

Garbled RAM Revisited

Part II

Steve Lu * Rafail Ostrovsky †

February 5, 2014

Abstract

In EUROCRYPT 2013, Lu and Ostrovsky proposed the notion of Garbled RAM (GRAM) programs. These GRAM programs are analogous to the classic result of Yao’s garbled circuits: a large encrypted memory can first be provided to evaluator, and then a program can separately be garbled and sent to an evaluator to securely execute while learning nothing but the output of the program and its running time. The key feature of GRAM is that it harnesses the natural complexity-theoretic power that Random Access Memory (RAM) programs have over circuits, such as in binary search or randomized algorithms, where program can be executed in a garbled way without “unrolling” it into a circuit. The candidate Lu-Ostrovsky GRAM scheme proposed in that paper is based on the existence of one-way functions, but due to an implicit reliance on a complex “circular” use of Yao garbled circuits, the scheme requires an additional circularity assumptions and may not be secure otherwise.

In this paper, we show how to construct efficient GRAM without circularity and *based solely on the existence of any one-way function*. The novel approach that allows us to break the circularity is a modification of the Goldreich-Goldwasser-Micali (PRF) construction. More specifically, we modify the PRF to allow PRF-keys to be “adaptively revoked” during run-time at the additive cost of roughly $\log n$ per revocation. Then, to improve the overhead of this scheme, we apply a delicate recursion technique that bootstraps mini-GRAM schemes into larger, more powerful ones while still avoiding circularity in the hybrid arguments. This results in secure GRAM with overhead of $\text{poly}(\kappa)(\min(t, n^\epsilon))$ for any constant $\epsilon > 0$ where n is the size of memory and t is the running time.

In a companion work (Part I), Gentry, Halevi, Raykova, and Wichs show an alternative approach using identity-based encryption to solve the circularity problem¹. Their scheme achieves overhead of $\text{poly}(\kappa)\text{polylog}(n)$ assuming the existence of IBE.

Keywords: Secure Computation, Oblivious RAM, Garbled RAM, Garbled Circuits.

1 Introduction

The problem of securely storing, accessing and computing on encrypted data is of increasing importance as our resources are being outsourced, such as in cloud storage and computation. This can be modeled as a general problem of secure computation of some functionality \mathcal{F} . This problem has come from a long line of research, starting with the two-party case of Yao’s garbled circuits [22] and the multi-party case of the

*Department of Computer Science, UCLA. Email: stevelu@cs.ucla.edu. Research supported in part by NSF grants CCF-0916574; IIS-1065276; CCF-1016540; CNS-1118126; CNS-1136174.

†Department of Computer Science and Department of Mathematics, UCLA. Email: rafail@cs.ucla.edu. Research supported in part by NSF grants CCF-0916574; IIS-1065276; CCF-1016540; CNS-1118126; CNS-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 also 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.

¹A merged version of the two works appears in Eurocrypt 2014.

Goldreich-Micali-Wigderson [9] paradigm. In both of these approaches, the functionality is represented as a circuit. Many of the subsequent works, including secure computation via fully homomorphic encryption, requires representing \mathcal{F} as a circuit, and there have only been a handful of results in other models such as branching programs, Turing Machines, RAM programs, and so forth.

On the other hand, many algorithms are naturally represented as RAM programs and simply cannot be represented by circuits without a large blowup. These include programs that run in poly-logarithmic time on large data (e.g. binary search on large data), or Las Vegas style algorithms that can have exponential running time but are efficient on average (e.g. simplex algorithm from linear programming). Performing secure RAM computation was first suggested in the work of Ostrovsky and Shoup [20], and a line of works [17, 12, 15, 16, 10] have considered the model of secure RAM computation.

We reiterate some of the main motivation of previous works of the advantages of using the RAM model of secure computation. In the case of Yao’s garbled circuits [22], you can non-interactively first garble the circuit and send it over during an “offline” phase, then when you have online inputs you can garble the input separately and send it over to the evaluator. The evaluator can then online run the garbled circuit and learn only the output. The Lu-Ostrovsky [16] GRAM construction achieves the same benefits of this non-interactive online/offline garbling paradigm but for RAM programs: you can garble a large memory and send it over, then separately garble a program and inputs that can be run online on this garbled memory. In this paper, we show how to construct a solution that also enjoys these same non-interactive garbling properties from *any one-way function* while avoiding circularity.

Of course, many complex real-world programs are more suitable for the RAM model instead of circuits: unrolling programs with multiple execution paths, recursion, loops, etc. into a circuit is often inefficient and prohibitive. The power of random access means that a program with small running time and code size can still access a large memory (such as in database search), but when represented as a circuit it must be at least as large as the size of the database and thus even in the online/offline model of secure computation, the online cost is still proportional to the database size. The online/offline model is still important since one can offload the cost of pre-processing a large database in preparation for secure computation.

In EUROCRYPT 2013, the work of [16] introduced the notion of Garbled RAM (GRAM) programs and presented a candidate scheme under the minimal assumption of one-way functions. However, there is a subtle “circularity” problem with the given construction, which prevents a proof of security from going through as-is, and would require an additional circularity assumption (which seems difficult to prove) to be secure. In a companion work (Part I), Gentry, Halevi, Raykova, and Wichs [6] proposed a modification to the construction that avoids the circularity in the argument and is secure under the existence of IBE schemes and asymptotically matches the efficiency of the [16] candidate.

In this work, we show how to construct a GRAM scheme that is secure based only on OWFs with a mild drop in overhead compared to the original construction of [16]. We patch the [16] scheme by directly breaking the circularity that lies in the PRF that is built into the construction. The main technique we use to break the circularity is to use a notion of *revocable PRFs* that allows for the “adaptive revocation” of PRF-keys and values. This notion is related to a well-studied theme of removing PRF values in the literature that appear in various forms such as delegatable PRFs [13], functional PRFs [5], or punctured PRFs [21]. That is to say, given a set X of values, and given a PRF key k , there is a “revoke” transformation $k \rightarrow k_X$ such that $\forall y \notin X, F_{k_X}(y) = F_k(y)$ and $\forall x \in X, F_k(x)$ is pseudorandom even given k_X . When k is revoked in such a fashion, we obtain a new key k_X , which as stated above is powerful enough to compute F_k on any value not in X , but still reveals no information about F_k on elements in X . The additional important property is that we can further revoke values, i.e. given any X and k_X , and any $Y \supset X$, there is a transformation $k_X \rightarrow k_Y$.

This construction will be based on the Goldreich-Goldwasser-Micali [8] PRF, and it will allow keys to be revoked at the cost of roughly $\log n$ per revocation. Then, to improve the overhead of this secure scheme, we apply a recursion technique that bootstraps mini-GRAM schemes into larger, more powerful ones. This results in our scheme having an overhead cost which is $\text{poly}(\kappa)(\min(t, n^\epsilon))$ for any constant $\epsilon > 0$ where n is the size of memory and t is the running time.

1.1 Related Work

The notion of Oblivious RAM, which is used to hide the contents and access pattern (namely, the sequence of locations that is read from and written to during the execution of a RAM program) from memory, was introduced in the context of software protection by Goldreich [7] and Ostrovsky [18, 19]. In the original work by Goldreich [7], a solution was given with $O(\sqrt{n})$ and communication overhead where lookups could be done in a single round and $O(2^{\sqrt{\log n \log \log n}})$ communication overhead for a recursive solution. Subsequently, [18, 19] gave a solution with only poly-log overhead and constant client memory (the so-called “hierarchical solution”). These works will be referenced in our construction of our GRAM scheme in order to introduce regularity properties in the access pattern of our programs.

The work of [20] suggested how to perform secure RAM computation based on an oblivious reading and writing scheme. Secure RAM computation was explored in the work of Naor and Nissim [17] using circuits with lookup tables. Gordon et al. [12] proposed a solution in the case of RAM programs with sublinear amortized cost. Namely, consider a client that holds a small input x , and a server that holds a large database D , and the client wishes to repeatedly perform private queries $f(x, D)$. In this model, an expensive initialization (depending only on D) is first performed (say, in an offline phase). Afterwards, if f can be computed in time T with space S with a RAM machine, then there is a secure two-party protocol computing f in time $O(T) \cdot \text{polylog}(S)$ with the client using $O(\log S)$ space and the server using $O(S \cdot \text{polylog}(S))$ space. The interactive secure RAM computation solution of Lu and Ostrovsky [15] gave a construction with lower concrete complexity and can also be viewed as a generalization of the [20] model where servers must also perform sublinear work. Recently, the works [16] and [10] show how to construct non-interactive garbled RAM solutions. The former construction was based off any one-way function (and the additional circularity assumption), and the latter construction achieved full compactness under stronger assumptions.

1.2 Our Results

Our main theorem covers the entire construction which is secure given any one-way function and does not require additional circularity assumptions. We allow for the garbling of memory, inputs, program code², and program CPU steps, and the evaluator will evaluate this garbled information.

Main Theorem (Informal). *Assume one-way functions exist, and let the security parameter be κ . Then, for any initial RAM contents D of size n , programs P_1, \dots, P_ℓ that run on D in at most t CPU steps, there exists a Garbled RAM scheme with $\text{polylog}(\kappa, n)$ overhead in the storage size of the garbled memory contents, and $\text{poly}(\kappa)(\min(t, n^\epsilon))$ overhead in the size of the garbled program and its running time.*

1.3 Remarks

Polynomial vs Exponential Size/Time RAM. Since the RAM CPU is of size at least κ , it can theoretically index up to 2^κ locations. While typically we consider κ , t and n to be polynomially related, we obtain interesting regimes in the case of exponentially (or superpolynomially) large running-times or memory sizes. In particular, if n is, say, exponentially large, we can strengthen our Main Theorem to drive the overhead to be sub-polynomial in n , i.e. $\text{poly}(\kappa)(\min(t, 2^{\sqrt{\log n}}))$.

Reactive Functionalities and Reusability. One of the nice properties that we retain from the construction of [16] is that one can “reuse” the memory once initially garbled. I.e. a sequence of different programs and different inputs can be garbled in sequence that all operate on a dynamically changing memory store that is garbled once (which mimics a reactive functionality). Note that these changes, once applied, cannot be rewound by the evaluator and the way to ensure this to have the garbler maintain an ever-increasing counter. Note that we require only the garbler to choose these new programs and inputs to

²We mention that the program code and input can be thought of as a single object, but we separate them so that we can run multiple programs on the same input or vice versa.

avoid any adaptivity issues that arise in garbling schemes (see Bellare-Hoang-Rogaway [2] for a discussion on this issue). A formal treatment of this notion is provided in Part I [6], and we retain the notation.

An interesting concept introduced in the STOC 2013 paper of Goldwasser et al. [11] is the notion of *token-based* obfuscation, where a reusable garbled circuit can be evaluated on multiple inputs as long as the garbler provides a new “token” for each input. Because our scheme can garble the code of a program and store it ahead of time, we achieve a slightly different notion of token-based obfuscation: we need to garble each input, but we also need to provide garbled CPU steps that allow the evaluator to run a single step of the program, thus resulting in tokens for both inputs and “time quota”.

Worst-case Versus Per-instance Running Time, Universal Programs, and Output Privacy.

As was noted in the CRYPTO 2013 work of Goldwasser et al. [10], the power of secure computation on Turing Machines and RAM programs over that of circuits is that for algorithms with very different worst-case and average-case running times, the circuit must be of worst-case size. Randomized algorithms such as Las Vegas algorithms or even heuristically good-on-average programs would benefit greatly if the online running time of the secure computation ran in time proportional to that particular instance. In our solution, though we have an upper bound T on the number of execution steps of the algorithm which affects the offline time and space, the online evaluation can have a CPU step output “halt” in the clear when the program has halted and the evaluator will then only run in time depending on this particular input.

In order to further mask the program, one can consider a T time-bounded universal program u_T , which takes as input the code of a program π and an input for that program. One can also provide an auxiliary mask so that the output of P is blinded by this value (such a modification has appeared in the literature, see, e.g. [1]).

2 Background

We follow the notational convention that was set forth in Part I [6]. A few of the subsections from this section are restated for the convenience of the reader.

2.1 RAM Model

Notation for RAM Computation. Before we describe *garbled* RAM, let us fix a notation for describing standard RAM computation. We will consider a program P that has random-access to a memory of size n which may initially contain some data $D \in \{0, 1\}^n$. In addition, the program gets a “short” input x , which we can alternatively think of as the initial state of the program. In general, the distinction between what to include in the program P , the memory data D and the short input x can be somewhat arbitrary. We use the notation $P^D(x)$ to denote the execution of such program. The program can read/write to various locations in memory throughout the execution. 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 a useful example to keep in mind throughout this work, 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. A useful representation of a RAM program P is through a small *CPU-Step Circuit* which executes a single CPU step:

$$C_{\text{CPU}}^P(\text{state}, b^{\text{read}}) = (\text{state}', i^{\text{read}}, i^{\text{write}}, b^{\text{write}})$$

This circuit takes as input the current CPU **state** and a bit b^{read} residing in the the last read memory location. It outputs an updated **state'**, the next location to read $i^{\text{read}} \in [n]$, a location to write to $i^{\text{write}} \in [n] \cup \{\perp\}$ (where \perp values are ignored), a bit b^{write} to write into that location.

The computation $P^D(x)$ starts in the initial state $\text{state}_1 = x$, corresponding to the “short input” and by convention we will set the initial read bit to $b_1^{\text{read}} := 0$. In each step j , the computation proceeds by running $C_{\text{CPU}}^P(\text{state}_j, b_j^{\text{read}}) = (\text{state}_{j+1}, i^{\text{read}}, i^{\text{write}}, b^{\text{write}})$. We first read the requested location i^{read} by setting $b_{j+1}^{\text{read}} := D[i^{\text{read}}]$ and, if $i^{\text{write}} \neq \perp$, we write to the location by setting $D[i^{\text{write}}] := b^{\text{write}}$. The value $y = \text{state}$ output by the last CPU step serves as the output of the computation.

We say that a program P has **read-only** memory access, if it never overwrites any values in memory. In particular, using the above notation, the outputs of C_{CPU}^P always set $i^{\text{write}} = \perp$.

2.2 Garbled Circuits

Garbled circuits was first suggested by Yao [22] and subsequently proven secure by Lindell and Pinkas [14]. We review the concept of garbled circuits and refer the reader to the work of Bellare et al. [3] for a thorough treatment of garbling schemes. We continue with the notation of Part I and work with a projective circuit garbling scheme which is a tuple of PPT algorithms $(\text{GCircuit}, \text{Eval})$. Since the scheme is projective, it is helpful to think of individual wire labels for the input so that wire w of the circuit is associated with two labels $\text{lbl}_0^w, \text{lbl}_1^w$ corresponding to the bit-values 0, 1. Finally, since one can apply a generic transformation (see, e.g. [1]) to blind the output, we allow output wires to also have arbitrary labels associated with them. We take the following definitions and notation from Part I:

- $(\tilde{C}, \{ (j, b, \text{lbl}_b^{\text{input},j}) \}) \leftarrow \text{GCircuit}(1^\kappa, C, \{ (i, b, \text{lbl}_b^{\text{output},i}) \})$: Given a circuit C with input size v_{input} and output size v_{output} , and a set of *output labels* $\text{lbl}_b^{\text{output},i}$ for all output wires $i \in [v_{\text{output}}]$ and $b \in \{0, 1\}$, outputs a *garbled circuit* \tilde{C} and a set of *input labels* $\text{lbl}_b^{\text{input},j}$ for every input wire $j \in [v_{\text{input}}]$ and $b \in \{0, 1\}$.
- $(\text{lbl}^{\text{output},1}, \dots, \text{lbl}^{\text{output},v_{\text{output}}}) = \text{Eval}(\tilde{C}, (\text{lbl}^{\text{input},1}, \dots, \text{lbl}^{\text{input},v_{\text{input}}}))$: Given a garbled circuit \tilde{C} and a sequence of input labels $\text{lbl}^{\text{input},j}$, outputs a sequence of output labels $\text{lbl}^{\text{output},i}$. Intuitively, if the input labels correspond to some input $x \in \{0, 1\}^{v_{\text{input}}}$ then the output labels should correspond to $y = C(x)$.

Correctness. For correctness, we require that for any circuit C and any input $x \in \{0, 1\}^{v_{\text{input}}}, x = (x[1], \dots, x[v_{\text{input}}])$ such that $y = (y[1], \dots, y[v_{\text{output}}]) = C(x)$ and any set of output labels $\{ (i, b, \text{lbl}_b^{\text{output},i}) \}$ we have

$$\Pr \left[\text{Eval}(\tilde{C}, (\text{lbl}_{x[1]}^{\text{input},1}, \dots, \text{lbl}_{x[v_{\text{input}}]}^{\text{input},v_{\text{input}}})) = (\text{lbl}_{y[1]}^{\text{output},1}, \dots, \text{lbl}_{y[v_{\text{output}}]}^{\text{output},v_{\text{output}}}) \right] = 1.$$

where $(\tilde{C}, \{ (j, b, \text{lbl}_b^{\text{input},j}) \}) \leftarrow \text{GCircuit}(C, \{ (i, b, \text{lbl}_b^{\text{output},i}) \})$.

Security. For security, we require that there is a PPT simulator Sim such that for any $C, x, \{ (i, b, \text{lbl}_b^{\text{output},i}) \}$ as above, we have

$$\left(\tilde{C}, \text{lbl}_{x[1]}^{\text{input},1}, \dots, \text{lbl}_{x[v_{\text{input}}]}^{\text{input},v_{\text{input}}} \right)^{\text{comp}} \approx \text{Sim}(1^\kappa, C, \text{lbl}_{y[1]}^{\text{output},1}, \dots, \text{lbl}_{y[v_{\text{output}}]}^{\text{output},v_{\text{output}}})$$

where $(\tilde{C}, \{ (j, b, \text{lbl}_b^{\text{input},j}) \}) \leftarrow \text{GCircuit}(C, \{ (i, b, \text{lbl}_b^{\text{output},i}) \})$, $y = C(x)$.

2.3 Goldreich-Goldwasser-Micali Pseudorandom Function

We review the GGM [8] PRF construction and make an observation that will be useful for us later on. Let G be a PRG that stretches from λ bits to 2λ bits, and we can write G_0 to denote the left half of the output

and G_1 to denote the right half. We can write G_{00} to denote $G_0 \circ G_0$, and similarly for G_{01}, G_{10}, G_{11} . Indeed for any bitstring x , we can write G_x to denote repeatedly going left/right using x as the indicator.

Then the GGM construction of a PRF F from G is as follows. Suppose k is the seed of the PRF, then $F_k(x)$ is defined to be $G_x(k)$. This can be viewed as evaluating G on the “root” of a tree consisting of k , and each bit of the input x defines whether to go left or right on the tree. The security property is that any PPT algorithm \mathcal{A} given oracle access to $F_k(\cdot)$ should behave as if it were interacting with a random oracle $R(\cdot)$:

$$\left| \Pr[\mathcal{A}^{F_k(\cdot)}(1^\lambda) = 1; k \leftarrow \{0, 1\}^\lambda] - \Pr[\mathcal{A}^{R(\cdot)}(1^\lambda) = 1] \right| < \epsilon$$

with ϵ being negligible.

2.4 Garbled RAM

We will right-away consider a scenario where 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 stress that each garbled program \tilde{P}_i can only be executed on a *single* garbled input \tilde{x}_i . In other words, although the garbled data is reusable and allows for the execution of many programs, the garbled programs are *not* reusable. The programs can only be executed in the specified order and are not “interchangeable”. Therefore, they cannot be garbled completely independently. In our case, we will assume that the garbling procedure of each program P_i gets t^{init} which is the total number of CPU steps executed so far by P_1, \dots, P_{i-1} and t^{cur} which is the number of CPU steps to be executed by P_i .

Syntax & Efficiency. A *garbled RAM* scheme consists of four procedures: (GData, GProg, GInput, GEval) with the following syntax:

- $\tilde{D} \leftarrow \text{GData}(D, k)$: Takes memory data $D \in \{0, 1\}^n$ and a key k . Outputs the garbled data \tilde{D} .
- $(\tilde{P}, k^{\text{input}}) \leftarrow \text{GProg}(P, k, n, t^{init}, t^{cur})$: Takes a key k and a description of a RAM program P with memory-size n and run-time consisting of t^{cur} CPU steps. In the case of garbling multiple programs, we also provide t^{init} indicating the cumulative number of CPU steps executed by all of the previous programs. Outputs a garbled program \tilde{P} and an input-garbling-key k^{input} .
- $\tilde{x} \leftarrow \text{GInput}(x, k^{\text{input}})$: Takes an input x and input-garbling-key k^{input} 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 computes the output $y = P^D(x)$. We model GEval itself as a RAM program that can read and write to arbitrary locations of its memory initially containing \tilde{D} .

For **efficiency**, we require that the run-time of GProg, and GEval is $|C_{\text{CPU}}^P| \cdot t^{cur} \cdot \text{poly}(\kappa) \cdot \text{polylog}(n)$, 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 $n \cdot \text{poly}(\kappa)$, which also serves as an upper bound on the size of \tilde{D} .

Correctness & Security. To define the correctness and security requirements of garbled RAMs, let P_1, \dots, P_ℓ be any sequence of programs with polynomially-bounded run-times t_1, \dots, t_ℓ . Let $D \in \{0, 1\}^n$ 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.

Consider the following experiment: choose a key $k \leftarrow \{0, 1\}^\kappa$, $\tilde{D} \leftarrow \text{GData}(D, k)$ and for $i = 1, \dots, \ell$:

$$(\tilde{P}_i, k_i^{\text{input}}) \leftarrow \text{GProg}(P_i, n, t_i^{init}, t_i, k), \tilde{x}_i \leftarrow \text{GInput}(x_i, k_i^{\text{input}})$$

where $t_i^{init} := \sum_{j=1}^{i-1} t_j$ denotes the run-time of all programs prior to P_i . Let

$$(y'_1, \dots, y'_\ell) = (\text{GEval}(\tilde{P}_1, \tilde{x}_1), \dots, \text{GEval}(\tilde{P}_\ell, \tilde{x}_\ell))^{\tilde{D}},$$

denotes the output of evaluating the garbled programs sequentially over the garbled memory.

We require that the following properties hold:

- **Correctness:** We require that $\Pr[y'_1 = y_1, \dots, y'_\ell = y_\ell] = 1$ in the above experiment.
- **Security:** we require that there exists a universal simulator Sim such that:

$$(\tilde{D}, \tilde{P}_1, \dots, \tilde{P}_\ell, \tilde{x}_1, \dots, \tilde{x}_\ell) \stackrel{\text{comp}}{\approx} \text{Sim}(1^\kappa, \{P_i, t_i, y_i\}_{i=1}^\ell, n).$$

Our security definition is non-adaptive: the data/programs/inputs are all chosen ahead of time. This makes our definitions/analysis simpler and also matches the standard definitions for our building blocks such as ORAM. However, there does not seem to be any inherent hurdle to allowing each subsequent program/input (P_i, x_i) to be chosen adaptively after seeing $\tilde{D}, (\tilde{P}_1, \tilde{x}_1), \dots, (\tilde{P}_{i-1}, \tilde{x}_{i-1})$.

Security with Unprotected Memory Access (UMA). We also consider a weaker security notion, which we call security with *unprotected memory access* (UMA). In this variant, the attacker may learn the initial contents of the memory D , as well as the complete memory-access pattern throughout the computation including the locations being read/written and their contents. In particular, we let $\text{MemAccess} = \{(i_j^{\text{read}}, i_j^{\text{write}}, b_j^{\text{write}}) : j = 1, \dots, t\}$ correspond to the outputs of the CPU-step circuits during the execution of $P^D(x)$. For security with unprotected memory access, we give the simulator the additional values $(D, \text{MemAccess})$. Using the notation from above, we require:

$$(\tilde{D}, \tilde{P}_1, \dots, \tilde{P}_\ell, \tilde{x}_1, \dots, \tilde{x}_\ell) \stackrel{\text{comp}}{\approx} \text{Sim}(1^\kappa, \{P_i, t_i, y_i\}_{i=1}^\ell, D, \text{MemAccess}, n).$$

2.5 Review of the Lu-Ostrovsky Construction

A thorough review and intuition of the candidate construction of [16] is given in Part I. Here, we restate the relevant technical details which will be used in our new construction. This construction is for GRAM with UMA-security on programs with so-called “predictably timed writes” which we define below.

Predictably Timed Writes. As a first step, we describe how to incorporate a limited form of writing to memory, which we call *predictably timed writes* (ptWrites). On a high level, this means that whenever we want to *read* some location i in memory, it is easy to figure out the time (i.e., CPU step) j in which that location was last *written* to, given only the current state of the computation and without reading any other values in memory. We will later describe how to upgrade a solution for ptWrites to one that allows arbitrary writes. We give a formal definition of ptWrites below:

Definition 2.1 (Predictably Timed Writes (ptWrites)). *A program execution $P^D(x)$ has predictably timed writes (ptWrites) if there exists a poly-size circuit WriteTime such that the following holds for every CPU step $j = 1, \dots, t$. Let the inputs/outputs of the j th CPU step be $C_{\text{CPU}}^P(\text{state}_j, b_j^{\text{read}}) = (\text{state}_{j+1}, i_j^{\text{read}}, i_j^{\text{write}}, b_j^{\text{write}})$. Then, $u = \text{WriteTime}(j, \text{state}_j, i_j^{\text{read}})$ is the largest value of $u < j$ such that the CPU step u wrote to location i_j^{read} ; i.e., $i_u^{\text{write}} = i_j^{\text{read}}$. We also define a ptWrites property for a sequence of program executions $(P_1(x_1), \dots, P_\ell(x_\ell))^D$ if the above property holds for each CPU step in the sequence.*

Garbled Data. The garbled data \tilde{D} consists of n secret keys for some symmetric-key encryption scheme. For each bit $i \in [n]$ of the original data D , the garbled data \tilde{D} contains a secret key sk_i . The secret keys are chosen pseudo-randomly using a pseudo-random function (PRF) family F_k via $\text{sk}_i = F_k(u, i, D[i])$ (where u is the last CPU step in which it was written). Initially $u = 0$, so given k , there are two possible values $\text{sk}_{(i,0)} = F_k(0, i, 0)$ and $\text{sk}_{(i,1)} = F_k(0, i, 1)$ that can initially reside in $\tilde{D}[i]$ depending on the bit $D[i]$ of the original data, and we set $\tilde{D}[i] = \text{sk}_{(i,D[i])}$.

Garbled Program (Technical). We define an augmented CPU-step circuit $C_{\text{CPU}^+}^P$ which gets as input $(\text{state}, b^{\text{read}})$ and outputs $(\text{state}', i^{\text{read}}, i^{\text{write}}, \text{sk}^{\text{write}}, \text{translate})$. It contains some hard-coded parameters $(j, k, r_0, r_1, \text{lbl}_0^{(\text{read})}, \text{lbl}_1^{(\text{read})})$, where j is the current step, and performs the following computation:

- $(\text{state}', i^{\text{read}}, i^{\text{write}}, b^{\text{write}}) = C_{\text{CPU}}^P(\text{state}, b^{\text{read}})$ are the outputs of the basic CPU-step circuit.
- $\text{translate} = (\text{ct}_0, \text{ct}_1)$ consists of two ciphertexts, computed as follows. For $b \in \{0, 1\}$, first compute $\text{sk}_{(i,b)} := F_k(u, i, b)$ for $i = i^{\text{read}}$, where $u \leftarrow \text{WriteTime}$. Then set $c_b = \text{Enc}_{\text{sk}_{(i,b)}}(\text{lbl}_b^{(\text{read})}; r_b)$ where Enc is a symmetric key encryption and r_b is the encryption randomness.
- If $i^{\text{write}} \neq \perp$, set $\text{sk}^{\text{write}} = F_k(j, i^{\text{write}}, b^{\text{write}})$.

The garbled program \tilde{P} consists of t garbled copies of this augmented CPU-step circuit $\tilde{C}_{\text{CPU}^+}^P(j)$. We start garbling from the end $j = t$. Each garbled circuit $\tilde{C}_{\text{CPU}^+}^P(j)$ outputs the values $i^{\text{read}}, \text{translate}$ in the clear and the updated state' is garbled with the same labels as the input state in the next circuit $\tilde{C}_{\text{CPU}^+}^P(j+1)$; the last circuit outputs state' in the clear as the output of the computation. Each garbled circuit $\tilde{C}_{\text{CPU}^+}^P(j)$ contains hard-coded values $(j, k, r_0^{(j)}, r_1^{(j)}, \text{lbl}_0^{(\text{read}, j+1)}, \text{lbl}_1^{(\text{read}, j+1)})$ which are used to compute the translation mapping translate as described above. The key k is the PRF key which was used to garbled the memory data. The values $r_0^{(j)}, r_1^{(j)}$ are fresh encryption random coins, and $\text{lbl}_0^{(\text{read}, j+1)}, \text{lbl}_1^{(\text{read}, j+1)}$ are the labels of the input-wire for the bit b^{read} in the garbled circuit $\tilde{C}_{\text{CPU}^+}^P(j+1)$.

Garbled Input & Evaluation. The garbled input \tilde{x} consists of the wire-labels for the value $\text{state}_1 = x$ for the garbled circuit $\tilde{C}_{\text{CPU}^+}^P(j=1)$. The evaluator simply evaluates the garbled augmented CPU-step circuits one by one starting from $j = 1$. It can evaluate the first circuit using only \tilde{x} , and gets out a garbled output state_2 along with the values $(i^{\text{read}}, \text{translate} = (c_0, c_1))$ in the clear. The evaluator looks up the secret key $\text{sk} := \tilde{D}[i^{\text{read}}]$ and attempts to use it to decrypt c_0 and c_1 to recover a label $\text{lbl}^{(\text{read}, j=2)}$. In case of a write operation, it writes the new sk^{write} to the indicated location. The evaluator then evaluates the second garbled circuit $\tilde{C}_{\text{CPU}^+}^P(j=2)$ using the garbled input state_2 and the wire-label $\text{lbl}^{(\text{read}, j=2)}$ for the wire corresponding to the bit b^{read} . This process continues until the last circuit $j = t$ which outputs state' in the clear as the output of the computation.

Circularity in the Security Analysis. There is a complex circularity as seen below:

1. In order to argue that the evaluator does not learn anything about the “other” label $\text{lbl}_1^{\text{read}}$, we need to rely on the security of the ciphertext ct_1 .
2. In order to rely on the security of the ciphertext ct_1 we need to argue that the attacker does not learn the decryption key $\text{sk}_{(i,1)} = F_k(i, 1)$, which requires us to argue that the attacker does not learn the PRF key k . However, the PRF key k is contained as a hard-coded value of the second garbled circuit $\tilde{C}_{\text{CPU}^+}^P(j=2)$ and all future circuits as well. Therefore, to argue that the attacker does not learn k we need to (at the very least) rely on the security of the second garbled circuit.
3. In order to use the security of the second garbled circuit $\tilde{C}_{\text{CPU}^+}^P(j=2)$, we need to argue that the evaluator only gets one label per wire, and in particular, we need to argue the the evaluator does not have the “other” label $\text{lbl}_1^{\text{read}}$. But this is what we wanted to prove in the first place!

For the rest of this paper, we will show how to circumvent this issue while still using only the minimal assumption of the existence of one-way functions.

3 Warm-up Read-only Construction

The main problem that arises in the circularity is that there is only one PRF key, and that this key when embedded in any future time step is able to decode anything the circuit does in the current time step. The

intuitive way to circumvent this is to iteratively weaken the PRF key. In order to do so, we introduce the following notion of revocable PRFs.

3.1 revocable PRFs

We define the notion of (adaptively) revocable PRFs and we explain how it differs from existing notions such as [4, 5, 13, 21]. The idea is that we can revoke values from the key so that the PRF cannot be evaluated on these values, and given an already-revoked key, one can further revoke new values.

Definition 3.1. *A revocable PRF is a PRF F equipped with an additional revoke algorithm Rev . Given any key k_X , where X is the set of revoked values (with $X = \emptyset$ if k is a fresh key), $k_Y \leftarrow \text{Rev}(k_X, Y)$ is a new revoked key relative to the set of values $Y \supset X$ satisfying the following properties:*

Correctness: $F_{k_Y}(x) = F_{k_X}(x)$ for all $x \notin Y$, and \perp otherwise.

Pseudorandomness: Given any set of keys $\{k_{Y_1}, \dots, k_{Y_m}\}$, $F_k(x)$ is pseudorandom for all $x \notin \bigcap_i Y_i$.

Note that this definition appears similar to constrained PRFs [4]; however, we do not require that the revoked set to be hidden in any way, and we allow for keys to be adaptively revoked further after the initial fresh key has been operated on. Despite these differences, the concrete GGM-based instantiation that we use is congruent to existing works, and we remind the reader of how it goes.

We start with the simple case of removing a single value x from the domain of F_k . Suppose we write $x = x^0x^1 \dots$ as the bits of x . Instead of giving out k , we can give out the set of values $G_{1-x^0}(k), G_{x^0||1-x^1}(k), \dots$, which can be viewed as the siblings of all the nodes in the path from the root to the leaf x in the GGM tree. Note that any value $y \neq x$ can still be evaluated by this key since we have the root of the subtree corresponding to the first bit on which x and y differ.

Let $X = \{x_1, x_2, \dots, x_s\}$ be a set of s distinct values that we want to “revoke” from this PRF. We can produce a new key k_X relative to the set X that contains at most $s \cdot |x|$ values (corresponding to the values on the vertices of the minimal tree cover of the set $\{0, 1\}^{|x|} \setminus X$). Furthermore, if we are given k_X and $Y \supset X$, we can apply another revocation to this key to get the new key k_Y by removing the roots of subtrees powerful enough to compute $y \in Y$, but adding in the siblings of the subtrees on paths to y .

3.2 Garbling a Read-Only, Read-Once Program

We begin with a warmup construction that gives intuition as to how our full solution works. Consider a RAM program that is read-only, and it only reads each memory location at most once. Then for each CPU step, after we publish the translation table `translate`, the augmented CPU revokes the two values corresponding exactly to that table, namely $\{(u, i, 0), (u, i, 1)\}$ where i is the memory location to be read (and since this is read-only, u is always 0). Indeed, let F be a revocable PRF, and let $(\text{GCircuit}, \text{CircEval})$ be a projective circuit garbling scheme. Let k be the master PRF key that we choose upon initiation.

Garbled Data. The garbled data \tilde{D} is an array of secret keys, which are just outputs of F as we will describe.

- $\tilde{D} \leftarrow \text{GData}(D, k)$: For each $i \in [n]$, set $\tilde{D}[i] := \text{sk}$ where $\text{sk} = F_k(0, i, D[i])$.

Garbled Program. We describe the augmented-CPU-step circuit $C_{\text{CPU}+}^P$ for the program P in Figure 1.

The garbled RAM program for P will consist of t copies of a garbled augmented-CPU-Step circuit $\tilde{C}_{\text{CPU}+}^P(j)$. As before, the labels for the output wires in each circuit are chosen carefully so that some wire values are revealed in the clear while others remain garbled for the next circuit.

Input: $(\mathbf{state}, b^{\text{read}}, X, k_X)$ **Output:** $(\mathbf{state}', i^{\text{read}}, \text{translate}, X', k_{X'})$
Hard-Coded Parameters: $j, \text{lbl}_0^{\text{read}}, \text{lbl}_1^{\text{read}},$

The circuit $C_{\text{CPU}+}^P$ performs the following computation:

- $(\mathbf{state}', i^{\text{read}}) := C_{\text{CPU}}^P(\mathbf{state}, b^{\text{read}})$ are the outputs of the basic CPU-step circuit.
- Set $u := 0$. For $b \in \{0, 1\}$, compute $\text{ct}_b := F_k(u, i^{\text{read}}, b) \oplus \text{lbl}_b^{\text{read}}$. Revoke/puncture $(u, i^{\text{read}}, 0)$ and $(u, i^{\text{read}}, 1)$ from k_X and insert them into X to obtain $k_{X'}$ and X' respectively. Set $\text{translate} := (\text{ct}_0, \text{ct}_1)$.

Figure 1: The Augmented CPU-Step Circuit

- $(\tilde{P}, k^{\text{input}}) \leftarrow \text{GProg}(P, k, n, t^{\text{init}}, t^{\text{cur}})$: Let $t^{\text{max}} := t^{\text{init}} + t^{\text{cur}}$. We let j count down from $j = t^{\text{max}}$ to $t^{\text{init}} + 1$. For each j , we garble $C_{\text{CPU}+}^P$ by calling $\tilde{C}_{\text{CPU}+}^P(j) \leftarrow \text{GCircuit}(1^\kappa, C_{\text{CPU}+}^P, \overline{\text{lbl}})$ where the output labels $\overline{\text{lbl}}$ are chosen as follows:

- The outputs $i^{\text{read}}, \text{translate}, X'$ are given out in the clear (in fact, X' can be inferred simply from the access pattern). For the last circuit $j = t^{\text{max}}$, we do not provide these outputs, but instead provide the output \mathbf{state}' in the clear and this serves as the output of the computation. This completely fixes all of the output-wire labels for that circuit.
- For $j \neq t^{\text{max}}$, the labels of the output wires corresponding to \mathbf{state}', k', X' are set to match the labels of the input wires corresponding to \mathbf{state}, k, X in circuit $\tilde{C}_{\text{CPU}+}^P(j + 1)$.
- For the initial circuit $j = t^{\text{init}} + 1$, we also hard-code the input bit b^{read} to 0, and hard-code k to be the master PRF key k , and X to be the empty set.

For $j \neq t^{\text{max}}$, let $\text{lbl}_0^{(\text{read}, j+1)}, \text{lbl}_1^{(\text{read}, j+1)}$ be the labels of input wire for the bit b^{read} in the $(j + 1)$ st garbled circuit. The j th garbled circuit contains the hard-coded secret values:

$$(j, \text{lbl}_0^{(\text{read}, j+1)}, \text{lbl}_1^{(\text{read}, j+1)}).$$

Set $\tilde{P} := [\tilde{C}_{\text{CPU}+}^P(t^{\text{init}} + 1), \dots, \tilde{C}_{\text{CPU}+}^P(t^{\text{max}})]$, and set $k^{\text{input}} := \{(i, b, \text{lbl}_b^{\text{input}, i}) : i \in [v], b \in \{0, 1\}\}$ to consist of all of the input-wire labels for the v input wires corresponding to the input in the initial circuit $\tilde{C}_{\text{CPU}+}^P(t^{\text{init}} + 1)$.

Garbled Input. Finally, the garbled input \tilde{x} is created the same way as in garbled circuits. It simply consists of the subset of labels of $k^{\text{input}} := \{(i, b, \text{lbl}_b^{\text{input}, i}) : i \in [v], b \in \{0, 1\}\}$ corresponding to the bits of x .

- $\tilde{x} \leftarrow \text{GInput}(x, k^{\text{input}})$: Parse $k^{\text{input}} := \{(i, b, \text{lbl}_b^{\text{input}, i}) : i \in [v], b \in \{0, 1\}\}$ and output $\tilde{x} = (\text{lbl}_{x[1]}^{\text{input}, 1}, \dots, \text{lbl}_{x[v]}^{\text{input}, v})$ where $x[i]$ denotes the i th bit of x and $v := |x|$.

Evaluation. To run $y = \text{GEval}^{\tilde{D}}(\tilde{P}, \tilde{x})$: parse $\tilde{P} = [\tilde{C}_{\text{CPU}+}^P(1), \dots, \tilde{C}_{\text{CPU}+}^P(t)]$ as consisting of t garbled circuits. The evaluator evaluates the circuits one-by-one. Set $\tilde{x}_1 := \tilde{x}$. For $j = 1, \dots, t$:

- When $j = 1$, run $\text{CircEval}(\tilde{C}_{\text{CPU}+}^P(1), \tilde{x}_1)$ else run $\text{CircEval}(\tilde{C}_{\text{CPU}+}^P(j), (\tilde{x}_j, \text{lbl}^{\text{read}, j}, \tilde{k}, \tilde{X}))$ where \tilde{x}_j consists of labels for the garbled \mathbf{state}_j , \tilde{k} consists of labels for the garbled k , and \tilde{X} consists of labels for the garbled X , and $\text{lbl}^{\text{read}, j}$ is a label for the read-bit in the j th CPU circuit. This reveals the outputs $i_j^{\text{read}}, \text{translate}_j, X'$ in the clear. For $j < t$ it also reveals the garbled output corresponding to the labels of $k_{X'}$ and \mathbf{state}_{j+1} for circuit $j + 1$. For $j = t$ it reveals the output of the computation $y = \mathbf{state}_{t+1}$ in the clear.

- Look up $sk = \tilde{D}[i_j^{\text{read}}]$ and decrypt the corresponding row in translate_j to obtain $|\text{bl}|^{\text{read},j+1}$.

Theorem 3.2 (Read-only Warmup to Main Theorem). *Assume F is a revocable PRF, concretely the revocable GGM-PRF, and suppose that $(\text{GCircuit}, \text{CircEval})$ be a projective circuit garbling scheme with wire labels with security parameter κ (both of which can be constructed from any one-way function). Then the warm-up RAM construction above is a read-only, read-once GRAM satisfying UMA-security. Furthermore the space overhead of the garbled memory is $\text{poly}(\kappa)$, and the space overhead of the garbled CPU steps, and the running time overhead are $\text{poly}(\kappa)\text{polylog}(n) \cdot t$ where t is the running time.*

Proof Sketch.

The overhead of the garbled memory is obviously the size of the outputs of the PRF which is $\text{poly}(\kappa)$. However, because each revocation increases the size of the key by a factor of $\log n$, after t steps the size of this key will be $t \log n$.

We provide a sketch at a very high level to provide the intuition for the rest of our construction. This intuition will be used for our full construction, and parallels the intuition for the alternative scheme found in Part I [6]. Since we are dealing with UMA-security, the simulator is allowed to get both the contents of memory and the access pattern, the simulator can pick its own master PRF key k and almost generate an honest transcript by using the underlying garbled circuit simulator for every step. The only real difficulty is in simulating the translation table, which in a real execution contains the wire labels for both zero and one for the next CPU step. Since a simulated garbled circuit will only return one wire label, we must set that to be the one corresponding to the bit we want, and the other one must be fabricated. Indeed, the simulator sets this other dummy key to be all zeroes in the simulated execution.

Then in order to prove security, we set up a series of hybrids \mathbf{Hyb}_j , where in hybrid j , garbled circuits $1, \dots, j$ are created as in the simulation and garbled circuits $j + 1, \dots, t^{\text{max}}$ are created as in the real distribution. We then have a series of companion hybrids \mathbf{Hyb}'_j , where in \mathbf{Hyb}'_j the wire label unused row in the translation table is set to the correct label instead of all zeroes as in the simulation.

The proof then proceeds as follows. $\mathbf{Hyb}_j \stackrel{\text{comp}}{\approx} \mathbf{Hyb}'_{j+1}$ directly due to the security of garbled circuits. In order to show $\mathbf{Hyb}_j \stackrel{\text{comp}}{\approx} \mathbf{Hyb}'_j$, we rely on the security of revocable PRFs. Any distinguisher \mathcal{A} can be used to construct a pseudorandomness breaker \mathcal{A}' for F : since \mathbf{Hyb}_j can be generated using only a revoked key, \mathcal{A}' invoke the PRF oracle to populate the entries of \tilde{D} that have been revoked and plant the PRF challenge in the revoked, unused row of the translation table in CPU step j . If the PRF challenge was purely random, the two distributions are identical, and if it were pseudorandom, the two distributions are as in the hybrid, thus \mathcal{A}' conveys any advantage \mathcal{A} has. □

4 GRAM With Reading and Writing From OWF

Notice that our scheme is not only read-only, it is read-once. This is due to the fact that once an element has been read, it gets revoked, so no future time step could ever generate a valid symmetric-key encryption for that entry in memory again. A tantalizing solution to this problem is to use a key mechanism that allows the CPU to encrypt but not decrypt something, i.e. a public-key encryption scheme. The two main issues are that 1) there is a black-box separation between OWFs and basic public-key encryption, and 2) the public key needs to be compact and yet be able to encrypt to many possible secret keys. If we insist on assuming only OWFs, it seems that 1) is difficult to overcome, however there is a beautiful resolution using identity-based encryption for issue 2 in the companion work of Part I [6].

We resolve this issue of reading multiple times via repetition. By maintaining p copies of the database, we can read from any location up to p times. If we restrict ourselves to RAM programs that are read-only and read- p , we can now obtain GRAM with the overhead of another p . Note that this seems like a strong restriction, but when we include RAM writing, we can relax the restriction to reading at most p times *between* writes to a location. We can trivially relax this condition by overwriting a location after it has

been read p times. We define the following notion of bounded reads for RAM programs (with read and write).

Bounded Reads. We impose another restriction on the number of read instructions that can be performed on a given location before a new value is written to that location. We say that a program has p -bounded reads (p -bdReads) if no location is read more than p times before it is overwritten. Formally:

Definition 4.1 (p -Bounded Reads (p -bdReads)). *A program execution $P^D(x)$ has p -Bounded Reads (p -bdReads) if the following holds for each location $i \in [n]$. For any set of $p+1$ CPU instructions $j_1 < \dots < j_{p+1}$, let the inputs/outputs of the j -th CPU step be $C_{\text{CPU}}^P(\text{state}_j, b_j^{\text{read}}) = (\text{state}_{j+1}, i_j^{\text{read}}, i_j^{\text{write}}, b_j^{\text{write}})$. If $i_{j_\ell}^{\text{read}} = i$ for all $\ell = 1, \dots, p+1$ then there must exist a CPU instruction j^* (not necessarily equal to any j_ℓ) with $j_1 \leq j^* < j_{p+1}$ such that $i_{j^*}^{\text{write}} = i$. Similarly, we define a p -bdReads property for a sequence of program executions $(P_1(x_1), \dots, P_\ell(x_\ell))^D$ if the above property holds across all the conjoined CPU steps in the sequence.*

In order to enable writing, we again rely on the ptWrites restriction on RAM. Whenever we want to write a bit b to a location i at time j , we simply write F evaluated on (j, i, b) under p different keys to \tilde{D} . Whenever we want to read from some location i , we discover the time u in which it was last written to, and proceed as before.

4.1 Overview of Construction

We describe at a high level how to construct a garbled RAM scheme for programs with ptWrites and p -bdReads that satisfies UMA security. As in the read-only warmup construction, in order to break the circularity, we must weaken the PRF key so the unopened entry in the translation table computationally hides the information of the unopened wire label. To do this, instead of using any arbitrary PRF, we use the above revocable GGM-PRF that allows us to adaptively constrain the PRF.

The main idea is that we maintain a “key schedule” K that is fed as garbled input to each garbled augmented CPU step, which then processes it and outputs an updated garbled K' . This key schedule is an array of p (possibly with different revocation sets) PRF keys. Each entry of the key schedule has a revocation set X associated with it that is known in the clear (we omit writing this X every time for each key, and implicitly allow a key k_X to reveal X).

Whenever the j -th CPU step wants to issue a read from a location i which was last written to in step u and this is the α -th read to that location since u , then it picks the key $K[\alpha]_X$, since we know $(u, i, 0), (u, i, 1) \notin X$. This is the key by which it will output (among other things) `translate`, and it updates K to K' by updating $K[\alpha]_X$ to $K[\alpha]_{X'}$ by puncturing $(u, i, 0)$ and $(u, i, 1)$ from the key (so $X' = X \cup \{(u, i, 0), (u, i, 1)\}$). This is where we avoid the circularity: the translation table can only be decoded by this key, but this key will have the critical values punctured, so all future keys in the key schedule still computationally hide the table. Since there are p values in K , a location can be read up to p times before it is overwritten.

Whenever the j -th CPU step wants to write a bit b to a location i , it computes $F_k(j, i, b)$ for every $k = K[1], \dots, K[p]$. Since j is brand new, these values could not have possibly been revoked.

4.2 Detailed Construction

We now give the detailed construction of our GRAM scheme under any one-way function for programs with ptWrites and p -bdReads satisfying UMA (unprotected memory access) security for multi-program execution with persistent memory. Let F be the revocable GGM-PRF described above, which can be based on any one-way function, and let \mathcal{K} be the keyspace for F . Let $(\text{GCircuit}, \text{CircEval})$ be a projective circuit garbling scheme with wire labels. The master key will be the fresh key schedule $K[1] \leftarrow \mathcal{K}, \dots, K[p] \leftarrow \mathcal{K}$. We also keep an array of sets corresponding to revoked values $X[1], \dots, X[p]$ (implicitly), and we write $K[\ell]_X$

whenever we want to alert the reader that this is not necessarily a fresh PRF key. One minor drawback is that these sets X need to be carried across program boundaries in the case of multiple program executions in order to ensure we do not run into circularity.

Garbled Data. The garbled data \tilde{D} is a $n \times p$ matrix of “secret keys”, which are just outputs of F as we will describe. We have a fresh key schedule, K , an array of p PRF keys as our master secret.

- $\tilde{D} \leftarrow \text{GData}(D, K)$: For each $i \in [n]$, $j \in [p]$ set $\tilde{D}[i, j] := \text{sk}$ where $\text{sk} = F_{K[j]}(0, i, D[i])$.

Garbled Program. We describe the augmented-CPU-step circuit $C_{\text{CPU}^+}^P$ for the program P in Figure 2. The garbled RAM program for P will consist of t copies of a garbled augmented-CPU-Step circuit

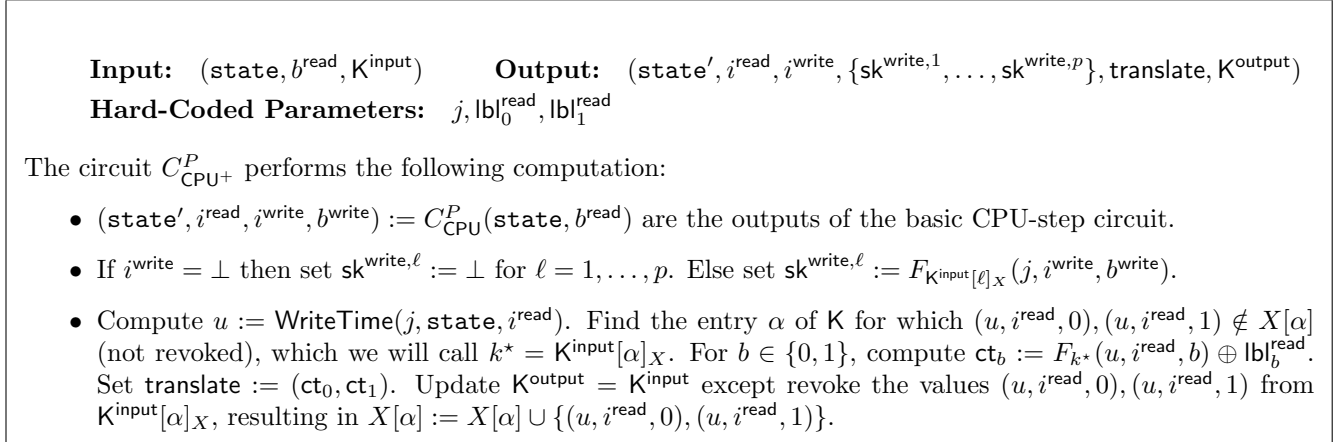


Figure 2: The Augmented CPU-Step Circuit

$\tilde{C}_{\text{CPU}^+}^P(j)$. As before, the labels for the output wires in each circuit are chosen carefully so that some wire values are revealed in the clear while others remain garbled for the next circuit.

- $(\tilde{P}, k^{\text{input}}) \leftarrow \text{GProg}(P, K, n, t^{\text{init}}, t^{\text{cur}})$: Let $t^{\text{max}} := t^{\text{init}} + t^{\text{cur}}$. We let j count down from $j = t^{\text{max}}$ to $t^{\text{init}} + 1$. For each j , we garble $C_{\text{CPU}^+}^P$ by calling $\tilde{C}_{\text{CPU}^+}^P(j) \leftarrow \text{GCircuit}(1^\kappa, C_{\text{CPU}^+}^P, \overline{\text{lbl}})$ where the output labels $\overline{\text{lbl}}$ are chosen as follows:

- The outputs $i^{\text{read}}, i^{\text{write}}, \text{sk}^{\text{write}}, \text{translate}$ are given out in the clear. For the last circuit $j = t^{\text{max}}$, we do not provide these outputs, but instead provide the output state' in the clear and this serves as the output of the computation. This completely fixes all of the output-wire labels for that circuit.
- For $j \neq t^{\text{max}}$, the labels of the output wires corresponding to state' are set to match the labels of the input wires corresponding to state in circuit $\tilde{C}_{\text{CPU}^+}^P(j+1)$.
- For the initial circuit $j = t^{\text{init}} + 1$, we also hard-code the input bit b^{read} to 0, and $K^{\text{input}} = K$ but with the appropriate “carried over” X values revoked.

For $j \neq t^{\text{max}}$, let $\text{lbl}_0^{(\text{read}, j+1)}, \text{lbl}_1^{(\text{read}, j+1)}$ be the labels of input wire for the bit b^{read} in the $(j+1)$ st garbled circuit. The j th garbled circuit contains the hard-coded secret values:

$$(j, \text{lbl}_0^{(\text{read}, j+1)}, \text{lbl}_1^{(\text{read}, j+1)}).$$

Set $\tilde{P} := [\tilde{C}_{\text{CPU}^+}^P(t^{\text{init}} + 1), \dots, \tilde{C}_{\text{CPU}^+}^P(t^{\text{max}})]$, and set $k^{\text{input}} := \{(i, b, \text{lbl}_b^{\text{input}, i}) : i \in [v], b \in \{0, 1\}\}$ to consist of all of the input-wire labels for the v input wires corresponding to the input state in the initial circuit $\tilde{C}_{\text{CPU}^+}^P(t^{\text{init}} + 1)$.

Garbled Input. Finally, the garbled input \tilde{x} is created the same way as in garbled circuits. It simply consists of the subset of labels of $k^{\text{input}} := \{(i, b, \text{lbl}_b^{\text{input}, i}) : i \in [v], b \in \{0, 1\}\}$ corresponding to the bits of x .

- $\tilde{x} \leftarrow \text{GInput}(x, K, k^{\text{input}})$: Parse $k^{\text{input}} := \{(i, b, \text{lbl}_b^{\text{input}, i}) : i \in [v], b \in \{0, 1\}\}$ and output $\tilde{x} = (\text{lbl}_{x[1]}^{\text{input}, 1}, \dots, \text{lbl}_{x[v]}^{\text{input}, v})$ where $x[i]$ denotes the i th bit of x and $v := |x|$.

Evaluation. To run $y = \text{GEval}^{\tilde{D}}(\tilde{P}, \tilde{x})$: parse $\tilde{P} = [\tilde{C}_{\text{CPU}^+}^P(1), \dots, \tilde{C}_{\text{CPU}^+}^P(t)]$ as consisting of t garbled circuits. The evaluator evaluates the circuits one-by-one. Set $\tilde{x}_1 := \tilde{x}$. For $j = 1, \dots, t$:

- When $j = 1$, run $\text{CircEval}(\tilde{C}_{\text{CPU}^+}^P(1), \tilde{x}_1)$ else run $\text{CircEval}(\tilde{C}_{\text{CPU}^+}^P(j), (\tilde{x}_j, \tilde{K}_j, \text{lbl}^{\text{read}, j}))$ where \tilde{x}_j consists of labels for the garbled state_j , \tilde{K}_j consists of labels for the garbled K_j , and $\text{lbl}^{\text{read}, j}$ is a label for the read-bit in the j th CPU circuit. This reveals the outputs $i_j^{\text{read}}, i_j^{\text{write}}, \{\text{sk}_j^{\text{write}, 1}, \dots, \text{sk}_j^{\text{write}, p}\}$, $\text{translate}_j = (\text{ct}_0^{(j)}, \text{ct}_1^{(j)})$ in the clear. For $j < t$ it also reveals the garbled output \tilde{x}_{j+1} and \tilde{K}_{j+1} corresponding to the labels of state_{j+1} for circuit $j + 1$. For $j = t$ it reveals the output of the computation $y = \text{state}_{t+1}$ in the clear.
- For $\ell = 1, \dots, p$, look up $\text{skread}, \ell_j = \tilde{D}[i_j^{\text{read}}, \ell]$. WLOG we can efficiently tell when we have a valid wire label, so we try all $2p$ possibilities: there is some unique $\ell^* \in [p]$ and $b^* \in \{0, 1\}$, such that $\text{sk}_j^{\text{read}, \ell^*} \oplus (\text{ct}_{b^*}^{(j)})$ is a valid label, which we will call $\text{lbl}^{\text{read}, j+1}$.
- If $i_j^{\text{write}} \neq \perp$, update the row $\tilde{D}[i_j^{\text{write}}] := \{\text{sk}_j^{\text{write}, 1}, \dots, \text{sk}_j^{\text{write}, p}\}$.

4.3 Proof of Security

We now state a technical theorem.

Theorem 4.2. *Given any OWF and a secure projective circuit garbling scheme (which can be built from the OWF), the above construction is a UMA (unprotected memory access) secure garbled RAM scheme for all program executions with pt Writes and p -bdReads. Furthermore, it is also secure in the setting of multi-program execution with persistent memory. If p is $O(\text{polylog}(n))$, the garbled memory size is $\tilde{O}(n)$, the garbled program size is $\tilde{O}(t^2)$, and running time of evaluating the garbled program is $\tilde{O}(t^2)$ where \tilde{O} denotes big-O up to $\text{poly}(\kappa)\text{polylog}(n)$ terms.*

The full proof appears in Appendix A.

5 Recursive Construction

In the previous section, we showed how to achieve a garbled data size of $\tilde{O}(n)$, a garbled program size of $\tilde{O}(t^2)$ (again observing that there are t CPU steps, and K can grow to size $\tilde{O}(t)$) and online running time of $\tilde{O}(t^2)$ where \tilde{O} denotes big-O up to $\text{poly}(\kappa)\text{polylog}(n)$ terms. In this section, we describe a way to use a recursive construction that gives an garbled program size of size of $\tilde{O}(t \cdot n^c)$ and running time of $\tilde{O}(t \cdot n^c)$ where c is a constant depending on the security parameter as we describe below.

We consider a balancing technique to further lower the overhead: observe that we can re-encode everything under new PRF keys more frequently if we do not want to wait t steps and let the keys K grow to size $t \log n$. Suppose, for example, after \sqrt{n} steps we overwrite the entire contents of memory with a *single* big “refresh” garbled circuit with fresh PRF keys. This ensures that each CPU step is no larger than $\tilde{O}(\sqrt{n})$, but now we incur an extra $\tilde{O}(n)$ -sized circuit every \sqrt{n} steps, a total of t/\sqrt{n} such refreshes across the full execution. Thus, the total size of the garbled program we calculate as follows: there are t CPU steps, each of size at most $\tilde{O}(\sqrt{n})$, and we incur also an additional $\tilde{O}(t/\sqrt{n})$ refresh steps, each of size $\tilde{O}(n)$, giving a total of $\tilde{O}(t\sqrt{n})$.

Pretend we have a CPU of size $f(n)$, then we only need a refresh of size n every $f(n)$ steps, but we also need to simulate this large CPU with a smaller CPU. In order to do so, we must emulate the CPU as a RAM program. We can simply convert the circuit into a RAM program generically, and then apply ORAM to guarantee the properties we need: `ptWrites` and `polylog(n)-bdReads`, which incurs a poly-log overhead. We can then simulate this CPU using our GRAM construction. This simulation loses a factor of $\text{poly}(\kappa)\text{polylog}(n)$ so it needs to be carefully balanced against the CPU size f . If $C(n)$ denotes the *average* cost per CPU step, then we have a recurrence relation

$$C(n) = \frac{n}{f(n)} + \text{poly}(\kappa)\text{polylog}(n) \cdot C(f(n)).$$

The left summand is for the the average running time for the refresh steps, and the right summand is average running time of emulating an ordinary CPU step with a smaller GRAM of size $f(n)$. In the case where κ is $\text{polylog}(n)$ (for exponential size RAM), this can be solved can be solved by setting $C(n) = 2^{\sqrt{2 \log n \log \text{poly}(\kappa)\text{polylog}(n)}}$ and f appropriately. Thus the amortized overhead is less than n^c for any $c > 0$. On the other hand, if κ is polynomially related to n , then suppose $\text{poly}(\kappa)\text{polylog}(n)$ is $O(n^d)$, then we solve the recurrence by setting f to be $2n^a$ and $C(n)$ to be n^ϵ as follows:

$$n^\epsilon = \frac{n}{2n^a} + n^d \cdot (2n)^{\epsilon a}$$

which we balance as

$$\begin{aligned} n^\epsilon/2 &= \frac{n}{2n^a} \\ \text{and } n^\epsilon/2 &= n^d \cdot (2n)^{\epsilon a} \end{aligned}$$

Which is then solved as $a = 1 - \epsilon$ and $d = \epsilon^2 + \frac{\epsilon^2 - \epsilon - 1}{\log n}$. By appropriately choosing the security parameter, ϵ can be made arbitrarily small. We therefore obtain the following theorem.

Theorem 5.1. *Given any OWF and a secure projective circuit garbling scheme (which can be built from the OWF), the above construction is a UMA (unprotected memory access) secure garbled RAM scheme for all program executions with `ptWrites` and `p-bdReads`. Furthermore, it is also secure in the setting of multi-program execution with persistent memory. If p is $O(\text{polylog}(n))$, the garbled data size is $\tilde{O}(n)$, the garbled program size is $\tilde{O}(t \cdot n^\epsilon)$, and running time of evaluating the garbled program is $\tilde{O}(t \cdot n^\epsilon)$ where \tilde{O} denotes big-O up to $\text{poly}(\kappa)\text{polylog}(n)$ terms.*

Combining this with the following lemma, we obtain our main theorem.

Lemma 5.2. *Any garbled RAM scheme G that provides UMA security and supports programs with `ptWrites` and `p-bdReads` (where p is $\text{polylog}(n)$) can be extended to a scheme G' with full security supporting arbitrary programs with at most $\text{poly}(\kappa)\text{polylog}(n)$ overhead.*

The proof of this lemma is of the same spirit of the security upgrade found in Part I [6]. For completeness, we include the proof in Appendix B.

We then obtain our main theorem, where we balance the running time with n^ϵ .

Theorem 5.3 (Main Theorem). *Given any OWF and a secure projective circuit garbling scheme (which can be built from the OWF), there exists a garbled RAM scheme that is fully secure supporting arbitrary programs, and also secure in the setting of multi-program execution with persistent memory. The resulting garbled memory size is $\tilde{O}(n)$, the garbled CPU steps size is $\tilde{O}(t \cdot \min(t, n^\epsilon))$, and running time of evaluating the garbled program is $\tilde{O}(t \cdot \min(t, n^\epsilon))$ where \tilde{O} denotes big-O up to $\text{poly}(\kappa)\text{polylog}(n)$ terms.*

6 Conclusions and Open Problems

In this paper, we showed how to construct Garbled RAM programs from one-way functions where, ignoring $\text{poly}(\kappa)\text{polylog}(n)$ factors, the garbled memory size stays the same, and the garbled CPU size and running time grow by a factor of $\min(t, n^\epsilon)$. It remains open whether or not $\text{poly}(\kappa)\text{polylog}(n)$ can be achieved just under one-way functions.

In our construction, we note that although the program code can be compactly garbled, the CPU steps that are needed to run the program grow with the number of steps. The recent work of Goldwasser et al. [10] have shown how to compactly transmit a Turing Machine that can be ran on encrypted data, though under stronger assumptions. We leave as an open problem the question of constructing a compact Garbled RAM where the size does not grow in the number of steps executed under just one-way functions.

References

- [1] 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.
- [2] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT*, volume 7658 of *Lecture Notes in Computer Science*, pages 134–153. Springer, 2012.
- [3] 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.
- [4] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT (2)*, volume 8270 of *Lecture Notes in Computer Science*, pages 280–300. Springer, 2013.
- [5] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. *IACR Cryptology ePrint Archive*, 2013:401, 2013.
- [6] Craig Gentry, Shai Halevi, Mariana Raykova, and Daniel Wichs. Garbled ram revisited (part i). volume 2014. To appear in EUROCRYPT 2014.
- [7] Oded Goldreich. Towards a theory of software protection and simulation by oblivious rams. In Alfred V. Aho, editor, *STOC*, pages 182–194. ACM, 1987.
- [8] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions (extended abstract). In *FOCS*, pages 464–479. IEEE Computer Society, 1984.
- [9] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [10] 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, *CRYPTO (2)*, volume 8043 of *Lecture Notes in Computer Science*, pages 536–553. Springer, 2013.
- [11] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *STOC*, pages 555–564. ACM, 2013.

- [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.
- [13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM Conference on Computer and Communications Security*, pages 669–684. ACM, 2013.
- [14] Yehuda Lindell and Benny Pinkas. A proof of security of yao’s protocol for two-party computation. *J. Cryptology*, 22(2):161–188, 2009.
- [15] Steve Lu and Rafail Ostrovsky. Distributed oblivious RAM for secure two-party computation. In *TCC*, pages 377–396, 2013.
- [16] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 719–734. Springer, 2013.
- [17] Moni Naor and Kobbi Nissim. Communication preserving protocols for secure function evaluation. In Jeffrey Scott Vitter, Paul G. Spirakis, and Mihalis Yannakakis, editors, *STOC*, pages 590–599. ACM, 2001.
- [18] Rafail Ostrovsky. Efficient computation on oblivious rams. In Harriet Ortiz, editor, *STOC*, pages 514–523. ACM, 1990.
- [19] Rafail Ostrovsky. *Software Protection and Simulation On Oblivious RAMs*. PhD thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 1992.
- [20] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In Frank Thomson Leighton and Peter W. Shor, editors, *STOC*, pages 294–303. ACM, 1997.
- [21] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. *IACR Cryptology ePrint Archive*, 2013:454, 2013.
- [22] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *FOCS*, pages 160–164, 1982.

A Proof of Theorem 5.3

Correctness follows from the definition. We calculate the efficiency as follows: \tilde{D} stores $n \cdot p$ “keys” which are PRF outputs. Each CPU step must be large enough to keep the key schedule which has p PRF keys, and since each time a key gets punctured it grows by $\log n$, after t steps the key schedule may take up to $p \cdot t$ keys. Then during evaluation, the running time is as large as the space since the entire circuit is evaluated. If p is $\text{polylog}(n)$, then our claim follows. Now we prove security.

Let *CircSim* be the simulator of the circuit garbling scheme. We describe a simulator *Sim* for our garbled program:

Input: Under UMA-security, *Sim* gets the inputs $\{P_a, t_a, y_a\}_{a=1}^\ell, D, \text{MemAccess} = \{(i_j^{\text{read}}, i_j^{\text{write}}, b_j^{\text{write}})\}_{j=1}^{t^{\max}}$, where program P_a executes t_a CPU steps and output y_a , the initial memory contents are $D \in \{0, 1\}^n$ and *MemAccess* describes the entire memory access throughout all $t^{\max} = \sum_{a=1}^\ell t_a$ CPU steps executed.

Output: The simulator *Sim* outputs: $\tilde{D}, \tilde{P}_1, \dots, \tilde{P}_\ell, \tilde{x}_1, \dots, \tilde{x}_\ell, \tilde{K}_1, \dots, \tilde{K}_\ell$.

Initialization: The simulator *Sim* samples the master key schedule $K[1] \leftarrow \mathcal{K}, \dots, K[p] \leftarrow \mathcal{K}$.

Garbled Data: It creates \tilde{D} using the honest process.

Garbled Programs/Inputs: The simulator Sim processes each program $a \in [\ell]$ separately. Let $j_a^{init} = \sum_{b=1}^{a-1} t_a + 1$ be the first CPU step executed by the program P_a and let $j_a^{max} = j_a^{init} + t_a - 1$ be the last step. The simulator starts from $j = j_a^{max}$ and counts down to $j = j_a^{init}$. For each j :

- For the last circuit $j = j_a^{max}$, create $\tilde{C}_{\text{CPU}^+}^P(j)$ by calling CircSim on the circuit $C_{\text{CPU}^+}^P$ with the output-labels of \mathbf{state}_{j+1} set to the value y_a in the clear. This produces some set of input labels for the input \mathbf{state}_j , they key schedule \mathbf{K}_j , and the bit b_j^{read} .
- For any other $j \neq j_{max}$, create $\tilde{C}_{\text{CPU}^+}^P(j)$ by calling CircSim on the circuit $C_{\text{CPU}^+}^P$ where the output-labels for $i_j^{\text{read}}, i_j^{\text{write}}, \{\mathbf{sk}_j^{\text{write},1}, \dots, \mathbf{sk}_j^{\text{write},p}\}, \text{translate}$ are given “in the clear” and the output-labels of of the updated \mathbf{state}_{j+1} are set to match the input labels for \mathbf{state}_{j+1} given by the circuit-simulator for the circuit $j + 1$. The actual values $\{\mathbf{sk}_j^{\text{write},1}, \dots, \mathbf{sk}_j^{\text{write},p}\}, \text{translate}$ are computed via:
 - If $i_j^{\text{write}} = \perp$ then set $\mathbf{sk}_j^{\text{write}} := \perp$. Otherwise, set $\mathbf{sk}_j^{\text{write},\ell} = F_{\mathbf{K}[\ell]}(j, i_j^{\text{write}}, b_j^{\text{write}})$.
 - Let $u < j$ be the last write-time to location i_j^{read} and let $b = b_u^{\text{write}}$ be the bit written to the location at time u . Let α be the α -th time this location has been read since u . The values u, b, α can be easily computed given MemAccess . Set $k^* = \mathbf{K}[\alpha]$ (which corresponds to the entry in the key schedule from which the CPU will read), and set:

$$\mathbf{ct}_b := F_{k^*}(u, i_j^{\text{read}}, b) \oplus \mathbf{lbl}^{\text{read},j+1}, \quad \mathbf{ct}_{1-b} := F_{k^*}(u, i_j^{\text{read}}, 1-b) \oplus 0^{v_{\text{lbl}}}$$

where $\mathbf{lbl}^{\text{read},j+1}$ is the label of the “read-bit” wire given by the circuit-simulator for the circuit $j + 1$ and v_{lbl} is the label-length. Set $\text{translate} := (\mathbf{ct}_0, \mathbf{ct}_1)$.

Set \tilde{x} to be the input labels created by the circuit-simulator for the input \mathbf{state} of the initial circuit, and similarly for \tilde{K} .

Our proof follows the overall structure of [6] where we argue indistinguishability of the real output and the output of Sim by using a series of hybrid distributions \mathbf{Hyb}_j for $j = 1, \dots, t^{max}$. In the hybrid j , garbled circuits $1, \dots, j$ are created as in the simulation and garbled circuits $j + 1, \dots, t^{max}$ are created as in the real distribution.

We also define a hybrid distribution \mathbf{Hyb}'_j which is like \mathbf{Hyb}_j except for the simulation of the j th CPU-step circuit (except for $j = t^{max}$ for which we set $\mathbf{Hyb}'_j = \mathbf{Hyb}_j$). Instead of choosing translate as in the simulation described above, we choose $\text{translate} = (\mathbf{ct}_0, \mathbf{ct}_1)$ to both be encryptions of the correct label of the next circuit:

$$\mathbf{ct}_0 := F_{k^*}(u, i_j^{\text{read}}, 0) \oplus \mathbf{lbl}_0^{\text{read},j+1}, \quad \mathbf{ct}_1 := F_{k^*}(u, i_j^{\text{read}}, 1) \oplus \mathbf{lbl}_1^{\text{read},j+1}$$

where $\mathbf{lbl}_0^{\text{read},j+1}, \mathbf{lbl}_1^{\text{read},j+1}$ are the labels corresponding to the bits 0,1 for the wire b^{read} in circuit $j + 1$ which is still created using the real garbling procedure.

Notice that \mathbf{Hyb}_0 is equal to the real distribution and $\mathbf{Hyb}_{t^{max}}$ is equal to the simulated distribution. Therefore, we prove the theorem by showing that for each j , we have:

$$\mathbf{Hyb}_j \stackrel{\text{comp}}{\approx} \mathbf{Hyb}'_{j+1} \stackrel{\text{comp}}{\approx} \mathbf{Hyb}_{j+1}$$

We prove this in the following two claims.

Claim A.1. For each $j \in \{0, \dots, t^{max}\}$ we have $\mathbf{Hyb}_j \stackrel{\text{comp}}{\approx} \mathbf{Hyb}'_{j+1}$.

Proof. This follows directly from the security of the circuit-garbling scheme applied only to the garbled augmented-CPU circuit $j + 1$. \square

Claim A.2. For each $j \in \{1, \dots, t^{max}\}$ we have $\mathbf{Hyb}'_j \stackrel{\text{comp}}{\approx} \mathbf{Hyb}_j$.

Proof. This follows from the security of our revocable GGM-PRF scheme. The only difference between \mathbf{Hyb}'_j and \mathbf{Hyb}_j is the value of $\text{translate}_j = (\text{ct}_0, \text{ct}_1)$ used to simulate the j th garbled circuit. Let $b = b_{j+1}^{\text{read}}$ be the *value* of the read-bit in location i_j^{read} in the real computation. Then, in \mathbf{Hyb}'_j we set

$$\text{ct}_b := F_{k^*}(u, i_j^{\text{read}}, b) \oplus \text{lbl}_b^{\text{read}, j+1} \quad , \quad \text{ct}_{1-b} := F_{k^*}(u, i_j^{\text{read}}, 1-b) \oplus \text{lbl}_{1-b}^{\text{read}, j+1}$$

whereas in \mathbf{Hyb}_j we set

$$\text{ct}_b := F_{k^*}(u, i_j^{\text{read}}, b) \oplus \text{lbl}_b^{\text{read}, j+1} \quad , \quad \text{ct}_{1-b} := F_{k^*}(u, i_j^{\text{read}}, 1-b) \oplus 0^{\text{lbl}}$$

where $u < j$. Observe the following, where we consider the key schedule and how it interacts with \tilde{D} :

- When b was written to \tilde{D} in CPU step u , the PRF was evaluated on $(u, i_j^{\text{read}}, b)$ under $K[1], \dots, K[p]$, one of which is k^* .
- For the real CPU steps after j , $(u, i_j^{\text{read}}, b)$ has been revoked from k^* .
- If this location is written to again, it will be under some different time $u' \neq u$, where it can then be further read again.

Suppose on the contrary that some adversary \mathcal{A} is able to distinguish the two hybrids. Then we devise an adversary \mathcal{A}' that will break the pseudorandomness of F as follows. \mathcal{A}' produces an input and memory that it wishes to be challenged upon. Using MemAccess , \mathcal{A}' will determine exactly which k^* will be used and what values have already been revoked from K . It asks for these partially-revoked keys from the PRF challenger, and therefore can construct all subsequent real steps after j . In order to obtain the sk output by some earlier CPU step in order to simulate it, \mathcal{A}' asks the PRF challenger to evaluate F at those points. Note by the definition of k^* it will never be the case that \mathcal{A}' asks the PRF oracle for $(u, i_j^{\text{read}}, 0)$ or $(u, i_j^{\text{read}}, 1)$. Then in order to generate translate , \mathcal{A}' asks the oracle to evaluate $F_{k^*}(u, i_j^{\text{read}}, b)$ on the correct read bit b , and asks to be challenged on $(u, i_j^{\text{read}}, 1-b)$. \mathcal{A}' gets from the PRF challenger a value R which is either real or random, and \mathcal{A}' completes the translation table by picking a challenge bit c and if $c = 0$ sets the ciphertext to be $R \oplus \text{lbl}_{1-b}^{\text{read}, j+1}$ and R otherwise. Clearly, the advantage of \mathcal{A} is then inherited by \mathcal{A}' since these two distributions perfectly indistinguishable if R is random, and identical to the hybrids if R is real. Since the PRF key k^* is punctured after step j , indistinguishability follows from the pseudorandomness of F . \square

B Upgrading to Full Security/Functionality

This section is largely a restatement from Part I [6] with the key difference of bdReads .

We now describe a general transformation from any garbled RAM scheme that only provides UMA security and only supports program executions with ptWrites and bdReads into a fully secure garbled RAM scheme for arbitrary programs. This transformation uses oblivious RAM (ORAM) to first compile the original program P into a new program P^* that stores/accesses its memory using ORAM. This ensures that the memory contents and access pattern of the compiled program do not reveal anything about those of the original program. For simplicity, we assume the ORAM scheme already ensures that the compiled program satisfies the ptWrites and $p\text{-bdReads}$ property. Indeed, many ORAM scheme do satisfy these with $p = \text{polylog}(n)$, e.g. Ostrovsky's hierarchical scheme [18]. Now we can simply apply our original UMA-secure garbled RAM scheme for ptWrites on this compiled program to get a fully secure solution.

The Compiler. Let $G = (\text{GData}, \text{GProg}, \text{GInput}, \text{GEval})$ be any garbled RAM scheme that provides UMA security and only supports program executions with ptWrites and $p\text{-bdReads}$. Let $O = (\text{OData}, \text{OProg})$ be any ORAM that guarantees ptWrites and $p\text{-bdReads}$. We define a garbled RAM scheme G' which first applies the ORAM O to the program to make it oblivious and satisfy ptWrites and $p\text{-bdReads}$, and then uses G to garble it. In detail, we define $G' = (\text{GData}', \text{GProg}', \text{GInput}', \text{GEval})$ as follows:

- $\text{GData}'(D, k = (k_1, k_2))$: Call $D^* \leftarrow \text{OData}(D, k_1), \tilde{D}^* \leftarrow \text{GData}(D^*, k_2)$. Output \tilde{D}^* .
- $\text{GProg}'(P, k = (k_1, k_2), n, t_{\text{init}}, t_{\text{cur}})$: Call $P^* \leftarrow \text{OProg}(P)$ and $(\tilde{P}^*, k_2^{\text{input}}) \leftarrow \text{GProg}(P^*, k_2, n, t_{\text{init}}^*, t_{\text{cur}}^*)$ where $t_{\text{init}}^*, t_{\text{cur}}^*$ are the updated times with the overhead of the ORAM scheme. Output $\tilde{P}^*, k^{\text{input}} = (k_1, k_2^{\text{input}})$.
- $\text{GInput}'(x, k^{\text{input}} = (k_1, k_2^{\text{input}}))$: Output $\tilde{x}^* \leftarrow \text{GInput}((x, k_1), k_2^{\text{input}})$.

Theorem B.1. *If G is a garbled RAM scheme that provides UMA security and supports programs with ptWrites and $p\text{-bdReads}$ and O is an ORAM with ptWrites then G' is a garbled RAM with full security and supporting arbitrary programs.*

Proof. It is clear that the use of ORAM with ptWrites and $p\text{-bdReads}$ in the above construction ensures that G is only used on program executions that satisfy ptWrites . Therefore, we only need to prove that G' provides security. Let Sim_1 be the ORAM simulator and let Sim_2 be the simulator for G . Then we define the simulator Sim' for G' which first calls $(D_{\text{sim}}^*, \text{MemAccess}_{\text{sim}}) \leftarrow \text{Sim}_1(1^\kappa, n, \{t_i\}_{i=1}^\ell)$ to compute the simulate data and access pattern, and then outputs $\text{Sim}_2(1^\kappa, \{P_i^*, t_i^*, y_i\}_{i=1}^\ell, D_{\text{sim}}^*, \text{MemAccess}_{\text{sim}})$ where t_i^* are the updated running times after applying ORAM. Security follows from two simple hybrid arguments:

- By the security of G' , the real distribution is indistinguishable from

$$\text{Sim}_2(1^\kappa, \{P_i^*, t_i^*, y_i\}_{i=1}^\ell, D^*, \text{MemAccess})$$

where $D^*, \text{MemAccess}$ are the “real” data and access pattern produced by the oblivious RAM scheme.

- By the security of the ORAM scheme O , we know that $(D^*, \text{MemAccess})$ is indistinguishable from the simulated $(D_{\text{sim}}^*, \text{MemAccess}_{\text{sim}})$.

This completes the proof. □