitive is still open, e.g. multi-statement non-interactive zero-knowledge [BFM88, FLS99, GOS06] and oblivious transfer extension [Bea96].² From a complexity perspective, black-box constructions are very appealing as they often lead to conceptually simpler and qualitatively more efficient constructions.³

Motivated by low round complexity and black-box constructions, in this work, we ask if we can achieve the best of both worlds. In particular:

Can we construct garbled RAM programs with only polylogarithmic overhead making only black-box use of cryptographic primitives?

1.1 Our Results

In this paper, we provide the first construction of a fully black-box garbled RAM, i.e. both the construction and the security reduction make only black-box use of underlying cryptographic primitives (one-way functions). The security of our construction is based on the existence of one-way functions alone. We state this as our main theorem:

Main Theorem (Informal). *Assuming only the existence of one-way functions, there exists a secure black-box garbled RAM scheme, where the size of the garbled database is $\tilde{O}(|D|)$, size of the garbled input is $\tilde{O}(|x|)$ and the size of the garbled program and its evaluation time is $\tilde{O}(T)$ where T is the running time of program P . Here $\tilde{O}(\cdot)$ ignores $\text{poly}(\log T, \log |D|, \kappa)$ factors where κ is the security parameter⁴. Since garbled RAM implies one-way functions, this can be stated as an equivalence of primitives.*

Just as in previous works on garbled RAM [LO13b, GHL⁺14, GLOS14], our construction allows for maintaining persistent database across execution of multiple programs on the garbled memory. Also as in [GKP⁺13, LO13b, GHL⁺14, GLOS14], if one is willing to disclose the exact running time of a specific execution, then the running time of a garbled RAM computation can be made input specific which could be much faster than the worst-case running time.

Secure RAM computation. We also obtain the first one-round secure computation protocol for RAM programs that makes only black-box use cryptography in the OT-hybrid model. A very unique feature of this construction is that it allows for asymmetric load in terms of storage costs, i.e., only one party stores the encrypted database. To the best of our knowledge no previous solutions allowed for an encrypted database based of private information of both parties to be stored on just one party, and yet allow secure RAM computation on it using black-box methods alone. This makes our constructions particularly relevant in the context of secure outsourced computation. Our garbled circuit generation algorithms all simply rely on a key for a pseudo-random function, and therefore can be also outsourced and generated by an external party holding the key in a manner similar to the work of Ananth et al. [ACG⁺14].

2 Our Techniques

We start by recalling briefly the high level idea behind the previous garbled RAM constructions [LO13b, GHL⁺14, GLOS14] and its follow up works. This serves as a good starting point in explaining the technical challenges that come up in realizing garbled RAM making only *black-box* use of cryptographic primitives.

²Interestingly for oblivious transfer extension we do know black-box construction based on stronger assumptions [IKNP03].

³Additionally, black-box constructions enable implementations agnostic to the implementation of the underlying primitives. This offers greater flexibility allowing for many optimizations, scalability, and choice of implementation.

⁴Although it is typically assumed that κ is polynomially related to M , one can redefine the security parameter to be as small as $\omega(1) \log^2 M$ and still efficiently achieve correctness and security that is all but negligible in M directly.

What makes black-box garbled RAM hard? We view the program P , to be garbled, as a sequence of T CPU steps. Each of these CPU steps is represented as a circuit. Each CPU step reads or writes one bit of the RAM, which stores some dataset D (that can grow dynamically at run-time though for simplicity we consider a D with M data elements). At a high level, known garbled RAM construction proceed in two steps. First a garbled RAM scheme is constructed under the weaker security requirement of unprotected memory access (UMA) in which we do not try to hide the database being stored or the memory locations being accessed (only the program and input is hidden). Next this weaker security guarantee is amplified to get full security by using oblivious RAM. Both these steps introduce a non-black-box use of cryptographic primitives. Besides some technical details, the second step can actually be made black-box just by using statistical oblivious RAM [Ajt10, DMN11, SCSL11, SvDS⁺13, CLP14], though these statistical schemes do not protect the memory contents which will need to be addressed. Next we describe the challenges we need to overcome in realizing a black-box construction with UMA security.

At a very high level, known garbled RAM constructions with UMA security construct the garbled memory in the following way. For each memory location i , containing value b_i the value $F_s(i, b_i)$ is stored in the “garbled” memory, where s is a secret key for the PRF F . Let’s consider that a CPU step that wants to read memory location i that needs to be fed into the next CPU step. Note that both these CPU step circuits will be independently garbled using Yao’s garbled circuit technique. Let label^0 and label^1 be the garbled input wire labels corresponding to the wire for the read bit of the second circuit. In order to enable evaluation of the second garbled circuit, we need to reveal *exactly one* of these two labels, corresponding to b_i , to the evaluator. Note that the first garbled circuit needs to do this without knowing i and b_i at the time of garbling. The idea for enabling the read is for the first garbled circuit to produce a translation gadget: the first garbled circuit outputs encryptions of labels label^0 and label^1 under keys $F_s(i, 0)$ and $F_s(i, 1)$ respectively. Since the evaluator holding the garbled memory only has one of the two values $F_s(i, 0)$ or $F_s(i, 1)$ at his disposal, he can only obtain either label^0 or label^1 . This enables the evaluator to feed the proper bit into the next CPU step and continue the evaluation. Since the location to be read, i , is generated dynamically at run time the values $F_s(i, 0)$ and $F_s(i, 1)$ must be computed inside the garbled circuit. This is exactly where non-black-box use of the PRF is made.

More generally, there appears to be a fundamental barrier in this technique. Note that the data in the memory has to be stored in some encrypted form. In the above case it was stored as the output of a PRF evaluation. Whenever a bit is to be read from memory, it will need to be internally decrypted to recover what the value being read is, this makes the need for non-black-box use of cryptography essential. In fact, if we limit ourselves to black-box constructions then we do not know any garbled RAM solutions that are any better than the trivial solution of converting the RAM program to a circuit.

Our starting idea: dynamic memory. Our starting idea to recover from the above problem is to replace “static” memory made of various ciphertexts, as has been done in previous works; with a “dynamic” memory consisting of various garbled circuits. More specifically, in our new solution, we envision reading from “memory” in a new way. Unlike previous work, we envision reads from memory are achieved by passing the control flow of the program itself to the memory. We explain how garbled memory is activated: a small number of garbled circuits in the garbled memory are evaluated in sequence. These garbled circuits are connected in the sense that one garbled circuit will output a valid garbled input for the next garbled circuit by having both the zero and one labels hardwired within it. Eventually the control reaches back to the main program and additionally carries the read bit along with. Note that the actual garbled circuits that are fired inside the garbled memory are dynamically decided based on the location being read. Looking ahead, in our final construction they will also depend on the previous state of the garbled memory. We describe this later.

Next, we describe a plausible arrangement of garbled circuits in the garbled memory for realizing the above intuition. Let M be the size of memory. Imagine a tree of garbled circuits where the root garbled circuit has hardcoded in it the input labels for both its children. Similarly the left child garbled circuit has

the input labels of its two children and so on. Finally the leaf garbled circuits, which are M in number, are such that each contains a bit of the database hardcoded in them. Our root garbled circuit takes as input two labels label^0 and label^1 and a location to be read. The root garbled circuit based on the location that needs to be read can activate its left or its right child garbled circuit, ultimately leading to the activation of a leaf garbled circuit which outputs either label^0 or label^1 based on whether the stored bit in it is 0 or 1. This enables a black-box way of reading a bit from memory.

However, the key challenge in realizing the above setup is that after only one read a sequence of garbled circuits from the root to a leaf in the garbled memory have been consumed/destroyed. Hence if we are to continue using the garbled memory further then we must provide “replacements” for the used garbled circuits. To get better insight into the issues involved, we start by describing a very natural dynamic replacement solution for this problem which actually fails because of rather subtle reasons. We believe that this solution highlights the technical issues involved.

Providing “replacements” dynamically. Our first attempted solution for overcoming the above challenge is to provide “generic” garbled circuits that can replace the specific garbled circuits that are used during a read. As mentioned earlier, during a read, garbled circuits corresponding to a path from the root to a leaf are fired and in the process consumed. It is exactly these circuits that we need to replace. So corresponding to every read we could provide a sequence of garbled circuits to exactly replace these consumed garbled circuits. These replacement garbled circuits could be prepared to have the input labels of their new child already hardcoded in them, though some information needs to be provided at run-time.

Unfortunately this attempted approach has a very subtle bug — relating to a circularity in the parameter sizes. The problem is that the size of the additional inputs of the “replacements” provides the input labels for the “regular” input wires, but this information must be passed by the “regular” wires. In other words, if this scheme is to work, then we need to have the “replacement” garbled circuits be smaller than the garbled circuit that are being consumed which will lead to a blow up in the size of the first circuit. This appears to be a fundamental problem with this approach. We believe that black-box techniques cannot be used to fix this problem if only dynamic “replacements” are provided.

Providing “replacements” statically. Our second stab at the problem is to include for each node of the tree not just one garbled circuit but instead a sequence of garbled circuits. Of course we still need to respect the relationship that each garbled circuit needs to have the ability to activate its left and right child garbled circuits. Now that we have a sequence of garbled circuits for every node it is not clear which garbled circuits in its children sequences should a garbled circuit be connected with. A very simple strategy would be to have T garbled circuits in each sequence corresponding to each node, where T is the number of reads the garbled memory is designed for. We can connect the i^{th} garbled circuit in each sequence with the i^{th} garbled circuit in its children sequences. However, this leads to a garbled memory of size $T \cdot M$, something that is much larger than what we want, and defeating the purpose of this approach.

It is clear that if we want to have T reads then we must have at least T garbled circuits at the root node but we can hope to have fewer garbled circuits in the children sequences, finally having much fewer garbled circuits in any sequence at the leaf nodes. We will next describe a strategy for generating and connecting these garbled circuits such that the overall number of garbled circuits needed grows only linearly in $T + M$. The key insight is to notice that if the access pattern were uniformly random, each node will be visited half as often as the parent, in expectation. For now, we make this simplifying assumption of uniformity and it later turns out that when we apply statistical Oblivious RAM, we almost get this property for free.

Consider the following attempt, the root node will have a sequence of T garbled circuits and its children sequences will have $T/2$ garbled circuits with the leaf nodes sequences containing T/M garbled circuits. Since the reads in memory are uniform, at least in expectation we can expect that having these many garbled circuits should suffice. Here we will connect the i^{th} garbled circuit in any sequence with the $i/2^{\text{th}}$ garbled circuits in its left and right children sequences. The key technical problem in this solution is that even with

uniform access, there is no strong guarantee of “balance” between the nodes. During the execution, the garbled circuits in a parent node would need to know exactly which child circuit to activate, but this is not known in advance and will depend on the read patterns. So instead of having just one left and one right child key hardcoded in a garbled circuit we will have hardware keys for a window of garbled circuits in its left and right child sequences. Unfortunately the windows sizes needed for this solution is too large. In particular, if we have a total of T reads happen at the root node then with a constant probability we expect a discrepancy of \sqrt{T} between the minimum number and maximum number of reads that go left. Similarly, the right child has the same issue. This means that we now need a window of size \sqrt{T} which is, which is still prohibitively large.

Defeating imbalances — having more circuits and fast-forwarding. Our main idea for dealing with the imbalances (which causes large window size) is to have more circuits in every sequence for each node. Indeed, these extra circuits serve two purposes. First, these extra circuits serve as a buffer in case we go beyond expectation. Secondly, when we are too far behind expectation, these extra circuits will be consumed *faster* to enforce that we are always within the window of keys that the parent node has. The key insight is that instead of having a fixed additive factor, the child pointers dynamically moves beyond the expectation (and enough standard deviations to achieve exponentially small probability of failure) *relative* to the current location. As such, in earlier time steps there is less of a “stretch”, whereas in later time steps there is more stretch. This resolves the tension between having too many stretch circuits yet still having enough to make sure you do not run out.

More specifically, our goal is to shrink the key window size down from \sqrt{T} to a value that grows only with the security parameter. We shrink the window size using the following strategy: keep the window always well ahead of number of garbled circuits that could possibly be consumed (by the Chernoff bound), and provide a method to move into the window when lagging behind. In order to achieve this we introduce two new ideas, which when combined accomplish this strategy. First, each circuit has the option to “fast forward” or “burn” by passing on data to its *successor* circuit in the same node, specifically the next garbled circuit in the sequence. This “fast forwarding” is enough to ensure that the children garbled circuits always remain in the appropriate window. The second idea allows a parent garbled circuit to be able to evaluate old circuits that have fallen out of this window. Note that we only need the parent to be able to evaluate a *single* old circuit since whenever the child node is activated it will burn garbled circuits within its own sequence pushing it back into the window. Thus, we pass from circuit to circuit the keys to the first unused left and right child garbled circuits (and when consumed, will be replaced by the next key inside the window).

Complications arise as this argument is needed for each node of the tree. We need to give more circuits than the expectation of a half to each child, but this causes a chain reaction as now each of them need to give their children another push beyond expectation and so forth. In particular, we need to provide a larger and larger ratio of extra garbled circuits away from expectation as we go down the tree. This causes a tension in the parameters: even though the number of circuits is roughly halving each time we go down, the factor by which we must push it back up by is also growing geometrically. Setting this growth rate to be even constant, say $c > 1$, is problematic. We run into difficulties since there will be $T \frac{1}{2}^i c^i$ circuits per node, and 2^i nodes per level, all the way up to $\log M$ — resulting in $TM^{\log c}$ circuits, which is polynomial overhead in the size of the dataset. It turns out that even a very slow growth rate suffices and allows up to get desired efficiency properties. With a careful analysis, it turns out this is both efficient and will only run out of garbled circuits with negligible probability. The exact details will become apparent in the construction and proof.

Getting provable security. Although the previous techniques achieve correctness and efficiency, it is not immediately obvious why it should be secure. Indeed, inputs keys for one garbled circuits are actually hardwired at multiple places, and if we do not carefully account for all of these locations, we could be running into circularity issues in our solution when using security of garbled circuits. To accommodate this, we will have a key hardwired only at the *first* place it was expected to be hardwired and “passed”

dynamically to later circuits. Each garbled circuit will still maintain the same window of keys as were available to it earlier but now they are dynamically passed by it to its successor (the next garbled circuit in a sequence). Moving forward, new keys are collected and the old keys will be dropped so that the total number of keys being passed remains small. Using this mechanism we can ensure that a garbled circuit can dynamically “drop” keys that correspond to a child garbled circuit whose security needs to be relied on in the proof. Based on this, in the hybrid argument, we argue that whenever a garbled circuit is replaced by a simulated version, we have all instances of its keys have been “dropped.”

A technical issue also arises with the fact at the end some of the unevaluated garbled circuits remain. The problem lies in the fact that some of their input keys have also been revealed. We handle this issue by providing a generic transformation that ensures that garbled circuits are indistinguishable from noise as long as input keys for even one wire are not disclosed.

Final touches. In explaining the technical ideas above made several simplifying assumptions. We now provides some ideas on how to remove these limitations.

- **Arbitrary Memory Access.** As mentioned above, we can achieve a GRAM solution for programs with arbitrary access pattern by first compiling it with an ORAM that has uniform access pattern. Programs compiled with statistical ORAM do not actually have uniform memory access, but rather a *leveled* uniform access pattern, where the accesses in each level is uniformly distributed. We deal with this technicality by breaking our memory down into levels where access in each separate memory is uniform. Alternatively, we can *bias* the distribution where a leveled ORAM structure is flattened into memory: for example, we know that the a memory location corresponding to some tree node is accessed twice as often as its children, thus when we build our circuits we can incorporate and absorb this distribution into our scheme.
- **Replenishing.** Since we generate a fixed amount of garbled circuits for the garbled memory, this places a bound on the number of reads the memory can be used for. We observe that is we sent the number of reads to be equal the size of memory then this give us enough reads in parallel to with the entire memory can be replenished to allow for another size of memory number of reads and so on. In our construction the garbled circuits are generated in a highly independent fashion and so more garbled circuits can be provided on the fly. Furthermore, this can be seamlessly amortized (the amortized overhead can be absorbed into the polylog factors) where the garbling algorithm for a T -time program can generate enough garbled circuits to support T more steps in the future. Finally, this strategy can also be used to accommodate memory that is dynamically growing.
- **Writing.** Writing in our construction is achieved in a way very similar to the reading. Reading in our scheme involves having a leaf garbled circuit pass on the value to the main circuit and simultaneously pass on the stored data value in it to its successor, so that it could be read again. During writing a garbled circuits passes the value to be written to its successor instead of the value previously stored.

2.1 Roadmap

We now lay out a roadmap for the remainder of the paper. In Section 3, we give necessary background and definitions for the RAM model, garbled circuits, and garbled RAM. In Section 4 we give the warmup heuristic construction of our result. We analyze the cost and correctness of the solution in Section 5. We extend our construction to a secure one in Section 6 and prove the security in Section 7 (with the full proof in Section 8).

3 Background

In this section we fix notation for RAM computation and provide formal definitions for Garbled Circuits and Garbled RAM Programs. Parts of this section have been taken verbatim from [GHL⁺14].

3.1 RAM Model

Notation for RAM Computation. We start by fixing the notation for describing standard RAM computation. For a program P with memory of size M we denote the initial contents of the memory data by $D \in \{0, 1\}^M$. Additionally, the program gets a “short” input $x \in \{0, 1\}^n$, which we alternatively think of as the initial state of the program. We use the notation $P^D(x)$ to denote the execution of program P with initial memory contents D and input x . The program can P read from and write to various locations in memory D throughout its execution.⁵

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 a RAM program whose execution involves at most T CPU steps. We represent a RAM program P via a sequence of T small *CPU-Step Circuit* where each of them executes a single CPU step. In this work we will denote one CPU step by:

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

This circuit takes as input the current CPU state state and a block “data”. Looking ahead this block will be read from the memory location that was requested for a memory location requested for in the previous CPU step. The CPU step outputs an updated state state' , a read or write bit R/W , the next location to read/write $L \in [M]$, and a block z to write into the location ($z = \perp$ when reading). The sequence of locations and read/write values collectively form what is known as the *access pattern*, namely $\text{MemAccess} = \{(L^\tau, \text{R/W}^\tau, z^\tau, \text{data}^\tau) : \tau = 1, \dots, t\}$, and we can consider the weak access pattern $\text{MemAccess2} = \{L^\tau : \tau = 1, \dots, t\}$ of just the memory locations accessed.

Note that in the description above without loss of generality we have made some simplifying assumptions. First, we assume that the output z^{write} is written into the same location z^{read} was read from. Note that this is sufficient to both read from and write to arbitrary memory locations. Secondly we note that we assume that each CPU-step circuit always reads from and writes to some location in memory. This is easy to implement via a dummy read and write step. Finally, we assume that the instructions of the program itself are hardwired into the CPU-step circuits, and the program can first load itself into memory before execution. In cases where the size of the program vastly differs from its running time, one can suitably partition the program into two pieces.

Representing RAM computation by CPU-Step Circuits. The computation $P^D(x)$ starts with the initial state set as $\text{state}_0 = x$ and initial read location $L_0 = 0$ as a dummy read operation. In each step $\tau \in \{0, \dots, T-1\}$, the computation proceeds by first reading memory location L^τ , that is by setting $b^{\text{read}, \tau} := D[L^\tau]$ if $\tau \in \{1, \dots, T-1\}$ and as 0 if $\tau = 0$. Next it executes the CPU-Step Circuit $C_{\text{CPU}}^P(\text{state}^\tau, b^{\text{read}, \tau}) = (\text{state}^{\tau+1}, L^{\tau+1}, b^{\text{write}, \tau+1})$. Finally we write to the location L^τ by setting $D[L^\tau] := b^{\text{write}, \tau+1}$. If $\tau = T-1$ then we set state to be the output of the program P and ignore the value $L^{\tau+1}$. Note here that we have without loss of generality assumed that in one step the CPU-Step the same location in memory is read from and written to. This has been done for the sake of simplifying exposition.

⁵In 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 that data D is large while the size of the program $|P|$ and input length n is small.

3.2 Garbled Circuits

Garbled circuits was first constructed by Yao [Yao82] (see e.g. Lindell and Pinkas [LP09] and Bellare et al. [BHR12] for a detailed proof and further discussion). A circuit garbling scheme is a tuple of PPT algorithms (GCircuit, Eval). Very roughly GCircuit is the circuit garbling procedure and Eval the corresponding evaluation procedure. Looking ahead, each individual wire w of the circuit will be associated with two labels, namely $\text{lab}_0^w, \text{lab}_1^w$. Finally, since one can apply a generic transformation (see, e.g. [AIK10]) to blind the output, we allow output wires to also have arbitrary labels associated with them. Indeed, we can classify the output values into two categories — *plain outputs* and *labeled outputs*. The difference in the two categories stems from how they will be treated when garbled during garbling and evaluation. The plain output values do not require labels provided for them and evaluate to cleartext values. On the other hand labeled output values will require that additional output labels be provided to GCircuit at the time of garbling, and Eval will only return these output labels and not the underlying cleartext. We also define a well-formedness test for labels which we call Test.

- $(\tilde{C}) \leftarrow \text{GCircuit}(1^\kappa, C, \{(w, b, \text{lab}_b^w)\}_{w \in \text{inp}(C), b \in \{0,1\}})$: GCircuit takes as input a security parameter κ , a circuit C , and a set of labels lab_b^w for all the input wires $w \in \text{inp}(C)$ and $b \in \{0,1\}$. This procedure outputs a *garbled circuit* \tilde{C} .
- It can be efficiently tested if a set of labels is meant for a garbled circuit.
- $y = \text{Eval}(\tilde{C}, \{(w, \text{lab}_{x_w}^w)\}_{w \in \text{inp}(C)})$: Given a garbled circuit \tilde{C} and a garbled input represented as a sequence of input labels $\{(w, \text{lab}_{x_w}^w)\}_{w \in \text{inp}(C)}$, Eval outputs an output y in the clear.

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

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

where $(\tilde{C}) \leftarrow \text{GCircuit}(1^\kappa, C, \{(w, b, \text{lab}_b^w)\}_{w \in \text{inp}(C), b \in \{0,1\}})$.

Security. For security, we require that there is a PPT simulator CircSim such that for any C, x , and uniformly random labels $(\{(w, b, \text{lab}_b^w)\}_{w \in \text{inp}(C), b \in \{0,1\}})$, we have that:

$$\left(\tilde{C}, \{(w, \text{lab}_{x_w}^w)\}_{w \in \text{inp}(C)} \right)^{\text{comp}} \approx \text{CircSim}(1^\kappa, C, C(x))$$

where $(\tilde{C}) \leftarrow \text{GCircuit}(1^\kappa, C, \{(w, \text{lab}_b^w)\}_{w \in \text{out}(C), b \in \{0,1\}})$ and $y = C(x)$.

3.3 Garbled RAM

Next we consider an extension of garbled circuits to the setting of RAM programs. In this setting the memory data D is garbled once and then many different garbled programs can be executed sequentially with the memory changes persisting from one execution to the next. We define both full security and a weaker variant known as Unprotected Memory Access 2 (UMA2) (similar to UMA security that appeared in [GHL⁺14]), and we show how UMA2-secure Garbled RAM can be compiled with statistical Oblivious RAM to achieve full security.

Definition 3.1. A (UMA2) secure single-program garbled RAM scheme consists of four procedures (GData, GProg, GInput, GEval) with the following syntax:

- $(\tilde{D}, s) \leftarrow \text{GData}(1^\kappa, D)$: Given a security parameter 1^κ and memory $D \in \{0,1\}^M$ as input GData outputs the garbled memory \tilde{D} .

- $(\tilde{P}, s^{in}) \leftarrow \text{GProg}(1^\kappa, 1^{\log M}, 1^t, P, s, m)$: Takes the description of a RAM program P with memory-size M as input. It also requires a key s and current time m . It then outputs a garbled program \tilde{P} and an input-garbling-key s^{in} .
- $\tilde{x} \leftarrow \text{GInput}(1^\kappa, x, s^{in})$: Takes as input $x \in \{0, 1\}^n$ and an input-garbling-key s^{in} , a garbled “tree root” key s and outputs a garbled-input \tilde{x} .
- $y = \text{GEval}^{\tilde{D}}(\tilde{P}, \tilde{x})$: Takes a garbled program \tilde{P} , garbled input \tilde{x} and garbled memory data \tilde{D} and output a value y . We model GEval itself as a RAM program that can read and write to arbitrary locations of its memory initially containing \tilde{D} .

Efficiency. We require the run-time of GProg and GEval to be $t \cdot \text{poly}(\log M, \log T, \kappa)$, which also serves as the bound on the size of the garbled program \tilde{P} . Moreover, we require that the run-time of GData should be $M \cdot \text{poly}(\log M, \log T, \kappa)$, which also serves as an upper bound on the size of \tilde{D} . Finally the running time of GInput is required to be $n \cdot \text{poly}(\kappa)$.

Correctness. For correctness, we require that for any program P , initial memory data $D \in \{0, 1\}^M$ and input x we have that:

$$\Pr[\text{GEval}^{\tilde{D}}(\tilde{P}, \tilde{x}) = P^D(x)] = 1$$

where $(\tilde{D}, s) \leftarrow \text{GData}(1^\kappa, D)$, $(\tilde{P}, s^{in}) \leftarrow \text{GProg}(1^\kappa, 1^{\log M}, 1^t, P, s, m)$, $\tilde{x} \leftarrow \text{GInput}(1^\kappa, x, s^{in})$.

Security with Unprotected Memory Access (Full vs UMA2). For full or UMA2-security, we require that there exists a PPT simulator Sim such that for any program P , initial memory data $D \in \{0, 1\}^M$ and input x , which induces access pattern MemAccess we have that:

$$(\tilde{D}, \tilde{P}, \tilde{x}) \stackrel{\text{comp}}{\approx} \text{Sim}(1^\kappa, 1^M, 1^t, y, \text{MemAccess})$$

where $(\tilde{D}, s) \leftarrow \text{GData}(1^\kappa, D)$, $(\tilde{P}, s^{in}) \leftarrow \text{GProg}(1^\kappa, 1^{\log M}, 1^t, P, s, m)$ and $\tilde{x} \leftarrow \text{GInput}(1^\kappa, x, s^{in})$, and $y = P^D(x)$. Note that unlike UMA security, the simulator does not have access to D . For full security, the simulator Sim does not get MemAccess as input.

4 The Construction

In this section we describe our construction for garbled RAM formally, namely the procedures (GData , GProg , GInput , GEval). In order to make the exposition simpler, in this section we will describe our construction making four simplifying assumptions, which will be all removed in our final construction.

1. **UMA2-security:** Here we will restrict ourselves to achieving UMA2-security alone Definition 3.1 (UMA2). We note that this construction can then be amplified to get full security satisfying Definition 3.1 (full) using Lemma A.1 in Appendix A. This is essentially the transformation from previous works [LO13b, GHL⁺14, GLOS14] except that we need to restrict ourselves to using statistical ORAMs [DMN11, SCSL11, SvDS⁺13]. Note that this transformation is information theoretic and preserves the black-box nature of our construction.
2. **Uniform memory accesses:** We assume that the distribution of memory accesses of the programs being garbled are uniform. Then in Section 8 we will describe how this restriction can be removed and construction achieved even give any arbitrary probability distribution on memory reads. This will be essential for using Lemma A.1.
3. **First Step: Heuristic proof:** The construction described in this section is “heuristic” in the sense that we do not know how to prove its security. However we do not know of any concrete attacks against it. At a high level it suffers from a sort of a circular security problem. However this issue is rather easy to solve in our context. We describe the issue and the fix in Section 6 to obtain a full security proof.

4. **Bounded reads:** We will describe our construction assuming that the total number of memory accesses (both read and write) made to the garbled memory is bounded by M , the size of the memory. In Section 8, we explain how this restriction can be removed. In particular we will describe a memory replenishing mechanism for refilling the garbled memory as it used. This replenishing will involve some additional communication for each garbled program, while ensuring that the overhead of this replenishing information sent with each garbled program is small.

Notation. We use the notation $[n]$ to denote the set $\{0, \dots, n - 1\}$. For any string L , we use L_i to denote the i^{th} bit of L where $i \in [|x|]$ with the 0^{th} bit being the highest order bit. We let $L_{0\dots j-1}$ denote the j high order bits of L . We will be using multiple garbled circuits and will need notation to represent bundles of input labels for garbled circuits succinctly. In particular, if $\text{lab} = \{\text{lab}^{i,b}\}_{i \in |x|, b \in \{0,1\}}$ describes the labels for input wires of a garbled circuit, then we let lab_x denote the labels corresponding to setting the input to x , i.e. the subset of labels $\{\text{lab}^{i,x_i}\}_{i \in |x|}$. Similarly we will sometimes consider a collection of garbled circuits and denote the collection of labels for input wires of these garbled circuits with $\overline{\text{lab}}$. Let i be an index of a garbled circuit in this collection then we let $\text{lab}[i]_x$ denote the labels corresponding to setting the input to x of the i^{th} garbled circuit. Looking ahead, throughout our construction the inputs to the circuits we consider will be partitioned into two parts, the red and the blue. We will use the colors red and blue to stress whether an input label corresponds to a red input wire or a blue input wire. We extend this coloring to collections of labels of the same color. We believe that this makes it much easier to read our paper and recommend reading it on a coloured screen or a colored printout.

4.1 Data Garbling: $(\tilde{D}, s) \leftarrow \text{GData}(1^\kappa, D)$

We start by providing an informal description of the data garbling procedure, which turns out to be the most involved part of the construction. The formal description of GData is provided in Figure 4. Our garbled memory consists of two parts.

1. **Garbled Circuits:** Intuitively our garbled memory will be organized as a binary tree and each node of this tree will correspond to a sequence of garbled circuits. For any garbled circuit its *successor* is defined as the next garbled circuit in the sequence of garbled circuits corresponding to that node. Similarly we define *predecessor* as the previous garbled circuit in the sequence. For a garbled circuit all the garbled circuits in its parent node are referred to as its *parents*. Analogously we define *children*. These garbled circuits are obtained by fresh garblings of two separate circuits, one corresponding to the leaf nodes and the other corresponding to non-leaf nodes.

For each of these garbled circuits, we will divide its input wires (and corresponding keys/labels) into two categories, the *red* input wires and the *blue* input wires.

Each garbled circuit will contain *all* input keys for its successor. Specifically this includes both the red and the blue input keys of its successor. Additionally each garbled circuit will contain a *subset* of input wires for a *subset* of its left and right children. Specifically, it will contain the blue input keys for a consecutive κ garbled circuits among its left children and a consecutive κ garbled circuits among its right children.

2. **Tabled garbled Information:** Additionally for each node in the tree as described above, the garbled memory consists of a table of information $\text{Tab}(i, j)$, where (i, j) denotes a node in the tree.

Looking ahead, as the memory is read from or written to these garbled circuits that constitute the garbled memory will actually be used. Furthermore if a garbled circuit corresponding to a node is being consumed then its predecessor must have previously already been consumed. The tabulated information will be the red input keys for the first unused garbled circuit for each node.

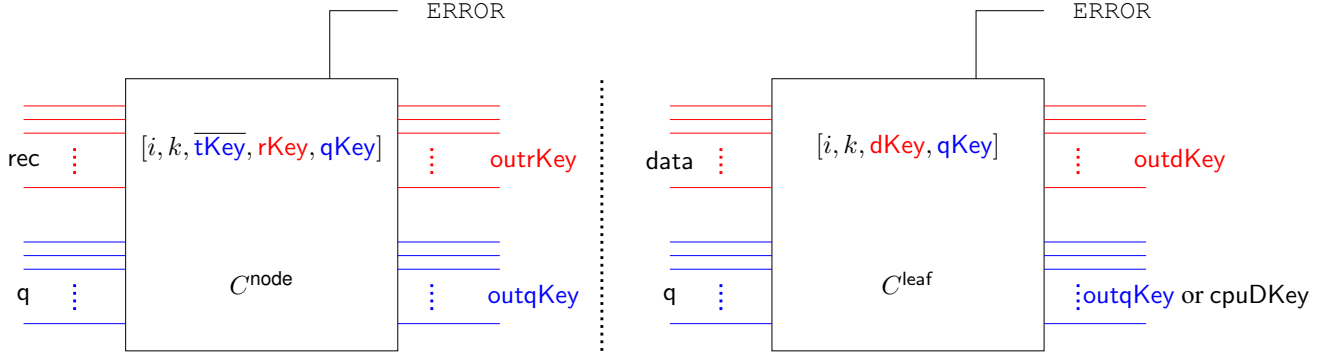


Figure 1: Memory circuits.

Circuits needed. Next we describe the two circuits, garblings of which will be used to generate the garbled memory. The circuits are described formally in Figures 2 and 3. The non-leaf node circuit takes as input some *recorded info* rec and a *query* q . Garbled labels for rec will be red and denoted as $rKey$ and garbled labels for q will be blue and denoted as $qKey$. Although every single circuit will have its own unique $rKey$ and $qKey$, when we refer to these in the context of some particular circuit, it will always refer to the keys of its successor and these values will be hard-coded in it. Additionally the circuit has hard-coded inside it its level i in the tree, its own position k within the sequence of garbled circuits at that node, garbled labels $rKey, qKey$ for its successor, and a collection of labels for the left and right child garbled circuits which we denote as $tKey$. Each $tKey$ is a vector of exactly 2κ $qKeys$, the first κ correspond to $qKeys$ for a contiguous block of κ of left child circuits (exactly which ones, we will describe later), and the last κ respectively correspond to labels for circuits in the right child.

The inputs are straightforward, rec contains indices to the first unused left and right child circuits as well as their $qKeys$. This allows us to either go left or right, although we will need to update the index and key as soon as we consume it, and we will replace it with something inside of $tKey$. The query q is simply a CPU query with one additional “goto” field $goto$ that informs where the first unused circuit in a node should be at to in order to fall inside the window of its parents $tKey$. If the current circuit $k < goto - 1$, we “burn” the circuit and pass our inputs onto our successor until it is precisely $goto - 1$, so that the first unused circuit is now indeed located at $goto$. In summary, we write $C^{node}[i, k, tKey, rKey, qKey](rec, q)$ for non-leaf circuits.

Similarly, leaf circuits takes as input some *memory data* $data$ and a *query* q . Here, the red key is $dKey$ which corresponds to the garbled labels of $data$. A leaf circuit will have hard-coded inside it the current level $i = d$ in the tree, its own position k within the sequence of garbled circuits at that leaf, garbled labels $dKey, qKey$ for its successor. Since a leaf node has no further children, there is no need for $tKey$. We write $C^{leaf}[i, k, dKey, qKey](data, q)$ for these circuits.

$C^{\text{node}}[i, k, \overline{\text{tKey}}, \overline{\text{rKey}}, \overline{\text{qKey}}]$

System parameters: ϵ (Will be set to $\frac{1}{\log M}$ as we will see later.)

Hardcoded parameters: $[i, k, \overline{\text{tKey}}, \overline{\text{rKey}}, \overline{\text{qKey}}]$

Input: $(\text{rec} = (\text{lidx}, \text{ridx}, \text{oldLKey}, \text{oldRKey}), \text{q} = (\text{goto}, \text{R/W}, L, z, \text{cpuDKey}))$.

Set $p := \text{goto}$ and $p' := \lfloor (\frac{1}{2} + \epsilon) k \rfloor$. Set $\text{lidx}' := \text{lidx}$ and $\text{ridx}' := \text{ridx}$. Set $\text{oldLKey}' := \text{oldLKey}$ and $\text{oldRKey}' := \text{oldRKey}$. We now have three cases:

1. If $k < p - 1$ then we output $(\overline{\text{outrKey}}, \overline{\text{outqKey}}) := (\overline{\text{rKey}}_{\text{rec}'}, \overline{\text{qKey}}_{\text{q}})$, where $\text{rec}' := (\text{lidx}', \text{ridx}', \text{oldLKey}', \text{oldRKey}')$.
2. If $k \geq p + \kappa$ then abort with output OVERCONSUMPTION-ERROR-I.
3. If $p - 1 \leq k < p + \kappa$ then:
 - (a) If $L_i = 0$ then,
 - i. If $\text{lidx} < p'$ then set $\text{lidx}' := p', \text{goto}' := p'$ and $\text{oldLKey}' := \overline{\text{tKey}}[0]$. Else set $\text{lidx}' := \text{lidx} + 1, \text{goto}' := \text{lidx}'$ and if $\text{lidx}' < p' + \kappa$ then set $\text{oldLKey}' := \overline{\text{tKey}}[\text{lidx}' - p']$ else abort with OVERCONSUMPTION-ERROR-II.
 - ii. Set $\overline{\text{outqKey}} := \text{oldLKey}_{\text{q}'}$, where $\text{q}' := \text{q}$ but with goto' replacing goto .
 - else
 - i. If $\text{ridx} < p'$ then set $\text{ridx}' := p', \text{goto}' := p'$ and $\text{oldRKey}' := \overline{\text{tKey}}[\kappa]$. Else set $\text{ridx}' := \text{ridx} + 1, \text{goto}' := \text{ridx}'$ and if $\text{ridx}' < p' + \kappa$ then set $\text{oldRKey}' := \overline{\text{tKey}}[\kappa + \text{ridx}' - p']$ else abort with OVERCONSUMPTION-ERROR-II.
 - ii. Set $\overline{\text{outqKey}} := \text{oldRKey}_{\text{q}'}$, where $\text{q}' := \text{q}$ but with goto' replacing goto .
 - (b) Set $\overline{\text{outrKey}} := \overline{\text{rKey}}_{\text{rec}'}$ where $\text{rec}' := (\text{lidx}', \text{ridx}', \text{oldLKey}', \text{oldRKey}')$ and output $(\overline{\text{outrKey}}, \overline{\text{outqKey}})$.

Figure 2: Formal description of the nonleaf Memory Circuit.

$C^{\text{leaf}}[i, k, \overline{\text{dKey}}, \overline{\text{qKey}}]$

System parameters: ϵ (Will be set to $\frac{1}{\log M}$ as we will see later.)

Hardcoded parameters: $[i, k, \overline{\text{dKey}}, \overline{\text{qKey}}]$

Input: $(\text{data}, \text{q} = (\text{goto}, \text{R/W}, L, z, \text{cpuDKey}))$.

Set $p := \text{goto}$ and $p' := \lfloor (\frac{1}{2} + \epsilon) k \rfloor$. We now have three cases:

1. If $k < p - 1$ then we output $(\overline{\text{outdKey}}, \overline{\text{outqKey}}) := (\overline{\text{dKey}}_{\text{data}}, \overline{\text{qKey}}_{\text{q}})$.
2. If $k \geq p + \kappa$ then abort with output OVERCONSUMPTION-ERROR-I.
3. If $p - 1 \leq k < p + \kappa$ then:
 - (a) If $\text{R/W} = \text{read}$ then output $(\overline{\text{dKey}}_{\text{data}}, \text{cpuDKey}_{\text{data}})$, else if $\text{R/W} = \text{write}$ then output $(\overline{\text{dKey}}_z, \text{cpuDKey}_z)$.

Figure 3: Formal description of the leaf Memory Circuit.

The algorithm $\text{GData}(1^\kappa, D)$ proceeds as follows. Without loss of generality we assume that $M = 2^d$ (where $M = |D|$) where d is a positive integer. We calculate $\epsilon = \frac{1}{\log M}$. We set $K_0 = M$, and for each $i \in [d+1]$ and set $K_i = \lfloor (\frac{1}{2} + \epsilon) K_{i-1} \rfloor + \kappa$.

1. Let $s \leftarrow \{0, 1\}^\kappa$.
2. Any $\mathbf{dKey}^{d,j,k}$ needed in the computation below is obtained as $F_s(\text{data}||d||j||k)$. Similarly for any i, j, k , $\mathbf{rKey}^{i,j,k} := F_s(\text{rec}||i||j||k)$ and $\mathbf{qKey}^{i,j,k} := F_s(\text{query}||i||j||k)$. Finally,

$$\overline{\mathbf{tKey}}^{i,j,k} := \left\{ \underbrace{\left\{ \mathbf{qKey}^{i+1,2j, \lfloor (\frac{1}{2} + \epsilon)k \rfloor + l} \right\}_{l \in [\kappa]}}_{\text{left}}, \underbrace{\left\{ \mathbf{qKey}^{i+1,2j+1, \lfloor (\frac{1}{2} + \epsilon)k \rfloor + l} \right\}_{l \in [\kappa]}}_{\text{right}} \right\}.$$

3. For all $j \in [2^d], k \in [K_d]$,
 $\tilde{C}^{d,j,k} \leftarrow \text{GCircuit}(1^\kappa, C^{\text{leaf}}[d, k, \mathbf{dKey}^{d,j,k+1}, \mathbf{qKey}^{d,j,k+1}], \mathbf{dKey}^{d,j,k}, \mathbf{qKey}^{d,j,k})$.
4. For all $i \in [d], j \in [2^i], k \in [K_i]$,
 $\tilde{C}^{i,j,k} \leftarrow \text{GCircuit}(1^\kappa, C^{\text{node}}[i, k, \overline{\mathbf{tKey}}^{i,j,k}, \mathbf{rKey}^{i,j,k+1}, \mathbf{qKey}^{i,j,k+1}], \mathbf{rKey}^{i,j,k}, \mathbf{qKey}^{i,j,k})$.
5. For all $j \in [2^d]$, set $\mathbf{Tab}(d, j) = \mathbf{dKey}_{D[j]}^{d,j,0}$.
6. For all $i \in [d], j \in [2^i]$, set $\mathbf{Tab}(i, j) := \mathbf{rKey}_{\text{rec}^{i,j,0}}^{i,j,0}$, where $\text{rec}^{i,j,0} := (0, 0, \mathbf{qKey}^{i+1,2j,0}, \mathbf{qKey}^{i+1,2j+1,0})$.
7. Output $\tilde{D} := \left(\left\{ \tilde{C}^{i,j,k} \right\}_{i \in [d+1], j \in [2^i], k \in [K_i]}, \left\{ \mathbf{Tab}(i, j) \right\}_{i \in [d+1], j \in [2^i]} \right)$ and s .

Figure 4: Formal description of GData.

Actual data garbling. At a high level, we generate our garbled memory by garbling multiple instances of circuits described in Figures 2 and 3. The formal construction is provided in Figure 4. As mentioned earlier, these garbled circuits actually correspond to the nodes of a tree. Specifically, if the size of the database is $M = 2^d$, then the root node will contain roughly $M\kappa$ circuits, each nodes in subsequent level will contain roughly half that amount. More specifically, any node at level i contains at most $K_i = \lfloor (\frac{1}{2} + \epsilon)^i (M + i\kappa) \rfloor + \kappa$ garbled circuits. In total, the garbled memory will consist of $\sum_{i=0}^d (1 + 2\epsilon)^i (M + i\kappa) + \kappa$ garbled circuits. Looking ahead ϵ will be set to $\frac{1}{d}$ so that the this number is linear in $M + d\kappa$. In order to simplify generation of garbled circuits, we generate all the labels needed for generation of these garbled circuits as the outputs of a PRF on appropriate input values under a fixed seed s . Looking ahead, this will be crucial in extending our construction to allow for generating memory replenishing information. This is elaborated upon in Section 8.

4.2 Program Garbling: $(\tilde{P}, s^{in}) \leftarrow \text{GProg}(1^\kappa, 1^{\log M}, 1^t, P, s, m)$

We start by defining a sub-circuit that will be needed in describing the program garbling in Figure 5. This circuit basically performs one step of the CPU and provides input labels for a root garbled circuit in the garbled memory. The formal description of program garbling itself is provided in Figure 6. The garbled program is obtained by garbling multiple CPU step circuits where very rough each circuit provides the input labels for the next CPU step and for the root circuit of the garbled memory, which then enables reading data

from memory.

$C^{\text{step}}[t, \text{rootqKey}, \text{cpuSKey}, \text{cpuDKey}]$

Hardcoded parameters: $[t, \text{rootqKey}, \text{cpuSKey}, \text{cpuDKey}]$

Input: (state, data).

Compute $(\text{state}', R/W, L, z) := C_{\text{CPU}}^P(\text{state}, \text{data})$. Set $q := (\text{goto} = t + 1, R/W, L, z, \text{cpuDKey})$ and output rootqKey_q and $\text{cpuSKey}_{\text{state}'}$, unless we are halting in which case only output state' in the clear.

Figure 5: Formal description of the step circuit.

The $\text{GProg}(1^\kappa, 1^{\log M}, 1^t, P, s, m)$ procedure proceeds as follows.

1. Any cpuSKey^τ needed in the computation below is obtained as $F_s(\text{CPUstate} || \tau)$, and any cpuDKey^τ is obtained as $F_s(\text{CPUdata} || \tau)$.
2. For $\tau = m, \dots, m + t - 1$ do:
 - (a) Set $\text{qKey}^{0,0,\tau} := F_s(\text{query} || 0 || 0 || \tau)$.
 - (b) $\tilde{C}^\tau \leftarrow \text{GCircuit}(1^\kappa, C^{\text{step}}[\tau, \text{qKey}^{0,0,\tau}, \text{cpuSKey}^{\tau+1}, \text{cpuDKey}^{\tau+1}], \text{cpuSKey}^\tau, \text{cpuDKey}^\tau)$
3. Output $\tilde{P} := (m, \{\tilde{C}^\tau\}_{\tau \in \{m, \dots, m+t-1\}}, \text{cpuDKey}_0^m)$, $s^{in} = \text{cpuSKey}^m$

Figure 6: Formal description of GProg.

4.3 Input Garbling: $\tilde{x} \leftarrow \text{GInput}(1^\kappa, x, s^{in})$

Informally, the GInput algorithm uses x as selection bits for the labels provided by s^{in} and outputs \tilde{x} , which is just the selected labels. A formal description of GProg is provided in Figure 7.

The algorithm $\text{GInput}(1^\kappa, x, s^{in})$ proceeds as follows.

1. Parse s^{in} as cpuSKey and output $\tilde{x} := \text{cpuSKey}_x$.

Figure 7: Formal description of GInput.

4.4 Garbled Evaluation: $y \leftarrow \text{GEval}^{\tilde{D}}(\tilde{P}, \tilde{x})$

The GEval procedure gets as input the garbled program $\tilde{P} = (m, \{\tilde{C}^\tau\}_{\tau \in \{m, \dots, m+t-1\}}, \text{cpuDKey})$, the garbled input $\tilde{x} = \text{cpuSKey}$ and random access into the garbled database $\tilde{D} = (\{\tilde{C}^{i,j,k}\}_{i \in [d+1], j \in [2^i], k \in [K_i]}, \{\text{Tab}(i, j)\}_{i \in [d+1], j \in [2^i]})$. Intuitively the GEval is very simple. It proceeds by executing a subset of the garbled circuits from the garbled program and the garbled memory in a specific order which is decided dynamically based on the computation. The labels needed to evaluate the first garbled circuit are provided as part of the garbled input and each evaluation of a garbled circuit reveals the labels for at most two distinct circuits. Among these two circuits, only one is such that all its input labels have been provided, and this circuit is executed next. The unused input labels are stored in memory table **Tab** to be used at a later point. Next we provide the formal description of GEval in Figure 8.

The algorithm $\text{GEval}^{\tilde{D}}(\tilde{P}, \tilde{x})$ proceeds as follows.

1. Parse \tilde{P} as $(m, \{\tilde{C}^\tau\}_{\tau \in \{m, \dots, m+t-1\}}, \text{cpuDKey})$, \tilde{x} as cpuSKey and \tilde{D} as $\left(\left\{ \tilde{C}^{i,j,k} \right\}_{i \in [d+1], j \in [2^i], k \in [K_i]}, \{\text{Tab}(i, j)\}_{i \in [d+1], j \in [2^i]} \right)$.
2. For $\tau \in \{m, \dots, m+t-1\}$ do:
 - (a) Evaluate $(\text{cpuSKey}, \text{qKey}) := \text{Eval}(\tilde{C}^\tau, (\text{cpuSKey}, \text{cpuDKey}))$. If an output y is produced by Eval instead, then output y and halt.
 - (b) Set $i = 0, j = 0, k = \tau$.
 - (c) Evaluate $\text{outputKey} := \text{Eval}(\tilde{C}^{i,j,k}, (\text{Tab}(i, j), \text{qKey}))$.
 - i. If outputKey is parsed as $(\text{rKey}, \text{qKey}^{i',j',k'})$ for some i', j', k' , then set $\text{Tab}(i, j) := \text{rKey}, \text{qKey} := \text{qKey}^{i',j',k'}, (i, j, k) = (i', j', k')$ and go to Step 2c.
 - ii. Otherwise, set $(\text{dKey}, \text{cpuDKey}) := \text{outputKey}$, and $\text{Tab}(i, j) := \text{dKey}$ and $\tau := \tau + 1$.

Figure 8: Formal description of GEval.

5 Cost and Correctness Analysis

5.1 Overall Cost

Before we analyze the cost of the main algorithms, we first calculate the sizes of all the constituent variables and circuits. The database D has size $|D| = M$ elements, and each data element is of $|\text{data}| = B$ bits. Garbled labels for each bit of an input wire are λ bits long. The complete garbled labels for n input bits takes up $2\lambda n$ bits. Furthermore, the current time step m or τ we upper bound by the total combined running time T . Of course, B, λ, T are all $\text{poly}(\kappa)$, and for the two former values we simply absorb them into the $\text{poly}(\kappa)$ term, whereas we keep T as a separate parameter for later use.

From this, we can compute $|\text{cpuSKey}| = 2\lambda|\text{state}|$ and $|\text{cpuDKey}| = 2\lambda B$. A query q has size $|\text{goto}| + |\text{R/W}| + |L| + |z| + |\text{cpuDKey}| \leq \log T + 1 + \log M + B + 2\lambda B$, and $|\text{qKey}| = 2\lambda|q|$. Since dKey are just the labels for memory data, $|\text{dKey}| = 2\lambda B$. Next, we compute the size of rec . Observe that oldLKey and oldRKey are simply qKeys , so we have $|\text{rec}| = |\text{lidx}| + |\text{ridx}| + |\text{oldLKey}| + |\text{oldRKey}| \leq 2(\log T + |\text{qKey}|)$. Finally, tKey consist of 2κ qKeys and therefore have size $|\text{tKey}| = 2\kappa|\text{qKey}|$.

Now we calculate $|C^{\text{node}}|$. Observe that the calculations within the circuit are primarily comparisons, and overall is at most polynomial in the size of the input and hardwired values. Thus $|C^{\text{node}}| = \text{poly}(|i| + |k| + |\text{tKey}| + |\text{rKey}| + |\text{qKey}| + |\text{rec}| + |q|) = \text{poly}(\log M, \log T, \kappa)$ and so is its garbled version.

Next, we calculate $|C^{\text{leaf}}|$. We have $|C^{\text{leaf}}[i, k, \text{dKey}, \text{qKey}](\text{data}, q)| = \text{poly}(|i| + |k| + |\text{dKey}| + |\text{qKey}| + |\text{data}| + |q|) = \text{poly}(\log M, \log T, \kappa)$ and so is its garbled version.

Finally, we calculate $|C^{\text{step}}|$. We assume that the plain CPU has size $\text{poly}(\log M, \log T, \kappa)$. Since the step circuit simply computes the CPU circuit and does a few selections, we have: $|C^{\text{step}}| = \text{poly}(\log T + |\text{qKey}| + |\text{cpuSKey}| + |\text{cpuDKey}| + |\text{state}| + |\text{data}| + |\text{CPU}|) = \text{poly}(\log M, \log T, \kappa)$.

We can now calculate the cost of the individual algorithms.

5.1.1 Cost of GData

The algorithm $\text{GData}(1^\kappa, D)$ first computes $O(M)$ $\text{dKey}, \text{rKey}, \text{qKey}$ values, taking $M \cdot \text{poly}(\log M, \log T, \kappa)$ steps. For each node at level $i < d$, it computes K_i garbled C^{node} circuits and tabulates an rKey for each of the 2^i nodes. At level $i = d$, it computes K_d garbled C^{leaf} circuits and tabulates $M = 2^d$ dKeys . The output is of size equal to all the garbled circuits plus the size of the tabulated values plus one PRF key. Let

$c = e^2$ and let $\epsilon = \frac{1}{\log M}$.

First, we show how to bound $K_i \leq \left(\frac{1}{2} + \epsilon\right)^i M + \sum_{j=0}^{i-1} \left(\frac{1}{2} + \epsilon\right)^j \kappa$. This can be shown by induction: $K_0 = M$, and by induction, $K_{i+1} = \left\lfloor \left(\frac{1}{2} + \epsilon\right) K_i \right\rfloor + \kappa \leq \left(\frac{1}{2} + \epsilon\right) K_i + \kappa \leq \left(\frac{1}{2} + \epsilon\right) \cdot \left[\left(\frac{1}{2} + \epsilon\right)^i M + \sum_{j=0}^{i-1} \left(\frac{1}{2} + \epsilon\right)^j \kappa \right] + \kappa \leq \left(\frac{1}{2} + \epsilon\right)^{i+1} M + \sum_{j=0}^{i+1-1} \left(\frac{1}{2} + \epsilon\right)^j \kappa$.

This bound can then be simplified to $K_i \leq \left(\frac{1}{2} + \epsilon\right)^i (M + i\kappa)$.

Thus, overall, we calculate the number of garblings of C^{node} as

$$\begin{aligned} \sum_{i=0}^{d-1} 2^i \cdot K_i &\leq \sum_{i=0}^{d-1} 2^i \cdot \left(\frac{1}{2} + \epsilon\right)^i (M + i\kappa) \\ &\leq \sum_{i=0}^{d-1} (1 + 2\epsilon)^i (M + d\kappa) \\ &\leq \frac{(1 + 2\epsilon)^d - 1}{(1 + 2\epsilon) - 1} (M + d\kappa) \\ &\leq \frac{e^{2\epsilon d} - 1}{2\epsilon} (M + d\kappa) \end{aligned}$$

Since $\epsilon = 1/d$, and garbling such circuits takes $\text{poly}(\log M, \log T, \kappa)$ time, this overall takes $M \cdot \text{poly}(\log M, \log T, \kappa)$ time and space. At the leaf level, it performs at most $2^d \cdot \left(\frac{1}{2} + \epsilon\right)^d (M + d\kappa)$ garblings of C^{leaf} . Again, this takes $\text{poly}(\log M, \log T, \kappa) \cdot M$ time and space. Finally, there are $O(M)$ of **rKey** and **dKey** values stored in **Tab**(i, j), which is again $\text{poly}(\log M, \log T, \kappa) \cdot M$.

5.1.2 Cost of GProg

The algorithm $\text{GProg}(1^\kappa, 1^{\log M}, 1^t, P, s, m)$ computes t **cpuSKeys**, **cpuDKeys**, and **qKeys**. It also garbles t C^{step} circuits and outputs them, along with a single **cpuSKey**. Since each individual operation is $\text{poly}(\log M, \log T, \kappa)$, the overall time and space cost is $\text{poly}(\log M, \log T, \kappa) \cdot t$.

5.1.3 Cost of GInput

The algorithm $\text{GInput}(1^\kappa, x, s^{\text{in}})$ selects labels of the state key based on the state as input. As such, the time and space cost is $|\text{cpuSKey}|$.

5.1.4 Cost of GEval

We first assume that an error does not occur in **GEval**. As we shall see in Section 5.2 that this occurs with all but negligible probability. We analyze how many circuits were consumed after T steps in order to obtain the amortized cost of **GEval**. We let k_i denote the maximum number of circuits consumed in some node at level i . At the root, exactly T circuits were consumed so $k_0 = T$, and in order for level i to not overflow, it must not have consumed more than $\left\lfloor \left(\frac{1}{2} + \epsilon\right) k_{i-1} \right\rfloor + \kappa$ circuits. By the same analysis of the bound of K_i , it must be the case that $k_i \leq \left(\frac{1}{2} + \epsilon\right)^i (T + i\kappa)$. Then no more than $\sum_{i=0}^d 2^i k_i$ circuits could have been consumed, each of which has evaluation cost at most $\text{poly}(\log M, \log T, \kappa)$. It turns out this bound is slightly insufficient, and this is due to the case when $T < M$, the 2^i term is an overestimate. Indeed, if there are only T accesses, then there can be at most T nodes that were ever touched at a level. Using $\min(2^i, T)$

as the bound on the number of nodes ever touched per level suffices:

$$\begin{aligned}
\sum_{i=0}^d \min(2^i, T) k_i &\leq \sum_{i=0}^d \min(2^i, T) \left(\frac{1}{2} + \epsilon\right)^i (T + i\kappa) \\
&\leq T \left(\sum_{i=0}^d \min(2^i, T) \left(\frac{1}{2} + \epsilon\right)^i + \frac{\min(2^i, T) i \kappa}{T} \right) \\
&\leq T \left(\sum_{i=0}^d (2^i) \left(\frac{1}{2} + \epsilon\right)^i + \frac{(T) i \kappa}{T} \right) \\
&\leq T \left(d((1 + 2\epsilon)^d + d\kappa) \right) \\
&\leq T (d(e^2 + d\kappa)).
\end{aligned}$$

When accounting for the cost of each of these circuits being evaluated, this means that the amortized cost is $T \cdot \text{poly}(\log M, \log T, \kappa)$ overall.

5.2 Correctness

Observe that as long as the memory data is correctly stored and passed on to the CPU step circuits, the scheme is correct. The only way this can fail to happen is if a query q fails to make it from the root to the leaf. In order to demonstrate this, we need to analyze two things. We must show that a parent circuit will always output the proper **qKey** for the first unused child circuit, and we also must show that the errors `OVERCONSUMPTION-ERROR-I` and `OVERCONSUMPTION-ERROR-II` do not occur except with a negligible probability.

Lemma 5.1. *Within C^{node} , $\text{lid}x$ always points to the first unused left child circuit which has **qKey** equal to oldLKey , and $\text{rid}x$ always opint to the first unused right child circuit which has **qKey** equal to oldRKey .*

Proof. WLOG we show this for the left child. We prove this by induction on the current CPU step. In the base case, this is true due to the way `GData` set up the keys. Now suppose we are consuming some parent circuit and it was true for the previous circuit, i.e. $\text{lid}x$ and oldLKey correctly point to the first unused left child circuit. Then it remains to show that $\text{lid}x'$ points to what will be the first unused left child circuit during the next CPU step, and that the updated old key $\text{oldLKey}'$ points to it. Recall $p' = \lfloor (\frac{1}{2} + \epsilon)k \rfloor$, and by definition of `GData`, this is precisely the child circuit that $\text{tKey}[0]$ is the **qKey** of. If $\text{lid}x < p'$ then the child circuit will burn until the goto' circuit, which is exactly what $\text{lid}x'$ is set to be, and $\text{oldLKey}'$ is set to $\text{tKey}[0]$ which is precisely what $\text{lid}x' = \text{goto}'$ is set to. On the other hand, if $\text{lid}x > \text{goto}'$ then by definition, $\text{goto}' = \text{lid}x' = \text{lid}x + 1$ and $\text{oldLKey}'$ holds the key for precisely the next circuit past $\text{lid}x$. But we know that the child node will consume exactly one circuit since goto' is precisely one past $\text{lid}x$ which by induction is the current child index, so $\text{lid}x'$ will point to the first unused child circuit and $\text{oldLKey}'$ is its key. \square

Lemma 5.2. *The errors `OVERCONSUMPTION-ERROR-I` and `OVERCONSUMPTION-ERROR-II` do not occur except with a negligible probability.*

Proof. Again, WLOG we show this for the left child. Note that an error can never occur at the root, and `OVERCONSUMPTION-ERROR-I` would occur if and only if an `OVERCONSUMPTION-ERROR-II` would have occurred just before it. Thus, we bound the probability that an `OVERCONSUMPTION-ERROR-I` could occur. Suppose an error first occurs at some node (i, j) at the m_1 -th circuit in this node. Then this means that the child $\text{lid}x'$ has become greater than $p' + \kappa = \lfloor (\frac{1}{2} + \epsilon)m_1 \rfloor + \kappa$. Since each time the left child is visited, many child circuits may be consumed due to burning, it might be difficult to figure out exactly how many child circuits were consumed if m_1 parent circuits were consumed. However, we can define a

synchronize event, which is namely that the parent is on circuit k and the child is on circuit $\lfloor (\frac{1}{2} + \epsilon)k \rfloor$, or more precisely, when $goto' = \lfloor (\frac{1}{2} + \epsilon)k \rfloor$. We let $m_0 < m_1$ be the latest point for the parent for which this synchronize occurred. We know that such an m_0 exists, since time $m_0 = 0$ is a valid solution.

Because there have been no more burns since that time, each time the left child was visited, exactly one circuit was also consumed. At m_0 , exactly $\lfloor (\frac{1}{2} + \epsilon)m_0 \rfloor$ child circuits were consumed, and at m_1 , more than $\lfloor (\frac{1}{2} + \epsilon)m_1 \rfloor + \kappa$ child circuits would have been consumed (if we did not break on error). During this time, $m_1 - m_0$ parent circuits were consumed, so the parent node was visited at most $m_1 - m_0$ times (it could be less due to burning), and we expect the child node to be visited $\mu = \frac{m_1 - m_0}{2}$ times. For $t = 0, \dots, m_1 - m_0$, let X_t denote the 0/1 random variable indicating that on time step $m_0 + t$ the left node was visited, and let $X = \sum_{t=0}^{m_1 - m_0} X_t$. We calculate the probability that $Pr[X > \lfloor (\frac{1}{2} + \epsilon)m_1 \rfloor + \kappa - \lfloor (\frac{1}{2} + \epsilon)m_0 \rfloor]$ which is the probability that this error would have occurred. This becomes $Pr[X > ((\frac{1}{2} + \epsilon)(m_1 - m_0)) + \kappa - 1]$. Note that we can trivially condition on the case where $m_1 - m_0 > \kappa$, because otherwise $X < \kappa$ with probability 1, so we can conclude $\mu > \kappa/2$.

Substituting $\delta = 2\epsilon + \frac{\kappa - 1}{\mu}$, this becomes $Pr[X > (1 + \delta)\mu]$. Then by the Chernoff bound for $\delta > 0$, $Pr[X > (1 + \delta)\mu] \leq \exp((\delta - (1 + \delta))\mu)$.

Then reorganizing terms and using a second-order log approximation:

$$\begin{aligned}
Pr[X > (1 + \delta)\mu] &\leq \exp \left[\left(\delta + (1 + \delta) \log \left(1 - \frac{\delta}{1 + \delta} \right) \right) \mu \right] \\
&\leq \exp \left[\left(\delta + (1 + \delta) \left(-\frac{\delta}{1 + \delta} - \frac{\delta^2}{(1 + \delta)^2} \right) \right) \mu \right] \\
&\leq \exp \left[-\frac{\delta^2 \mu}{1 + \delta} \right] \\
&\leq \exp \left[-\delta \mu \left(\frac{\delta}{1 + \delta} \right) \right] \\
&\leq \exp \left[-(2\epsilon\mu + \kappa - 1) \left(\frac{2\epsilon + (\kappa - 1)/\mu}{1 + 2\epsilon + (\kappa - 1)/\mu} \right) \right] \\
&\leq \exp \left[-(2\epsilon\mu + \kappa - 1) \left(\frac{2\epsilon}{1 + 2\epsilon} \right) \right] \\
&\leq \exp \left[-(2\epsilon\mu + \kappa - 1) \left(\frac{2\epsilon}{1 + 1} \right) \right] \\
&\leq \exp \left[-(2\epsilon^2\mu + \epsilon(\kappa - 1)) \right]
\end{aligned}$$

Since $\epsilon = \frac{1}{\log M}$, this is negligible. □

6 Secure Main Construction

We state the theorem that we attempt to prove and show where there is a barrier and how we work around it by modifying our solution.

We first provide the intuition of how we would like the proof would go through. Our goal is to construct a simulator Sim such that only given the access pattern and output (and not the database contents) it can simulate the view of the evaluator. The first observation is that the only point of the PRF F was to allow GProg to efficiently be able to compute the root keys and to replenish new circuits in nodes without having to remember all the existing labels. Since Sim can run in time propotional to the size of the database, it can

simply remember all these values internally, and therefore the first step is to replace F with a truly random table that Sim keeps track of.

Next, we must simulate the garbled circuits one at a time. In order to do so, we order the circuits by the order in which they were evaluated, and then define a series of hybrids where hybrid i has the first i garbled circuits simulated. The hybrids are constructed so that the circuits are simulated in reverse order, such that the i -th circuit is simulated first, and so on, until the first circuit \tilde{C}^0 is simulated. At first glance, this appears to work, but there is actually a subtle issue. Each \mathbf{qKey} of a circuit in a node circuit resides in several locations: the predecessor circuit in the same node, and inside several \mathbf{tKey} , $\mathbf{oldLKey}$, $\mathbf{oldRKey}$ of circuits in its parent node. Indeed, if a full \mathbf{qKey} appears anywhere (whether as a passed or hardwired value) in the current hybrid, then this causes a ‘‘circularity’’ issue. In order to overcome this barrier, we *pass* \mathbf{tKey} from circuit to circuit instead of hardwiring them as seen in Figures 9 and 10.

$C^{\text{node}}[i, k, \mathbf{newLtKey}, \mathbf{newRtKey}, \mathbf{rKey}, \mathbf{qKey}]$
System parameters: ϵ (Will be set to $\frac{1}{\log M}$ as we will see later.)
Hardcoded parameters: $[i, k, \mathbf{newLtKey}, \mathbf{newRtKey}, \mathbf{rKey}, \mathbf{qKey}]$
Input: $(\text{rec} = (\text{lidx}, \text{ridx}, \mathbf{oldLKey}, \mathbf{oldRKey}, \mathbf{tKey}), \mathbf{q} = (\text{goto}, \text{R/W}, L, z, \text{cpuDKey}))$.

Set $\mathbf{p} := \text{goto}$ and $\mathbf{p}' := \lfloor (\frac{1}{2} + \epsilon) k \rfloor$.
Set $\text{lidx}' := \text{lidx}$ and $\text{ridx}' := \text{ridx}$. Set $\mathbf{oldLKey}' := \mathbf{oldLKey}$ and $\mathbf{oldRKey}' := \mathbf{oldRKey}$.
Define $\text{ins}(\mathbf{tKey}, \mathbf{newLtKey}, \mathbf{newRtKey})$ to be the function that outputs \mathbf{tKey} with a possible shift: if $\lfloor (\frac{1}{2} + \epsilon) (k + 1) \rfloor > \lfloor (\frac{1}{2} + \epsilon) k \rfloor$, shift \mathbf{tKey} to the left by 1 and set $\mathbf{tKey}[\kappa - 1] = \mathbf{newLtKey}$, $\mathbf{tKey}[2\kappa - 1] = \mathbf{newRtKey}$.
We now have three cases:

1. If $k < \mathbf{p} - 1$ then we output $(\mathbf{outrKey}, \mathbf{outqKey}) := (\mathbf{rKey}_{\text{rec}'}, \mathbf{qKey}_{\mathbf{q}})$, where $\text{rec}' := (\text{lidx}', \text{ridx}', \mathbf{oldLKey}', \mathbf{oldRKey}', \mathbf{tKey}')$ where $\mathbf{tKey}' = \text{ins}(\mathbf{tKey}, \mathbf{newLtKey}, \mathbf{newRtKey})$.
2. If $k \geq \mathbf{p} + \kappa$ then abort with output OVERCONSUMPTION-ERROR-I.
3. If $\mathbf{p} - 1 \leq k < \mathbf{p} + \kappa$ then:
 - (a) If $L_i = 0$ then,
 - i. If $\text{lidx} < \mathbf{p}'$ then set $\text{lidx}' := \mathbf{p}'$, $\text{goto}' := \mathbf{p}'$ and $\mathbf{oldLKey}' := \mathbf{tKey}[0]$. Else set $\text{lidx}' := \text{lidx} + 1$, $\text{goto}' := \text{lidx}'$ and if $\text{lidx}' < \mathbf{p}' + \kappa$ then set $\mathbf{oldLKey}' := \mathbf{tKey}[\text{lidx}' - \mathbf{p}']$ else abort with OVERCONSUMPTION-ERROR-II.
 - ii. Set $\mathbf{tKey}[v] := \perp$ for all $v < \text{lidx}' - \mathbf{p}'$. Set $\overline{\mathbf{tKey}}' = \text{ins}(\overline{\mathbf{tKey}}, \mathbf{newLtKey}, \mathbf{newRtKey})$.
 - iii. Set $\mathbf{outqKey} := \mathbf{oldLKey}_{\mathbf{q}'}$, where $\mathbf{q}' := \mathbf{q}$ but with goto' replacing goto .
 - else
 - i. If $\text{ridx} < \mathbf{p}'$ then set $\text{ridx}' := \mathbf{p}'$, $\text{goto}' := \mathbf{p}'$ and $\mathbf{oldRKey}' := \mathbf{tKey}[\kappa]$. Else set $\text{ridx}' := \text{ridx} + 1$, $\text{goto}' := \text{ridx}'$ and if $\text{ridx}' < \mathbf{p}' + \kappa$ then set $\mathbf{oldRKey}' := \mathbf{tKey}[\kappa + \text{ridx}' - \mathbf{p}']$ else abort with OVERCONSUMPTION-ERROR-II.
 - ii. Set $\mathbf{tKey}[\kappa + v] := \perp$ for all $v < \text{ridx}' - \mathbf{p}'$. Set $\overline{\mathbf{tKey}}' = \text{ins}(\overline{\mathbf{tKey}}, \mathbf{newLtKey}, \mathbf{newRtKey})$.
 - iii. Set $\mathbf{outqKey} := \mathbf{oldRKey}_{\mathbf{q}'}$, where $\mathbf{q}' := \mathbf{q}$ but with goto' replacing goto .
 - (b) Set $\mathbf{outrKey} := \mathbf{rKey}_{\text{rec}'}$ where $\text{rec}' := (\text{lidx}', \text{ridx}', \mathbf{oldLKey}', \mathbf{oldRKey}', \overline{\mathbf{tKey}}')$ and output $(\mathbf{outrKey}, \mathbf{outqKey})$.

Figure 9: Formal description of the nonleaf Memory Circuit with key passing.

The algorithm $\text{GData}(1^\kappa, D)$ proceeds as follows. Without loss of generality we assume that $M = 2^d$ (where $M = |D|$) where d is a positive integer. We calculate $\epsilon = \frac{1}{\log M}$. We set $K_0 = M$, and for each $0 < i \in [d+1]$ and set $K_i = \lfloor (\frac{1}{2} + \epsilon) K_{i-1} \rfloor + \kappa$.

1. Let $s \leftarrow \{0, 1\}^\kappa$.
2. Any $\mathbf{dKey}^{d,j,k}$ needed in the computation below is obtained as $F_s(\text{data}||d||j||k)$. Similarly for any i, j, k , $\mathbf{rKey}^{i,j,k} := F_s(\text{rec}||i||j||k)$ and $\mathbf{qKey}^{i,j,k} := F_s(\text{query}||i||j||k)$.

Set

$$\overline{\mathbf{tKey}}^{i,j,0} := \left\{ \underbrace{\left\{ \mathbf{qKey}^{i+1,2j,l} \right\}_{l \in [\kappa]}}_{\text{left}}, \underbrace{\left\{ \mathbf{qKey}^{i+1,2j+1,l} \right\}_{l \in [\kappa]}}_{\text{right}} \right\}.$$

and if $\lfloor (\frac{1}{2} + \epsilon) (k+1) \rfloor > \lfloor (\frac{1}{2} + \epsilon) (k) \rfloor$, then set

$$\mathbf{newLtKey}^{i,j,k} = \mathbf{qKey}^{i+1,2j, \lfloor (\frac{1}{2} + \epsilon)(k+1) \rfloor + \kappa - 1}$$

$$\mathbf{newRtKey}^{i,j,k} = \mathbf{qKey}^{i+1,2j+1, \lfloor (\frac{1}{2} + \epsilon)(k+1) \rfloor + \kappa - 1}$$

, otherwise set $\mathbf{newLtKey}^{i,j,k} = \mathbf{newRtKey}^{i,j,k} = \perp$.

3. For all $j \in [2^d], k \in [K_d]$,
 $\tilde{C}^{d,j,k} \leftarrow \text{GCircuit}(1^\kappa, C^{\text{leaf}}[d, k, \mathbf{dKey}^{d,j,k+1}, \mathbf{qKey}^{d,j,k+1}], \mathbf{dKey}^{d,j,k}, \mathbf{qKey}^{d,j,k})$.
4. For all $i \in [d], j \in [2^i], k \in [K_i]$,
 $\tilde{C}^{i,j,k} \leftarrow \text{GCircuit}(1^\kappa, C^{\text{node}}[i, k, \mathbf{newLtKey}^{i,j,k}, \mathbf{newRtKey}^{i,j,k}, \mathbf{rKey}^{i,j,k+1}, \mathbf{qKey}^{i,j,k+1}], \mathbf{rKey}^{i,j,k}, \mathbf{qKey}^{i,j,k})$.
5. For all $j \in [2^d]$, set $\mathbf{Tab}(d, j) = \mathbf{dKey}_{D[j]}^{d,j,0}$.
6. For all $i \in [d], j \in [2^i]$, set $\mathbf{Tab}(i, j) := \mathbf{rKey}_{\text{rec}^{i,j,0}}^{i,j,0}$, where $\text{rec}^{i,j,0} := (0, 0, \mathbf{qKey}^{i+1,2j,0}, \mathbf{qKey}^{i+1,2j+1,0}, \overline{\mathbf{tKey}}^{i,j,0})$.
7. Output $\tilde{D} := \left(\left\{ \tilde{C}^{i,j,k} \right\}_{i \in [d+1], j \in [2^i], k \in [K_i]}, \left\{ \mathbf{Tab}(i, j) \right\}_{i \in [d+1], j \in [2^i]} \right)$ and s .

Figure 10: Formal description of GData with passed keys.

This way, we can control which keys are present in a circuit. In particular, we drop consumed \mathbf{qKeys} so that no future unevaluated circuit has them. However, we still need to hardwire new \mathbf{qKeys} , but we only do so as needed as part of two new hardwired variables $\mathbf{newLtKey}$ and $\mathbf{newRtKey}$. Then, whenever a circuit outputs some \mathbf{qKey} for a child circuit, it also drops all \mathbf{qKeys} inside \mathbf{tKey} that are older than or equal to \mathbf{qKey} . Note that we still maintain the exact same \mathbf{tKey} size of 2κ keys, so this passing method will never accumulate too many keys. We formalize this construction by modifying how GData and C^{node} work in Figures 9 and 10, and the remainder of the construction is unchanged. It is straightforward to see that this neither affects correctness (only keys corresponding to consumed or soon-to-be-burned circuits will be

dropped) nor asymptotic cost. We argue correctness as follows: within a node, the sequence of oldLKey and oldRKey within the circuits is an increasing sequence, and oldLKey strictly increases if we are going left, and oldRKey strictly increases if we are going right. Note that the $\overline{\text{tKey}}$ shifting corresponds precisely to the previous windows with the only difference being now there could be some keys set to \perp . Thus, the only way this scheme could be incorrect is if we attempt to assign a \perp to oldLKey or oldRKey. But all the \perp values correspond to indexes that are strictly less than the current oldLKey or oldRKey index, and therefore cannot be the new index.

We demonstrate a series of lemmas that ensures that when some circuit needs to be simulated, all appearances of its keys will have been in already been simulated or dropped. This strategy then allows the full simulation proof goes through as we will see in Section 7.

7 Security Proof

In this section we prove the UMA2-security of the black-box garbled RAM (GData, GProg, GInput, GEval).

Theorem 7.1 (UMA2-security). *Let F be a PRF and $(\text{GCircuit}, \text{Eval}, \text{CircSim})$ be a circuit garbling scheme, both of which can be built from any one-way function in black-box manner. Then our construction is a UMA2-secure garbled RAM scheme for uniform access programs running in total time $T < M$ making only black-box access to the underlying OWF.*

Proof. We first prove a lemma (Lemma 7.3) before proving our main theorem. For the lemma, we consider ourselves during the course of GEval, where we are about to evaluate some non-root node $\tilde{C}^{i,j,k}$. Our eventual goal is to show that all instances of $\text{qKey}^{i,j,k}$ are in previously evaluated (hence, will be simulated) circuits, and is not being passed as part of any $\overline{\text{tKey}}$.

Fact 7.2. *The rKey_{rec} to be consumed by $\tilde{C}^{i,j,k}$ was output by $\tilde{C}^{i,j,k-1}$, or initially stored in $\text{Tab}(i, j)$ in the case where $k = 0$. The qKey_q used to evaluate $\tilde{C}^{i,j,k}$ was either output by (Case 1) $\tilde{C}^{i,j,k-1}$ or (Case 2) $\tilde{C}^{i-1, \lfloor j/2 \rfloor, k'}$ for some k' .*

To further pinpoint where qKeys are stored, we group the circuits in the parent node into three groups. We let k'_{\min} be the smallest value such that $\lfloor (\frac{1}{2} + \epsilon)k'_{\min} \rfloor + \kappa - 1 = k$, and k'_{\max} be the largest value such that $\lfloor (\frac{1}{2} + \epsilon)k'_{\max} \rfloor = k$. For a parent circuit $\tilde{C}^{i-1, \lfloor j/2 \rfloor, k'}$, we call it a *past* circuit if $k' < k'_{\min}$, a *future* circuit if $k' > k'_{\max}$, and a *present* circuit if $k'_{\min} \leq k' \leq k'_{\max}$.

We now state our main lemma.

Lemma 7.3. *Suppose during the execution of GEval, we are about to evaluate garbled circuit $\tilde{C}^{i,j,k}$. Let qKey denote $\text{qKey}^{i,j,k}$. Then all instances of qKey exist only in previously evaluated circuits.*

Proof. We let k^* denote the index of the last parent circuit that evaluated prior to our current circuit, i.e. $\tilde{C}^* = \tilde{C}^{i-1, \lfloor j/2 \rfloor, k^*}$ was the last circuit to be evaluated at level $i - 1$. WLOG assume that the current circuit is the left child of the parent. Observe that qKey only occurs in the following locations: the predecessor circuit, inside newLtKey of the final “past” parent, or inside some “current” or “future” parent’s oldLKey or $\overline{\text{tKey}}$. Since the predecessor circuit must be evaluated already, we only need to check the existence of qKey inside one or more of my parent circuits.

Let lidx be the left index (implicitly) passed into \tilde{C}^* , and let lidx' be the left index (implicitly) output by it.

Observe that by definition, qKey is not in the $\overline{\text{tKey}}$ of any past or future parent. In particular, it can only be included inside $\overline{\text{tKey}}$ when being inserted as a newLtKey , and once it is removed it can never be present again in any future parent’s $\overline{\text{tKey}}$. Note that qKey may still be inside newLtKey of a past parent or oldLKey of a future parent. Furthermore, all parent circuits with index $k' \leq k^*$ have been evaluated, and thus we

only need to argue that no (unevaluated) parent circuit $k' > k^*$ contains qKey as either $\overline{\text{tKey}}$, newLtKey , or oldLKey .

We analyse the following six cases:

- [Case 1A] The predecessor circuit $\tilde{C}^{i,j,k-1}$ output my qKey_q , and k^* belongs to a past parent.
- [Case 1B] The predecessor circuit $\tilde{C}^{i,j,k-1}$ output my qKey_q , and k^* belongs to a present parent.
- [Case 1C] The predecessor circuit $\tilde{C}^{i,j,k-1}$ output my qKey_q , and k^* belongs to a future parent.
- [Case 2A] The parent circuit $\tilde{C}^{i-1,\lfloor j/2 \rfloor, k^*}$ output my qKey_q , and k^* belongs to a past parent.
- [Case 2B] The parent circuit $\tilde{C}^{i-1,\lfloor j/2 \rfloor, k^*}$ output my qKey_q , and k^* belongs to a present parent.
- [Case 2C] The parent circuit $\tilde{C}^{i-1,\lfloor j/2 \rfloor, k^*}$ output my qKey_q , and k^* belongs to a future parent.

CASE 1A. This case cannot occur. Since \tilde{C}^* was a past parent, by definition we must have $\lfloor (\frac{1}{2} + \epsilon)k^* \rfloor + \kappa - 1 < k$. Since qKey was passed from the predecessor circuit, it must have taken the branch where $k - 1 < \text{goto}' - 1 = \lfloor (\frac{1}{2} + \epsilon)k^* \rfloor - 1$. Combining these two inequalities yields $\kappa < 2$ which is a contradiction.

CASE 1B, 1C. Note that $k - 1 < \text{goto}' - 1$ still holds. We know that $\text{lid}x' \geq \text{goto}'$, so we have that $\text{lid}x' > k$. Since $\text{oldLKey}'$ is the key for circuit $\text{lid}x'$, it cannot be the qKey for k . Furthermore, all keys inside $\overline{\text{tKey}}$ with index less than $\text{lid}x'$ have been set to \perp by \tilde{C}^* so no unevaluated parent circuit can have the current qKey as part of $\overline{\text{tKey}}$ or oldLKey . Finally, qKey appearing as newLtKey can only occur in a past parent, which has already been evaluated in this case.

CASE 2A. This case cannot occur. By the definition of C^{node} , the only way the parent circuit could output my qKey_q directly is if it is held as oldLKey . However, the oldLKey is only assigned due to the value contained in an older $\overline{\text{tKey}}$ in some parent circuit $k' \leq k^*$. The indices k of any parent $k' \leq k^*$ parents is at most $\lfloor (\frac{1}{2} + \epsilon)k' \rfloor + \kappa - 1 \leq \lfloor (\frac{1}{2} + \epsilon)k^* \rfloor + \kappa - 1 < k$ by definition of k^* being a past parent. Therefore, qKey could not have been output by any past parent circuit.

CASE 2B. In this case, k^* belongs to a present circuit that was evaluated. Note that \tilde{C}^* replaced its $\text{oldLKey} = \text{qKey}$ with some new $\text{oldLKey}'$ which corresponds to the $\text{lid}x'$ -th circuit at level i . Since $k = \text{lid}x < \text{lid}x'$ and all $\overline{\text{tKey}}[v]$ is set to \perp for $v < \text{lid}x' - p'$, the current qKey was removed from $\overline{\text{tKey}}$ by \tilde{C}^* and hence all successor parent circuits' $\overline{\text{tKey}}$ do not contain qKey . Furthermore, oldLKey can only be updated by $\overline{\text{tKey}}$ and \tilde{C}^* does not set the updated $\text{oldLKey}'$ to qKey , and no parent circuit $k' > k^*$ can set oldLKey to qKey since it is no longer contained in any of their $\overline{\text{tKey}}$ values. Finally, qKey appearing as newLtKey can only occur in a past parent, which has already been evaluated in this case.

CASE 2C. Because k^* belongs to a future parent that was evaluated, it must be the case that all past and present parents have already been evaluated. We check that qKey does not exist in any unevaluated parent circuit's $\overline{\text{tKey}}$ or oldLKey : all parent circuits $k' \leq k^*$ have been evaluated, \tilde{C}^* was evaluated and it output and replaced the qKey sitting in oldLKey with some $\overline{\text{tKey}}$. Since \tilde{C}^* and all its successors are future parents, none of them have qKey inside its $\overline{\text{tKey}}$ and thus oldLKey would never contain qKey . Finally, qKey appearing as newLtKey can only occur in a past parent, which has already been evaluated in this case. □

We now proceed to prove the theorem. Let CircSim be the garbled circuit simulator. Suppose in the real execution, a total of w circuits are evaluated by GEval . We construct Sim and then give a series of hybrids $\hat{H}^0, H^0, \dots, H^w, \hat{H}^w$ such that the first hybrid outputs the $(\tilde{D}, \tilde{P}, \tilde{x})$ of the is the real execution and the last hybrid is the output of Sim , which we will define. H^0 is the real execution with the PRF F replaced with a

uniform random function (where previously evaluated values are tabulated). Since the PRF key is not used in evaluation, we immediately obtain $\hat{H}^0 \stackrel{\text{comp}}{\approx} H^0$.

Our goal is to build a garbled memory, program, and input that is indistinguishable from the real one. Since we know exactly the size and running time and memory access, we can allocate the exact correct amount of placeholder garbled circuits, initially set to \perp . The simulator considers the sequence of circuits (starting from 1) that would have been evaluated given MemAccess. This sequence is entirely deterministic and therefore we let S_1, \dots, S_w be this sequence of circuits, e.g. $S_1 = \tilde{C}^0$ (the first CPU circuit), $S_2 = \tilde{C}^{0,0,0}$ (the first root circuit), \dots . The idea is to have H^u simulate the first u of these circuits, and generate all other circuits as in the real execution.

Hybrid Definition: $(\tilde{D}, \tilde{P}, \tilde{x}) \leftarrow H^u$

The hybrid H^u proceeds as follows: For each circuit not in S_1, \dots, S_u , generate it as you would in the real execution, and for each circuit S_u, \dots, S_1 (in that order) we simulate the circuit using CircSim by giving it as output what it would have generated in the real execution or what was provided as the simulated input labels. Note that this may use information about the database D and the input x , and our goal is to show that at the very end, Sim will not need this information.

We now show $H^{u-1} \stackrel{\text{comp}}{\approx} H^u$. There are several cases: when S_u is a non-root node, when S_u is a root node, and when S_u is a CPU step circuit. In the first case, we must argue that one can replace S_u from a real distribution to one output by CircSim. In order to do so, we must show that its input keys are independent of the rest of the output. Its garbled inputs are rKey_{rec} and qKey_{q} . rKey only existed in its predecessor circuit, which was simulated since it was executed in some hybrid $u' < u$. Furthermore, by Lemma 7.3, all instances of qKey only exist in previously evaluated circuits, hence they were also simulated out in some earlier hybrid. Therefore, any distinguisher of H^{u-1} and H^u can be used to distinguish between the output of CircSim and a real garbling.

When S_u is a root node, the only circuit that has its rKey was its predecessor, and the only circuits that have its qKey are its predecessor or the CPU step circuit that invoked it. Both of these circuits were simulated in a earlier hybrid, and so once again any distinguisher of H^{u-1} and H^u can be used to distinguish between the output of CircSim and a real garbling.

Finally, if S_u is a CPU step circuit, the only circuit that has its cpuSKey was its predecessor (or the initial garbled input), but its cpuDKey was passed around across the entire tree starting from its predecessor. However, again, these were all simulated in an earlier hybrid, so again, any distinguisher of H^{u-1} and H^u can be used to distinguish between the output of CircSim and a real garbling.

Finally, we mention how to handle unevaluated circuits in hybrid \hat{H}^w . Note that the security definition of CircSim does not deal with partial inputs, though this can be handled generically as follows. We encrypt each circuit using a semantically secure symmetric key-encryption scheme with a fresh key for each circuit and secret share the key into two portions. We augment each rKey/dKey by giving it one share (in the clear), and the qKey will have the other share. In unevaluated circuits, the qKey never appears, so the secret encryption key is information theoretically hidden, and thus by the semantic security of the encryption scheme, we can replace all unevaluated circuits with encryptions of zero. This is formally stated and proven as Lemma B.1 in Appendix B, and this is precisely what \hat{H}^w utilizes to convert unused circuits into encryptions of zero.

Then our simulator $\text{Sim}(1^\kappa, 1^M, 1^t, y, 1^D, \text{MemAccess} = \{L^\tau, z^{\text{read},\tau}, z^{\text{write},\tau}\}_{\tau=0,\dots,t-1})$ can output the distribution \hat{H}^w without access to D or x . We see this as follows: the simulator, given MemAccess can determine the sequence S_1, \dots, S_w . The simulator starts by simulating all unevaluated circuits by replacing them with encryptions of zero. It then simulates the S_u in reverse order, starting with simulating S_w using the output y , and then working backwards simulates further ones ensuring that their output is set to the appropriate inputs.

□

8 Full Security and Replenishing

In this section, we prove the full security by showing how to compile any UMA2-secure GRAM scheme with statistical ORAM. First, we show how to extend our construction to running times beyond M by replenishing.

8.1 Circuit Replenishing

Although we observed that “dynamic” circuit replenishing is potentially problematic, here we give a method of allowing GProg to replenish circuits. Note that in GData there was no need to generate some fixed amount of circuits at the root, but it was bounded to be proportional to M in order for it to not run too long. However, using the exact same template, an exponential amount of circuits could be generated: as long as the domain of the PRF is not exhausted, one can always generate more circuit labels that follow this exact pattern.

Using this observation, we can then view our circuit replenishing as a way to amortize this process rather than making dynamic replacements on the fly. That is to say, when we make new circuits, they will be concatenated on to the end of the sequence of circuits of each node. In order to replenish, we give a replenishment strategy so that GProg will be augmented to perform the following functionality. A program at time m running for t steps will replenish some $t \cdot \text{poly}(\log M, \log T, \kappa)$ number of circuits in such a way that after M steps, a total of $S = \sum_{i=0}^d (2^i K_i)$ circuits will be replenished, where K_i is as defined in GData, thus providing us with a number of new circuits that is as large as the original number of original circuits provided.

Given a program that starts at time m (which now could be larger than M) and runs for t steps, we know exactly where we are going to replenish because the locations are relative only to $m, m+1, \dots, m+t-1$. For ease of exposition, we assume that these values do not cross a multiple boundary of M and there exists some c such that they are all in the range $cM, \dots, (c+1)M-1$. Specifically, we want to provide roughly $\frac{S \cdot t}{M} = t \cdot \text{poly}(\log M, \log T, \kappa)$ circuits, and so we consider the following array A of values $\left\lceil \frac{S \cdot (m \bmod M)}{M} \right\rceil, \dots, \left\lceil \frac{S \cdot ((m+t-1) \bmod M)}{M} \right\rceil$. This index tells us which node and circuit we should replicate as follows. Order the nodes from left to right, starting from the root down, then we know that the (i, j) -th node has $n(i, j) = j \cdot K_i + \sum_{i'=0}^{i-1} (2^{i'} K_{i'})$ circuits prior to it. Then GProg performs the following: for $l = 0, \dots, |A| - 1$, we will replenish node (i, j) for the maximal node such that $n(i, j) \leq A[l]$. It will set $k = (c+1) \cdot M + A[l] - n(i, j)$, and create a $\tilde{C}^{i,j,k}$ as in GData, generating the keys using the PRF key s .

Thus, after M steps, there will be $K'_i = 2K_i$ circuits at each node in level i , and in general after nM steps, there will be $K'_i = (n+1)K_i$ circuits.

Two facts follow from this replacement strategy. First, the cost of GProg is not asymptotically impacted since there are only $\frac{S}{M}$ is only $\text{poly}(\log M, \log T, \kappa)$ and hence for t steps, the number of new circuits replenished is $t \cdot \text{poly}(\log M, \log T, \kappa)$. Overall, this only adds $t \text{poly}(\log M, \log T, \kappa)$ additional work for GProg. Secondly, we will never run out of circuits using this strategy. Observe that although partial replacements only appear sequentially down the tree whereas the accesses could consume circuits anywhere in the tree, those yet-to-be-replenished leaves would have not run out in any case since the base scheme *without replenishing* does not run out of circuits. After M time steps, each node is now replenished with as many circuits as a fresh GData would have generated, which allows us to proceed indefinitely.

In some previous GRAM schemes [LO14, GHL⁺14], memory refreshes were also done after M reads and writes were done in memory. Unlike the previous constructions, our circuit replenishing works in an amortized way that smoothly adds more circuits rather than having to refresh the entire database at once.

8.2 Compiling with Statistical ORAM to get full security

In order to achieve this, we show how to compile our UMA2-secure GRAM scheme with a statistical ORAM that has uniform access pattern to achieve a secure GRAM scheme. However, schemes such as [SvDS⁺13] are tree-based and have uniform access pattern only on each level of their tree. That is to say, on each level,

the access pattern is uniform, though not necessarily on the entire tree.

Leveled Memory. In order to combat this issue, we make several *independent* copies of memory, each corresponding to a level in the ORAM tree. Then each CPU step will have the key corresponding to the memory block of one level of the ORAM tree. Indeed, this is a generic method of handling leveled memory, though one possible avenue of concrete benefits is “marrying” together the underlying ORAM tree with our GRAM tree. We give a more formal description as follows.

Suppose instead of a single dataset D , there are now datasets D_0, \dots, D_{N-1} that are accessed in order such that at CPU step m , a random location in dataset $D_{m \bmod N}$ is accessed. The GData algorithm now generates a tree for each dataset and the PRF evaluations now also take the dataset into account, namely we can generate s_0, \dots, s_{N-1} using $s_i = PRF_s(i)$, and then generating the keys in dataset D_i using s_i . The one additional change we need to make to GProg is to ensure each CPU step points to a different dataset, and will be consuming circuits at a $1/N$ rate due to the first N steps all referring to the first root circuit of each of the respective N memory trees. Specifically, it sets `qKey` for time step τ to be $F_{s_i}(\text{query}||0||0||\lceil\tau/N\rceil)$.

The last step is to combine ORAM with UMA2-secure GRAM to obtain fully secure GRAM. The proof is nearly identical to extending UMA-secure GRAM to fully secure GRAM, so we paraphrase previous works [LO13b, GH⁺14, GLOS14] and defer this to Appendix A.

Putting it all together, we obtain our main theorem.

Theorem 8.1 (Full security). *Assuming only the existence of one-way functions, there exists a secure black-box garbled RAM scheme for arbitrary RAM programs. The size of the garbled database is $\tilde{O}(|D|)$, size of the garbled input is $\tilde{O}(|x|)$ and the size of the garbled program and its evaluation time is $\tilde{O}(T)$ where T is the running time of program P . Here $\tilde{O}(\cdot)$ ignores $\text{poly}(\log T, \log |D|, \kappa)$ factors where κ is the security parameter. Furthermore, because garbled RAM trivially implies one-way functions, there is a black-box equivalence of the existence of one-way functions and garbled RAM.*

Additional observations. Instead of storing data only at the leaves, we can store the data at all levels of the tree and pull an entire path of values from the tree down into the CPU step. This is conducive to certain ORAM schemes (e.g. Path ORAM [SvDS⁺13]) which also follow this nature and can be used to obtain additional savings. Furthermore, in our construction, one can consider handling a non-uniform distribution of memory accesses. As long as the distribution of leaf accesses are data-independent and known in advance, we can assign to each leaf a probability it is accessed. Then each parent inherits a probability that is the sum of its two childrens’ probabilities. Based on this new distribution, one can provide a different number of circuits per node as according to this distribution. This would lead to a concrete efficiency improvement for GRAM in the case of certain oblivious algorithms with simple CPU steps and access patterns that are distributed in some known fashion.

Acknowledgments

We thank Alessandra Scafuro for her contribution during the initial stages of the project. Work supported in part by NSF grants 09165174, 1065276, 1118126 and 1136174, US-Israel BSF grant 2008411, 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. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014 -11 -1-0392. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

- [ACG⁺14] Prabhanjan Ananth, Nishanth Chandran, Vipul Goyal, Bhavana Kanukurthi, and Rafail Ostrovsky. Achieving privacy in verifiable computation with multiple servers - without FHE and

- without pre-processing. In Hugo Krawczyk, editor, *PKC 2014: 17th International Workshop on Theory and Practice in Public Key Cryptography*, volume 8383 of *Lecture Notes in Computer Science*, pages 149–166, Buenos Aires, Argentina, March 26–28, 2014. Springer, Berlin, Germany.
- [AIK10] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. From secrecy to soundness: Efficient verification via secure computation. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *ICALP (1)*, volume 6198 of *Lecture Notes in Computer Science*, pages 152–163. Springer, 2010.
- [Ajt10] Miklós Ajtai. Oblivious RAMs without cryptographic assumptions. In Leonard J. Schulman, editor, *42nd Annual ACM Symposium on Theory of Computing*, pages 181–190, Cambridge, Massachusetts, USA, June 5–8, 2010. ACM Press.
- [Bea96] Donald Beaver. Correlated pseudorandomness and the complexity of private computations. In *28th Annual ACM Symposium on Theory of Computing*, pages 479–488, Philadelphia, Pennsylvania, USA, May 22–24, 1996. ACM Press.
- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications. In *STOC*, pages 103–112, 1988.
- [BGT14] Nir Bitansky, Sanjam Garg, and Sidharth Telang. Succinct randomized encodings and their applications. Cryptology ePrint Archive, Report 2014/771, 2014. <http://eprint.iacr.org/2014/771>.
- [BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM Conference on Computer and Communications Security*, pages 784–796. ACM, 2012.
- [CHJV14] Ran Canetti, Justin Holmgren, Abhishek Jain, and Vinod Vaikuntanathan. Indistinguishability obfuscation of iterated circuits and RAM programs. Cryptology ePrint Archive, Report 2014/769, 2014. <http://eprint.iacr.org/2014/769>.
- [CLP14] Kai-Min Chung, Zhenming Liu, and Rafael Pass. Statistically-secure ORAM with $\tilde{O}(\log^2 n)$ overhead. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014, Part II*, volume 8874 of *Lecture Notes in Computer Science*, pages 62–81, Kaoshiung, Taiwan, R.O.C., December 7–11, 2014. Springer, Berlin, Germany.
- [CR73] Stephen A. Cook and Robert A. Reckhow. Time bounded random access machines. *J. Comput. Syst. Sci.*, 7(4):354–375, 1973.
- [DMN11] Ivan Damgård, Sigurd Meldgaard, and Jesper Buus Nielsen. Perfectly secure oblivious RAM without random oracles. In Yuval Ishai, editor, *TCC 2011: 8th Theory of Cryptography Conference*, volume 6597 of *Lecture Notes in Computer Science*, pages 144–163, Providence, RI, USA, March 28–30, 2011. Springer, Berlin, Germany.
- [FLS99] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs under general assumptions. *SIAM Journal of Computing*, 29(1):1–28, 1999.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st Annual ACM Symposium on Theory of Computing*, pages 169–178, Bethesda, Maryland, USA, May 31 – June 2, 2009. ACM Press.

- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 1–17, Athens, Greece, May 26–30, 2013. Springer, Berlin, Germany.
- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science*, pages 40–49, Berkeley, CA, USA, October 26–29, 2013. IEEE Computer Society Press.
- [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, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 405–422, Copenhagen, Denmark, May 11–15, 2014. Springer, Berlin, Germany.
- [GHRW14] Craig Gentry, Shai Halevi, Mariana Raykova, and Daniel Wichs. Outsourcing private RAM computation. In *55th Annual Symposium on Foundations of Computer Science*, pages 404–413, Philadelphia, PA, USA, October 18–21, 2014. IEEE Computer Society Press.
- [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 *CCS*, 2012.
- [GKP⁺13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. How to run turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 536–553, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Berlin, Germany.
- [GLOS14] Sanjam Garg, Steve Lu, Rafail Ostrovsky, and Alessandra Scafuro. Garbled RAM from one-way functions. Cryptology ePrint Archive, Report 2014/941, 2014. <http://eprint.iacr.org/2014/941>.
- [GLOV12] Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *53rd Annual Symposium on Foundations of Computer Science*, pages 51–60, New Brunswick, NJ, USA, October 20–23, 2012. IEEE Computer Society Press.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th Annual ACM Symposium on Theory of Computing*, pages 218–229, New York City, New York, USA, May 25–27, 1987. ACM Press.
- [GO96] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *J. ACM*, 43(3):431–473, 1996.
- [Gol87] Oded Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In Alfred Aho, editor, *19th Annual ACM Symposium on Theory of Computing*, pages 182–194, New York City, New York, USA, May 25–27, 1987. ACM Press.
- [GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for np. In *Proceedings of Eurocrypt 2006, volume 4004 of LNCS*, pages 339–358. Springer, 2006.

- [IKLP06] Yuval Ishai, Eyal Kushilevitz, Yehuda Lindell, and Erez Petrank. Black-box constructions for secure computation. In Jon M. Kleinberg, editor, *38th Annual ACM Symposium on Theory of Computing*, pages 99–108, Seattle, Washington, USA, May 21–23, 2006. ACM Press.
- [IKNP03] Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending oblivious transfers efficiently. In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 145–161, Santa Barbara, CA, USA, August 17–21, 2003. Springer, Berlin, Germany.
- [IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st Annual ACM Symposium on Theory of Computing*, pages 44–61, Seattle, Washington, USA, May 15–17, 1989. ACM Press.
- [IR90] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Shafi Goldwasser, editor, *Advances in Cryptology – CRYPTO’88*, volume 403 of *Lecture Notes in Computer Science*, pages 8–26, Santa Barbara, CA, USA, August 21–25, 1990. Springer, Berlin, Germany.
- [LO13a] Steve Lu and Rafail Ostrovsky. Distributed oblivious RAM for secure two-party computation. In Amit Sahai, editor, *TCC 2013: 10th Theory of Cryptography Conference*, volume 7785 of *Lecture Notes in Computer Science*, pages 377–396, Tokyo, Japan, March 3–6, 2013. Springer, Berlin, Germany.
- [LO13b] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 719–734, Athens, Greece, May 26–30, 2013. Springer, Berlin, Germany.
- [LO14] Steve Lu and Rafail Ostrovsky. Garbled RAM revisited, part II. Cryptology ePrint Archive, Report 2014/083, 2014. <http://eprint.iacr.org/2014/083>.
- [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.
- [LP14] Huijia Lin and Rafael Pass. Succinct garbling schemes and applications. Cryptology ePrint Archive, Report 2014/766, 2014. <http://eprint.iacr.org/2014/766>.
- [OS97] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In *29th Annual ACM Symposium on Theory of Computing*, pages 294–303, El Paso, Texas, USA, May 4–6, 1997. ACM Press.
- [Ost90] Rafail Ostrovsky. Efficient computation on oblivious RAMs. In *22nd Annual ACM Symposium on Theory of Computing*, pages 514–523, Baltimore, Maryland, USA, May 14–16, 1990. ACM Press.
- [PF79] Nicholas Pippenger and Michael J. Fischer. Relations among complexity measures. *J. ACM*, 26(2):361–381, 1979.
- [PW09] Rafael Pass and Hoeteck Wee. Black-box constructions of two-party protocols from one-way functions. In Omer Reingold, editor, *TCC 2009: 6th Theory of Cryptography Conference*, volume 5444 of *Lecture Notes in Computer Science*, pages 403–418. Springer, Berlin, Germany, March 15–17, 2009.

- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93, Baltimore, Maryland, USA, May 22–24, 2005. ACM Press.
- [SCSL11] Elaine Shi, T.-H. Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious RAM with $o((\log n)^3)$ worst-case cost. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 197–214, Seoul, South Korea, December 4–8, 2011. Springer, Berlin, Germany.
- [SvDS⁺13] Emil Stefanov, Marten van Dijk, Elaine Shi, Christopher W. Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. Path ORAM: an extremely simple oblivious RAM protocol. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13: 20th Conference on Computer and Communications Security*, pages 299–310, Berlin, Germany, November 4–8, 2013. ACM Press.
- [Wee10] Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *51st Annual Symposium on Foundations of Computer Science*, pages 531–540, Las Vegas, Nevada, USA, October 23–26, 2010. IEEE Computer Society Press.
- [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: 21st Conference on Computer and Communications Security*, pages 191–202, Scottsdale, AZ, USA, November 3–7, 2014. ACM Press.
- [Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 160–164, Chicago, Illinois, November 3–5, 1982. IEEE Computer Society Press.

A UMA2 to Full Security

Lemma A.1. *Assume there exists a UMA2-secure GRAM scheme for programs with (leveled) uniform memory access, and a statistically secure ORAM scheme with (leveled) uniform memory access that protects the access pattern but not the memory contents. Then there exists a fully secure GRAM scheme.*

Proof. We prove the theorem in the case of a single memory D with uniform access to it, and the proof naturally extends to the leveled case. We construct the new GRAM scheme in a black-box manner as follows. Let $(\widehat{\text{GData}}, \widehat{\text{GProg}}, \widehat{\text{GInput}}, \widehat{\text{GEval}})$ be a UMA2-secure GRAM and let $(\text{OData}, \text{OProg})$ be an ORAM scheme. We construct a new GRAM scheme $(\widehat{\text{GData}}, \widehat{\text{GProg}}, \widehat{\text{GInput}}, \widehat{\text{GEval}})$ as follows:

- $\widehat{\text{GData}}(1^\kappa, D)$: Execute $(D^*) \leftarrow \text{OData}(1^\kappa, D)$ followed by $(\tilde{D}, s) \leftarrow \widehat{\text{GData}}(1^\kappa, D^*)$. Output $\widehat{D} = \tilde{D}$ and $\widehat{s} = s$. Note that OData does not require a key as it is a statistical scheme.
- $\widehat{\text{GProg}}(1^\kappa, 1^{\log M}, 1^t, P, \widehat{s}, m)$: Execute $P^* \leftarrow \text{OProg}(1^\kappa, 1^{\log M}, 1^t, P)$ followed by $(\tilde{P}, s^{in}) \leftarrow \widehat{\text{GProg}}(1^\kappa, 1^{\log M}, 1^t, P^*, \widehat{s}, m)$. Output $\widehat{P} = \tilde{P}$, $\widehat{s}^{in} = s^{in}$.
- $\widehat{\text{GInput}}(1^\kappa, x, \widehat{s}^{in})$: Note that x is valid input for P^* . Execute $\tilde{x} \leftarrow \widehat{\text{GInput}}(1^\kappa, x, \widehat{s}^{in})$, and output $\widehat{x} = \tilde{x}$.
- $\widehat{\text{GEval}}^{\widehat{D}}(\widehat{P}, \widehat{x})$: Execute $y \leftarrow \widehat{\text{GEval}}^{\widehat{D}}(\widehat{P}, \widehat{x})$ and output y .

We show that the construction above given by $(\widehat{\text{GData}}, \widehat{\text{GProg}}, \widehat{\text{GInput}}, \widehat{\text{GEval}})$ is a fully secure GRAM scheme. For the remainder of the proof, we consider the following notation. Let P_1, \dots, P_ℓ be any sequence

of programs with polynomially-bounded run-times t_1, \dots, t_ℓ , and let m_j denote the sum of the running times of the first $j - 1$ programs. Let $D \in \{0, 1\}^M$ be any initial memory data, let x_1, \dots, x_ℓ be inputs and $(y_1, \dots, y_\ell) = (P_1(x_1), \dots, P_\ell(x_\ell))^D$ be the outputs given by the sequential execution of the programs on D . Let $(\widehat{D}_0, \widehat{s}) \leftarrow \widehat{\text{GData}}(1^\kappa, D)$, and for $i = 1 \dots \ell$: $(\widehat{P}_i, \widehat{s}_i^{in}) \leftarrow \widehat{\text{GProg}}(1^\kappa, 1^{\log M}, 1^{t_i}, P_i, \widehat{s}, m_i)$, $\widehat{x}_i \leftarrow \widehat{\text{GInput}}(1^\kappa, x_i, \widehat{s}_i^{in})$. Finally, we consider the sequential execution of the garbled programs for $i = 1 \dots \ell$: $y'_i \leftarrow \widehat{\text{GEval}}^{\widehat{D}_{i-1}}(\widehat{P}_i, \widehat{x}_i)$ which updates the garbled database to \widehat{D}_i .

Correctness. Our goal is to demonstrate that

$$\Pr[(y'_1, \dots, y'_\ell) = (y_1, \dots, y_\ell)] = 1.$$

Since $\widehat{\text{GEval}}$ in our construction directly calls the underlying GRAM scheme for evaluation, the correctness of the underlying scheme guarantees that $(y'_1, \dots, y'_\ell) = (P_1^*(x_1), \dots, P_\ell^*(x_\ell))^{D^*}$. Then by the correctness of the ORAM scheme, $(P_1^*(x_1), \dots, P_\ell^*(x_\ell))^{D^*} = (P_1(x_1), \dots, P_\ell(x_\ell))^D = (y_1, \dots, y_\ell)$.

Security. For any programs P_1, \dots, P_ℓ , database D , and inputs x_1, \dots, x_ℓ , let

$$\text{REAL}^{D, \{P_i, x_i\}} = (\widehat{D}_0, \widehat{P}_i, \widehat{x}_{i=1}^\ell)$$

Our goal is to construct a simulator Sim such that for all $D, \{P_i, x_i\}$, we have that $\text{REAL}^{D, \{P_i, x_i\}} \stackrel{\text{comp}}{\approx} \text{Sim}(1^\kappa, 1^M, \{1^{t_i}, y_i\})$. We let OSim be the ORAM simulator, and USim be the simulator for the UMA2-secure GRAM scheme. The procedure Sim proceeds as follows.

1. Compute $(M', \text{MemAccess}) \leftarrow \text{OSim}(1^\kappa, 1^M, \{1^{t_i}, y_i\}_{i=1}^\ell)$. Note that the statistical simulator can only simulate the MemAccess and not D^* (only its size). However, this is fine as USim does not need D^* to simulate since it is a UMA2-simulator.
2. Compute $(\widetilde{D}, \{\widetilde{P}_i, \widetilde{x}_i\}_{i=1}^\ell) \leftarrow \text{USim}(1^\kappa, 1^{M'}, \{1^{t'_i}, y_i\}_{i=1}^\ell, \text{MemAccess})$, where t'_i is the running time of the oblivious program i .
3. Output $(\widehat{D}_0, \widehat{P}_i, \widehat{x}_{i=1}^\ell) = (\widetilde{D}, \{\widetilde{P}_i, \widetilde{x}_i\}_{i=1}^\ell)$.

We now prove the output of the simulator is computationally indistinguishable from the real distribution. For any $D, \{P_i, x_i\}$, we define a series of hybrid distributions $\mathbf{Hyb}_0, \mathbf{Hyb}_1, \mathbf{Hyb}_2$ with $\mathbf{Hyb}_0 = \text{REAL}^{D, \{P_i, x_i\}}$, and $\mathbf{Hyb}_2 = \text{Sim}(1^\kappa, 1^M, \{1^{t_i}, y_i\}_{i=1}^\ell)$, and argue that for $j = 0, 1$ we have $\mathbf{Hyb}_j \stackrel{\text{comp}}{\approx} \mathbf{Hyb}_{j+1}$.

- **Hyb₀**: This is the real distribution $\text{REAL}^{D, \{P_i, x_i\}}$.
- **Hyb₁**: Use the correctly generated (D^*) from $\widehat{\text{GData}}$ and P_i^* from $\widehat{\text{GProg}}$ and execute $(P_1^*(x_1), \dots, P_\ell^*(x_\ell))^{D^*}$ to obtain $\{y_i\}$ and a sequence of memory accesses MemAccess . Run $(\widetilde{D}, \{\widetilde{P}_i, \widetilde{x}_i\}_{i=1}^\ell) \leftarrow \text{USim}(1^\kappa, 1^{M'}, \{1^{t'_i}, y_i\}_{i=1}^\ell, \text{MemAccess})$ and output $(\widehat{D}_0, \widehat{P}_i, \widehat{x}_{i=1}^\ell) = (\widetilde{D}, \{\widetilde{P}_i, \widetilde{x}_i\}_{i=1}^\ell)$.
- **Hyb₂**: This is the simulated distribution $\text{Sim}(1^\kappa, 1^M, \{1^{t_i}, y_i\}_{i=1}^\ell)$.

We now show that adjacent hybrid distributions are computationally indistinguishable.

$\mathbf{Hyb}_0 \stackrel{\text{comp}}{\approx} \mathbf{Hyb}_1$: Let \mathcal{A} be a PPT distinguisher between these two distributions for some $D, \{P_i, x_i\}$. We construct an algorithm \mathcal{B} that breaks the UMA2-security of the underlying GRAM scheme that proceeds as follows. First, \mathcal{B} runs $(D^*) \leftarrow \text{OData}(1^\kappa, D)$, $P_i^* \leftarrow \text{OProg}(1^\kappa, 1^{\log M}, 1^{t_i}, P_i)$ and declares $D^*, \{P_i^*, x_i\}$

as the challenge database, programs and inputs for the UMA2-security GRAM game. The UMA2-security challenger then outputs $(\tilde{D}', \{\tilde{P}'_i, \tilde{x}'_i\}_{i=1}^\ell)$ and \mathcal{B} must output a guess whether it is real or simulated. In order to do so, \mathcal{B} sets $(\hat{D}', \{\hat{P}'_i, \hat{x}'_i\}_{i=1}^\ell) = (\tilde{D}', \{\tilde{P}'_i, \tilde{x}'_i\}_{i=1}^\ell)$ and forwards this as the challenge to \mathcal{A} . \mathcal{B} then outputs the same guess as \mathcal{A} .

Observe that if the UMA challenger outputs the real values, then $(\hat{D}', \{\hat{P}'_i, \hat{x}'_i\}_{i=1}^\ell)$ is distributed identically as if it were generated from \mathbf{Hyb}_0 , and if the UMA challenger outputs simulated values, then $(\hat{D}', \{\hat{P}'_i, \hat{x}'_i\}_{i=1}^\ell)$ is distributed identically as if it were generated from \mathbf{Hyb}_1 . Therefore, \mathcal{A} distinguishes with the same probability as \mathcal{B} , which is negligible by the UMA2-security of the underlying GRAM scheme.

$\mathbf{Hyb}_1 \stackrel{\text{comp}}{\approx} \mathbf{Hyb}_2$: Let \mathcal{A} be a PPT distinguisher between these two distributions for some $D, \{P_i, x_i\}$. We construct an algorithm \mathcal{B} that breaks the security of the underlying ORAM scheme that proceeds as follows. First, \mathcal{B} announces $D, \{P_i, x_i\}$ as the challenge database, programs, and inputs for the ORAM security game. The ORAM challenger then outputs $(\text{MemAccess}')$ which is either real or simulated. Then, \mathcal{B} computes $(y_1, \dots, y_\ell) = (P_1(x_1), \dots, P_\ell(x_\ell))^D$ and runs the UMA2 simulator $(\tilde{D}', \{\tilde{P}'_i, \tilde{x}'_i\}_{i=1}^\ell) \leftarrow \text{USim}(1^\kappa, 1^{M'}, \{1^{t'_i}, y_i\}, \text{MemAccess}')$. Next, \mathcal{B} sets $(\hat{D}', \{\hat{P}'_i, \hat{x}'_i\}_{i=1}^\ell) = (\tilde{D}', \{\tilde{P}'_i, \tilde{x}'_i\}_{i=1}^\ell)$ and forwards this to \mathcal{A} . \mathcal{B} then outputs the same guess as \mathcal{A} .

Observe that if the ORAM challenger outputs the real values, then $(\hat{D}', \{\hat{P}'_i, \hat{x}'_i\}_{i=1}^\ell)$ is distributed identically as if it were generated from \mathbf{Hyb}_1 , and if the ORAM challenger outputs simulated values, then $(\hat{D}', \{\hat{P}'_i, \hat{x}'_i\}_{i=1}^\ell)$ is distributed identically as if it were generated from \mathbf{Hyb}_2 . Therefore, \mathcal{A} distinguishes with the same probability as \mathcal{B} , which is negligible by the security of the underlying ORAM scheme. \square

B Simulating Garbled Circuits With Partial Garbled Inputs

Lemma B.1. *Suppose there exists a garbling scheme $(\text{GCircuit}, \text{Eval})$ with the property that input wire labels are generated independently and uniformly at random, and a semantically-secure symmetric-key encryption scheme $(\mathcal{G}, \mathcal{E}, \mathcal{D})$. There exists a generic transformation into a garbling scheme $(\text{GCircuit}', \text{Eval}')$ with the additional feature that if \tilde{C}' is a garbled circuit and \tilde{x}' is a strict subset of input wire labels, then $(\tilde{x}', \tilde{C}') \stackrel{\text{comp}}{\approx} (r, \mathcal{E}(0))$ where r is uniformly random.*

Proof. We construct $\text{GCircuit}'$ as follows. First, sample a secret key $sk \leftarrow \mathcal{G}(1^\kappa)$ and samples random values r_i it for each input bit i of the circuit subject to $sk = \bigoplus r_i$. The algorithm then calls $\tilde{C} \leftarrow \text{GCircuit}$ and encrypts $\tilde{C}' := \mathcal{E}_{sk}(\tilde{C})$ and outputs \tilde{C}' , and attaches r_i to each zero and one label of lab (note that this slightly changes the definition that labels must be independently uniform, though does not impact the overall solution). Garbling an input bit remains the same. To evaluate a circuit, Eval' takes each r_i from each garbled input label \tilde{x}'_i and computes $sk = \bigoplus r_i$. It then decrypts $\tilde{C} := \mathcal{D}_{sk}(\tilde{C}')$ and then runs $\text{Eval}(\tilde{C}, \tilde{x})$ where \tilde{x} is \tilde{x}' with the additional r_i removed.

Clearly, this does not impact correctness. The simulator is also straightforward: $\text{CircSim}'$ runs the underlying simulator CircSim and encrypts the resulting simulated circuit by itself. Finally, we show that this additional property holds. Suppose we have (\tilde{x}', \tilde{C}') , where \tilde{x}' is now a strict subset of input labels for 0/1. Then each label is just a label of the underlying scheme which we assumed to be uniformly random along with r_i . However, without even a single r_j , the remaining r_i are uniformly random and independent of sk . Hence, \tilde{x}' is indeed uniformly random and independent of anything else, and by semantic security of the encryption scheme, $\tilde{C}' = \mathcal{E}_{sk}(\tilde{C}) \stackrel{\text{comp}}{\approx} \mathcal{E}_{sk}(0)$. \square